Powered by

SCIENTIFIC COMPUTING WORLD

SCW 75

2026

Celebrating the thought leaders in scientific computing

In association with

aws

SEAGATE

Smarter Storage for Scientific Research & Life Sciences

Designed for the **realities of modern research**: expanding **AI workloads**, high-value **data**, and the need for trusted performance at **scale**.

At Seagate, we collaborate with **research institutions** and **life sciences organisations** to address **real-world data challenges**. Our Seagate Systems and Exos solutions support the **full lifecycle of scientific data** — from acquisition and analysis to long-term archiving.

What We Offer:

+ **Expand research data capacity**
  High-capacity storage solutions that live alongside your instruments, devices, and lab systems — wherever your research takes them.
+ **Orchestrate data flow at the edge**
  Modular storage solutions that pair with scientific workflows to move large data sets faster, more securely, and with greater control.
+ **Maintain control and flexibility**
  Vendor-agnostic platforms that help ensure effective data flow, management, and infrastructure flexibility — without locking your research into a single ecosystem.
+ **Overcome exponential data growth**
  Hard drive capacities up to 44TB, with a clear roadmap toward 100TB. That gives you the long-term scalability needed to keep pace with larger datasets, AI workloads, and tomorrow's discoveries.

Speak to our team to explore how **Seagate Systems and Exos** can **support your research goals.**

Welcome to the SCW75 – class of 2026

It is with enormous pride that we introduce the inaugural cohort of the Scientific Computing World 75 – 75 exceptional individuals who are redefining what is possible at the intersection of computing and research.

Launching a new recognition programme is no small undertaking, and the response has surpassed all our expectations.

From the first call for nominations, it became clear that the scientific computing community had been waiting for precisely this kind of dedicated spotlight – one focused not on institutional prestige, but on the people actively in the 'trenches', deploying transformative technology to push research forward.

The 2026 list honours leaders across three distinct disciplines: High Performance Computing, Computational Engineering & Simulation and Laboratory Informatics & Data Management.

What unites them is a shared conviction that the right computing infrastructure, thoughtfully implemented, can accelerate research timelines and turn ever-growing volumes of data into genuine discovery.

So how did we arrive at our final 75? We asked ourselves a consistent set of questions for every nomination received: Is this person driving or facilitating strategic computing decisions at a research level? Are there concrete projects, deployments, or outcomes that have advanced their organisation's research goals? And is scientific computing genuinely at the heart of that work?

Those criteria helped us identify individuals who are not simply adopting new tools, but championing them – making the case internally, navigating budget constraints, mentoring colleagues and delivering results that matter to science.

The SCW75 broadly recognises individuals across three categories:

High Performance Computing – directors, managers, and R&D IT leaders in universities, national laboratories and commercial research organisations who are championing HPC adoption and expansion, whether cloud-based, on-premise or hybrid.

Computational Engineering & Simulation – engineering directors, simulation team leads and research computing managers advancing research through simulation technologies, from CFD and structural analysis to climate modelling and molecular simulation.

Laboratory Informatics & Data Management – informatics directors, data science leads and research IT managers building the data infrastructure that powers modern research through LIMS, ELNs, data lakes and analytics platforms.

As this community takes shape, we are committed to offering its members meaningful opportunities to connect, collaborate and be heard. Being named to the SCW75 comes with a featured profile on Scientific Computing World's website, www.scientific-computing.com, recognition across our social media channels and partner networks, an official digital badge and exclusive invitations to expert panel discussions and virtual roundtables. As the programme matures, we expect those opportunities to expand further – into conference speaking slots, site visits and elite in-person networking.

We would also like to extend our sincere thanks to everyone who took the time to submit a nomination. The quality and breadth of entries made the selection process genuinely difficult.

Congratulations to the SCW75 class of 2026. This community is only just beginning.

Robert Roe,
Editor, *Scientific Computing World*

Content and delivery team
Editor: Robert Roe *robert.roe@europascience.com*
Head of content: Finbarr O'Reilly
Senior graphic designer: Zoe Wade
Audience development manager: Andrew Knight
Production manager: Nick Clark

Sales
Senior account manager:
Tim Richters *tim.richters@europascience.com*

Corporate
Chief operating officer: Mark Elliott
Chief executive officer: Warren Clark

Registered address: Europa Science Ltd, Europa Science | St John's Innovation Centre | Cowley Road | Cambridge | CB4 OWS | The United Kingdom. Tel: +44 (0)1223 221 030. Fax: +44 (0)1223 213 385. ©2025 Europa Science Ltd. ● Tel: +44 (0) 1223 211170 ● Fax: +44 (0) 1223 213385●

Subscriptions: Free registration is available to qualifying individuals. Register online at www.scientific-computing.com

Thank you to our partners

The first Scientific Computing World 75 would not have been possible without the support of its sponsors and project partners. *Scientific Computing World* is proud to work with these inspiring organisations and to help showcase what they do

ACD/Labs
8 King Street East, Suite 107, Toronto, Ontario, M5C 1B5 Canada
Tel: +1 416 368-3435
Contact: info@acdlabs.com
www.acdlabs.com

ACD/Labs is a leading provider of scientific software, aiding decision-making and digital transformation in chemical and pharmaceutical R&D. Its science-centric solutions offer harmonised, AI-ready analytical, structural and molecular information. ACD/Labs' technologies enable proliferation of chemically-intelligent knowledge and workflow automation. As part of Revvity Signals, ACD/Labs will continue to partner with global organisations and accelerate scientific innovation.

AWS
410 Terry Ave. N, Seattle,
WA 98109-5210, United States
Contact: www.x.com/AWSNewsroom
aws.amazon.com

Amazon Web Services (AWS) is guided by customer obsession, pace of innovation, commitment to operational excellence and long-term thinking. By democratising technology for nearly two decades and making cloud computing and generative AI accessible to organisations of every size and industry, AWS has built one of the fastest-growing enterprise technology businesses in history. Millions of customers trust AWS to accelerate innovation, transform their businesses and shape the future.

COOLIT
CoolIT Systems, #200, 115 Quarry Park Road SE, Calgary, AB, T2C 5G9, Canada
Tel: +1 403-770-5766
Contact: info@coolitsystems.com
www.coolitsystems.com

CoolIT Systems specialises in scalable liquid cooling solutions for the world's most demanding high-performance computing environments. Through its broad portfolio of modular direct liquid cooling (DLC) products, CoolIT enables dramatic increases in rack density, component performance and power efficiency across any high-density computing scenario. CoolIT collaborates with the top semiconductor manufacturers and cloud service providers to develop cooling technologies that enable ever-greater computing performance and efficiency. Through its Liquid Lab™ R&D centres in Calgary and Taipei, global manufacturing operations in Canada, China and Vietnam, and on-site support in more than 80 countries, CoolIT serves customers at scale worldwide. Together, CoolIT and its partners are leading the way for widespread adoption of efficient high-density, high-performance computing.

REVVITY SIGNALS
77 4th Avenue, Waltham, Massachusetts 02451, USA
Tel: (800) 762-4000 or (+1) 203-925-4602
www.revvitysignals.com

Revvity Signals provides science-based software solutions that leverage big tech innovations to help improve lives everywhere. Its scientific SaaS platform rapidly delivers new capabilities, embraces the latest innovations to simplify science and promotes multidisciplinary collaboration to accelerate innovation. Its powerful informatics solutions are used across a spectrum of industries including pharma and biotech, specialty and agro-chemicals, energy and petrochemicals, flavours and fragrances, food and beverage, and electronics. With a 97% renewal rate, Revvity Signals is the SaaS business-to-business platform, trusted by the leading pharma and biotech companies worldwide: 100% of the top 20 biopharma corporations choose Revvity Signals.

Sapio Sciences
205 N George St, York, PA, 17401, USA
Contact: www.sapiosciences.com/contact-sapio/

Sapio Sciences delivers an AI-native laboratory informatics platform that unifies ELN and LIMS into a single, configurable environment. Sapio ELN is the AI-native electronic lab notebook, where scientists design, execute and analyse experiments using natural language while maintaining scientific rigour and compliance. Sapio ELN is enhanced by its ecosystem of trusted AI infrastructure, molecular modelling, scientific knowledge and biological insight partners, bringing in silico reasoning directly into the experimental record. Sapio LIMS provides robust workflow automation, sample and materials management, and operational control across the lab. Across both systems, Sapio Elain functions as an AI co-scientist, supporting reasoning, analysis and next-step decisions while keeping scientists in control.

Seagate Technology
Tupolevlaan 105, 1119 PA Schiphol-Rijk,
The Netherlands
Tel: +31 (0) 858 33587
Contact: www.seagate.com

As data becomes the foundation of modern scientific discovery, research organisations are under increasing pressure to manage, access and preserve ever-growing volumes of data. Seagate Technology plays a critical role in enabling this transformation, delivering the infrastructure needed to support data-intensive research at scale. From genomics and high-throughput sequencing to computational modelling and AI-driven analysis, Seagate helps institutions overcome the challenges of managing petabyte-scale datasets, ensuring data remains accessible, reliable and ready to drive insight. With a focus on long-term data integrity, efficiency and sustainability, Seagate supports researchers in accelerating innovation while building the data foundations needed for the next generation of breakthroughs.

The architects of scientific computing's new era

Budgets are rising and ambitions are high – but as some members of the inaugural SCW75 make clear, turning infrastructure investment into research impact is harder than it looks, writes **Robert Roe**

Scientific computing has moved out of the specialist computing centre and into the strategic planning of laboratories, engineering teams and research-led businesses. Behind that shift are people: infrastructure leaders, informatics directors, simulation specialists and research computing managers who are translating rapidly evolving technology into reliable, usable, trusted outcomes. The inaugural SCW75 brings together 75 of the most influential of them, and their collective experience reveals as much about the challenges ahead as it does about the progress already made.

Research communities that once appeared distinct are increasingly converging around similar infrastructure needs. A bioinformatician, automotive engineer, computational chemist and semiconductor researcher may use different tools and datasets, but they all depend on scalable compute, high-performance storage, workflow orchestration, robust data management and software that can bridge specialist domains.

Analyst data points to the scale of this shift.

'The most significant challenge in scientific computing today is not technical; it is architectural and strategic'

Hyperion Research has reported that the HPC, AI and technical computing market grew by 23.5% in 2024 and projected that the overall HPC and technical computing market will exceed $100 billion by 2028. It also reported that AI is now used by more than 78% of HPC sites worldwide. Intersect360 Research has indicated that the worldwide market for accelerated and high-performance data centre infrastructure serving AI workloads reached $193bn in 2024, up 121% year-on-year, with the HPC segment growing by 24.1%.

These figures suggest that scientific computing budgets are expanding upward into larger AI and HPC systems, and outward into the data, workflow, governance and software layers needed to translate that capacity into reliable scientific and engineering outcomes.

Spending is rising, but so are expectations

Of the respondents to our SCW75 survey, 23% report they are planning to increase spending significantly, by more than 20%, while 34% plan to increase spending moderately, by between 5% and 20%. The survey also indicates that many respondents are already responsible for sizeable projects or infrastructure, with 34% reporting that they manage projects or infrastructure of more than £5m and 23% reporting budgets of between £1m and £5m.

This spending will likely not be confined to traditional HPC centres. Hyperion Research's public cloud forecast indicates that HPC-AI cloud spending is projected to grow by almost 20% annually between 2025 and 2029, with biosciences and computer-aided engineering (CAE) among the key verticals.

For lab-based industries, the figures underline the growing role of computation in research workflows.

Pharmaceutical, biotechnology, chemical and materials companies increasingly rely on computational methods to manage complex data, guide experiments, accelerate candidate selection and support decision-making.

Their needs, however, extend beyond raw compute into laboratory informatics, data governance, secure access controls, auditability and integration with experimental workflows.

In engineering, simulation, CAE and digital engineering tools are being used to reduce physical prototyping, improve design confidence and shorten development cycles, with cloud-based deployment growing faster than the overall CAE market.

The SCW75 by country

The US claims 31 of the 75 spots in the inaugural SCW75 – a reflection of the scale and maturity of its scientific computing infrastructure and the density of research universities, national laboratories and tech companies driving the field forward.

The UK is a strong second, with 21 honourees, underlining the depth of British research computing talent. Germany is third, with 10 representatives, while the remainder of the list is drawn from 11 further countries across four continents, — from Europe: Finland, the Netherlands, Norway and Sweden in the north, to Croatia, Italy, Spain and Switzerland further south.

Beyond Europe and North America, the inaugural list includes honourees based

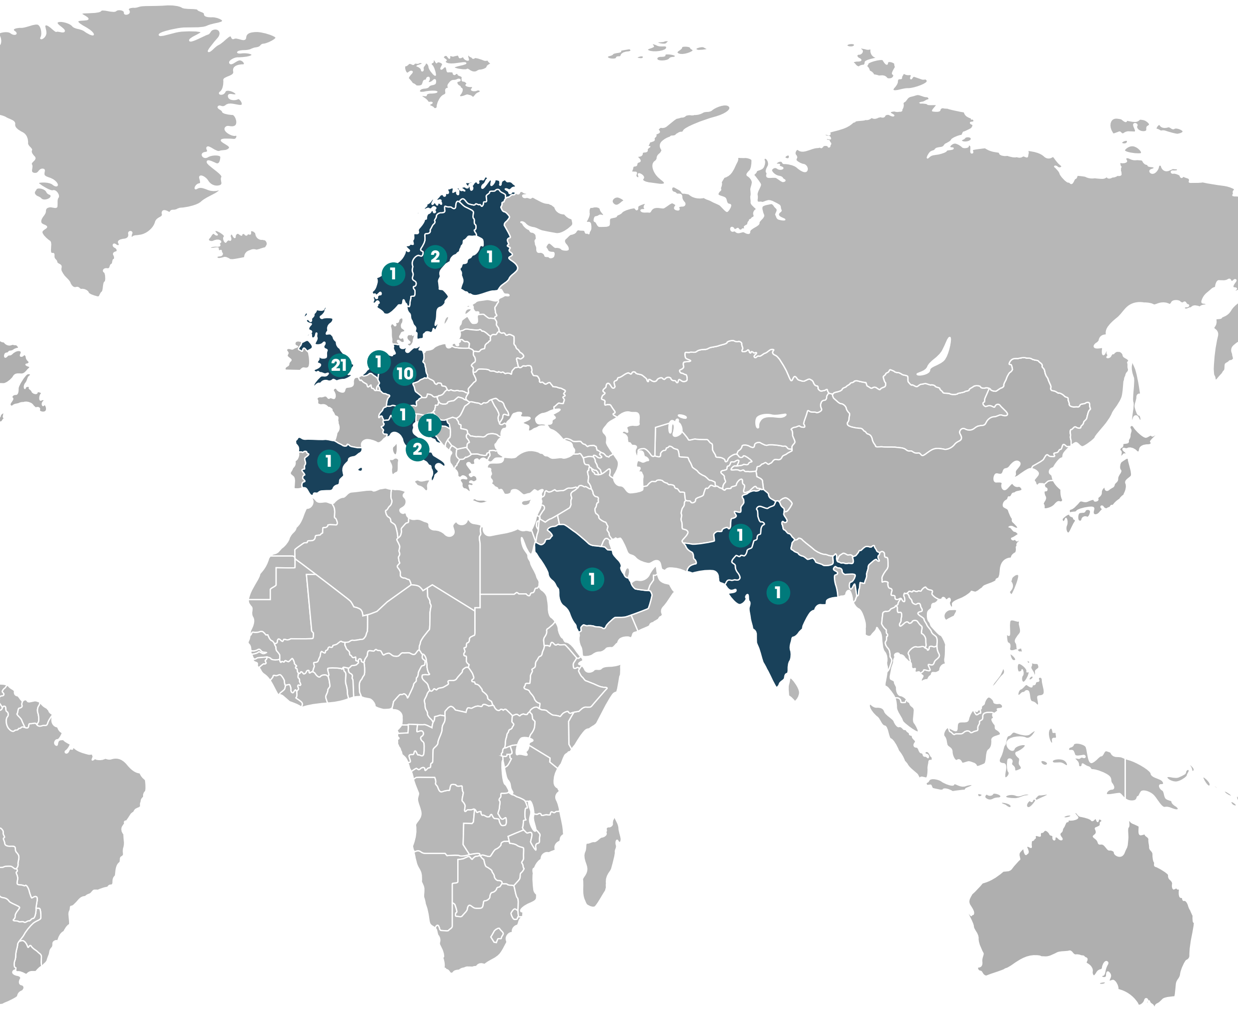

Country	# of SCW75 honourees
US	31
UK	21
Germany	10
Italy	2
Sweden	2
Croatia	1
Finland	1
India	1
The Netherlands	1
Norway	1
Pakistan	1
Saudi Arabia	1
Spain	1
Switzerland	1

in India, Saudi Arabia and Pakistan. That breadth matters. Scientific computing is a global endeavour, and the infrastructure challenges of scale, sustainability and accessibility are shared across research communities worldwide. As the SCW75 grows in future years, we hope that nominations from Asia, the southern hemisphere and emerging research economies will grow with it.

AI is changing how simulation is used

The most immediate force behind this transformation is AI computing. The rapid adoption of machine-learning, foundation models and data-intensive analytics has created strong demand for GPUs, high-bandwidth memory, specialist storage, fast networking and cloud-based compute capacity.

But in scientific computing, AI is not replacing established modelling and simulation methods. Instead, it is increasingly being combined with them. In engineering, AI is being used to accelerate design-space exploration, build surrogate models and reduce the cost of repeated simulation. In life sciences, it helps researchers interpret genomics, imaging, real-world data and experimental results. In materials science and chemistry, AI is being used to search vast chemical and materials spaces that would be impractical to explore experimentally.

That shared dependence is reshaping spending priorities. Organisations no longer invest only in faster machines or larger clusters. They are also spending on cloud platforms, data pipelines, observability tools, software licences, workflow systems, storage architectures, model validation frameworks and specialist staff. The pressure is not simply to run bigger jobs, but to make computational work more accessible, trusted and repeatable.

Samar Aseeri, Computational Scientist at the King Abdullah University of Science and Technology, noted that the practical adoption of advanced computing can be as important as the technology itself: "A significant challenge is bridging the gap between rapidly evolving computational technologies and their practical, scalable adoption in research environments."

Aseeri also highlighted the human and ecosystem challenges that can accompany technical change. "There is often a gap between advanced computational methods and the researchers who could benefit from them. Addressing this requires not only technical innovation, but also sustained efforts in education, community building and knowledge transfer," he said.

That point is central to the current investment cycle. The limiting factor is whether these systems can be embedded into everyday scientific workflows without creating new barriers for researchers and engineers.

Lab-based industries need trusted data as much as compute

The challenge is especially visible in lab-based industries. These increasingly depend on computational approaches, but their workflows are constrained by experimental validation, regulation, data quality and reproducibility. AI can help identify promising candidates, optimise processes or detect patterns in complex datasets, but scientific value still depends on whether those outputs can be tested, explained and trusted.

Our SCW75 survey responses repeatedly point to this issue. Respondents identified validation, reproducibility and trustworthiness as major challenges, particularly where computational models must be checked against incomplete, indirect or expensive real-world data. They also highlighted the difficulty of accessing high-quality, integrated and AI-ready multimodal data, especially in sensitive clinical and medical contexts.

That makes scientific computing investment in lab-based industries broader than merely the price of buying GPUs and servers. Organisations need electronic lab notebooks (ELNs), laboratory informatics management systems (LIMS), scientific data platforms, workflow systems, metadata standards and integration layers that can connect instruments, experiments and computational models. That means linking sample metadata, assay results, instrument files, experimental protocols, quality records and computational outputs in ways that can be searched, governed and reused.

Without those foundations, the cost of AI and HPC could rise, with no corresponding improvement in productivity. A model trained on incomplete or poorly described data may produce outputs quickly, but that has limited value if researchers cannot trace the source data, reproduce the workflow or understand the confidence attached to the result.

Engineering spending is being pulled towards simulation and digital design

For engineering industries, the same issue appears in a different form. Simulation and digital engineering can reduce the need for physical prototyping, shorten development cycles and improve design confidence. However, this places greater pressure on the credibility of models, the quality of input data and the ability of engineers to understand uncertainty. As AI enters these workflows, organisations must ensure that speed does not come at the expense of reliability.

The practical use cases are increasingly broad. Automotive and aerospace engineers use simulation for aerodynamics, crash analysis, thermal management and structural performance, while electronics, semiconductor, energy and machinery companies use digital models to explore system behaviour before physical assets are built or modified.

Recent consolidation in engineering software shows the strategic value being placed on simulation and AI-enabled design. Siemens completed its acquisition of Altair Engineering for an enterprise value of approximately $10bn, adding capabilities in mechanical and electromagnetic simulation, HPC, data science and AI. Synopsys completed its acquisition of Ansys in 2025, positioning the combined company in an expanded $31bn total addressable market. Cadence also completed its acquisition of Hexagon's design and engineering business in 2026, strengthening its position in physical AI, multiphysics and system design.

The SCW75 by gender

The inaugural SCW75 features 23 women among its 75 honourees, representing 31% of the list – a figure that compares reasonably well with the wider scientific computing landscape, though one we are determined to improve upon in future years.

For context, women remain underrepresented across the computing and research infrastructure workforce more broadly. According to US Department of Labor statistics for 2025, women hold just over 25% of 'computer and information systems managers' occupations in the US.

In HPC specifically, a paper entitled 'Representation of women in HPC conferences' from Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis in 2021, found that women represent only 10% of all HPC authors.

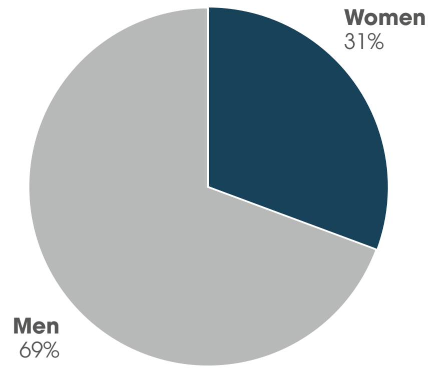

These deals perhaps indicate that major engineering software suppliers expect growing demand for integrated simulation and AI-enabled design workflows. For engineering companies, the supporting infrastructure increasingly includes cloud HPC, on-premises clusters, GPU capacity, high-speed storage and data management systems. The goal is not simply to replace physical testing, but to use computation earlier and more frequently in the design cycle, with physical validation reserved for where it adds the most value.

Strategy becomes a scientific challenge

As scientific computing becomes more central to research and industrial strategy, infrastructure planning becomes harder. The challenge is no longer only technical performance. It also involves sovereignty, procurement timing, access, governance, energy use and the ability to update systems.

Sadaf Alam, CTO and Director of Advanced Computing Strategy at the Bristol Centre for Supercomputing, University of Bristol, said: "The most significant challenge in scientific computing today is not technical; it is architectural and strategic: delivering sovereign, sustainable, and federated AI infrastructure at the pace science and policy demand, while the underlying technology landscape changes faster than any procurement or governance cycle can accommodate."

Alam's point captures a tension facing national facilities, universities and large industrial users. AI hardware generations can move faster than procurement, facilities planning and governance processes. Capacity may exist, but it still has to be made securely and equitably available to health researchers, government agencies, industry and early-career researchers. Energy usage and sustainability add a further constraint.

"These three challenges are interdependent. Solving any one without the others simply moves the bottleneck. The field needs infrastructure leaders who can hold all three in view and design systems that address them together," added Alam.

Complexity is now a cost issue

As organisations spend more on HPC, AI and cloud infrastructure, the cost of underutilisation also rises. A poorly optimised workload is no longer just a technical inconvenience; it can mean wasted GPU hours, higher cloud bills, increased energy consumption and delayed research.

Ayesha Afzal, Researcher at Erlangen National High Performance Computing Center, calls this a 'transparency crisis' – "the widening gap between the extreme complexity of exascale architectures and our ability to predict, interpret and optimise their behaviour".

Scientific computing environments are heterogeneous, combining CPUs, GPUs, specialist accelerators, distributed storage, cloud resources and complex software stacks creating powerful capabilities. But it also makes systems harder to optimise, debug and explain.

"We have reached a point where trial-and-error benchmarking is no longer a viable engineering strategy; it is a costly and unsustainable barrier to entry," said Afzal.

This argument is particularly relevant as energy and cloud costs become more visible. Organisations need better ways to predict performance, understand bottlenecks and quantify trade-offs between runtime, cost and power consumption before deployment. Without that visibility, scientific computing becomes harder to plan and justify.

For many domain scientists, a barrier to progress may therefore be usability rather than raw capacity. Advanced tools exist, but they remain difficult to use in routine research settings. If scientific computing can only be used effectively by a small group of specialists, the benefits of investment will be limited. ■

SCW75

High-performance computing

Ayesha Afzal

Organisation: Erlangen National High Performance Computing Center (NHR@FAU)
Role: Researcher
Based in: Erlangen, Germany

Ayesha Afzal develops predictive models that transform how supercomputing systems are used, replacing costly trial-and-error approaches with analytical insight. Her work aims to make large-scale computing more efficient, sustainable and accessible.

It is a perspective shaped by a career of deliberate evolution. "The black-box nature of supercomputing was a barrier to true scientific scalability," she explains. That realisation marked a turning point, shifting her work from using computational systems to fundamentally rethinking how their performance is designed and understood.

Dr Afzal's research sits at the intersection of performance modelling, energy efficiency and high-performance computing architecture. Her work focuses on developing predictive "white-box" models that allow scientists to anticipate how complex applications will behave on modern supercomputers, before they are ever run at scale.

She began with a foundation in electrical engineering at UET Lahore, followed by a pivotal move into computational engineering at FAU Erlangen-Nürnberg. It was there she began to question conventional approaches to performance optimisation, which often rely on costly trial-and-error benchmarking. During her doctoral research, she developed first-principles models for large-scale parallel systems, earning the highest distinction. This work laid the foundation for DisCostiC, a digital twin simulation framework designed to model and predict application performance across thousands of computing nodes.

At Friedrich-Alexander-Universität Erlangen-Nürnberg, and within the National High Performance Computing Center Erlangen (NHR@FAU), Afzal now leads efforts to address one of the defining challenges of the exascale era: complexity. As supercomputing systems become increasingly heterogeneous, with deep memory hierarchies and intricate interconnects, understanding performance behaviour has become both technically and economically demanding.

Her solution is to replace empirical tuning with predictive modelling. Through DisCostiC, she enables researchers to simulate the behaviour of parallel applications, capturing communication patterns, memory contention and load imbalance, without the need for full-scale execution.

This approach not only improves efficiency but also reduces the computational cost associated with large-scale experimentation. Alongside performance modelling, Afzal's work addresses the growing importance of energy efficiency in scientific computing. Within national initiatives, she develops mathematical frameworks to optimise performance-per-watt and translate abstract sustainability goals into actionable engineering strategies. In doing so, she is helping to ensure that large-scale computing remains both economically viable and environmentally responsible.

Her work is supported by a sophisticated ecosystem of tools, including system-level profiling frameworks such as LIKWID and Score-P, which are integrated into her predictive modelling pipeline.

These tools enable high levels of reproducibility and transparency, allowing complex workloads to be analysed with unprecedented clarity.

Yet Afzal emphasises that the challenges facing the field extend beyond technology.

She describes a "transparency crisis" in scientific computing, where the increasing complexity of systems outpaces our ability to interpret and optimise them. At the same time, the growing importance of energy constraints introduces a new dimension, requiring careful balancing of computational performance and power consumption.

For Afzal, addressing these challenges requires a fundamental shift – from observing performance to designing it. Her work on digital twin frameworks aims to bridge the gap between measurement and understanding, enabling a more analytical and sustainable approach to high-performance computing.

Looking ahead, she aims to establish digital twin simulation as a standard component of the exascale software development lifecycle.

By enabling researchers to pre-optimise applications and explore performance behaviour in virtual environments, she hopes to significantly reduce reliance on costly benchmarking and accelerate scientific discovery. A key milestone will be the integration of performance and energy modelling, making "sustained performance-per-watt" a central design principle.

Beyond technical contributions, Afzal is deeply engaged in building the human infrastructure of scientific computing.

She has led international collaborations, contributed to multinational research initiatives and founded the NHR Women in HPC chapter, promoting inclusivity and expanding access to the field. Her work reflects a belief that meaningful progress depends not only on technical innovation but also on strong, diverse communities.

Her advice to early-career researchers reflects this dual focus. Technical credibility, she argues, is essential, gained through direct engagement with real systems and challenges.

At the same time, lasting influence comes from addressing systemic problems and building collaborative ecosystems.

"Don't just optimise your code, optimise the ecosystem," she says.

Sadaf Alam

Organisation: Bristol Centre for Supercomputing (BriCS), University of Bristol
Role: Chief Technology Officer and Director of Advanced Computing Strategy
Based in: United Kingdom

As demand for artificial intelligence (AI) and high-performance computing accelerates, building scalable, efficient and accessible infrastructure has become a critical challenge. Operating at the intersection of technology and strategy, Dr Sadaf Alam is shaping the next generation of research computing infrastructure. Her work spans system architecture, data access and governance, with an emphasis on enabling large-scale, data-driven science.

Her path into high-performance computing began with a foundation in computer architecture, culminating in a PhD at the University of Edinburgh. There, her work on simulating the UKQCD computer established a guiding principle that has shaped her career: meaningful scientific outcomes depend on the co-design of hardware and software.

This perspective informed her early work at Oak Ridge National Laboratory, where she contributed to the development of the Oak Ridge Leadership Computing Facility at a time when GPU acceleration was still considered experimental. Her early commitment to GPU-based architectures, particularly for molecular dynamics and biomolecular simulation, placed her at the forefront of a transition that would later define mainstream HPC.

A major phase of her career unfolded at the Swiss National Supercomputing Centre (CSCS), where over more than a decade she progressed from Computer Scientist to Chief Technology Officer. During this period, CSCS evolved into a globally recognised centre for advanced computing. Among the defining achievements were leading the development of the world's first GPU-based operational numerical weather prediction system for MeteoSwiss and serving as Chief Architect of Piz Daint, which ranked third on the TOP500 and was the highest-ranked supercomputer in Europe at the time.

Her most ambitious undertaking to date has been the creation of the Bristol Centre for Supercomputing (BriCS) at the University of Bristol. Founded in 2022, BriCS was built from the ground up and has grown rapidly into a nationally significant institution. Under her leadership, the centre has delivered major infrastructure projects including Isambard-AI, a top-tier supercomputer recognised for both performance and energy efficiency, and AIRRPORT, a federated national portal providing access to AI research resources across academia, government and industry.

Today, Dr Alam's work focuses on a central challenge: designing and sustaining sovereign AI infrastructure. This involves not only building high-performance systems, but also addressing questions of access, governance, security and sustainability. Her work spans the full computational stack, from hardware architecture, such as GPU-accelerated systems based on advanced processor designs, to federated access platforms and open, cloud-native system management tools.

A key aspect of this effort is enabling secure and equitable access to computing resources. Through initiatives that connect supercomputing infrastructure with trusted research environments, her work supports applications such as health data analysis in areas including cardiac imaging and cancer research, where both computational power and data governance are critical.

At the same time, she is advancing open and secure approaches to system management through international collaborations, including work on open-source platforms that reduce reliance on proprietary technologies and improve resilience across national infrastructure. Her leadership also extends to international partnerships, including large-scale collaborations to establish shared approaches to sovereign AI systems.

Despite rapid technological progress, Dr Alam identifies the most significant challenge in the field as strategic rather than purely technical. The pace of change in AI hardware and software far exceeds traditional infrastructure planning cycles, necessitating modular, adaptable systems that can evolve over time. This is compounded by the need to balance accessibility, governance and sustainability, three interconnected challenges that must be addressed simultaneously.

Looking ahead, her ambitions include establishing scalable models for federated access to national computing resources, advancing international collaboration in AI infrastructure, and preparing for the convergence of AI and quantum computing. The aim is to ensure that future research systems are not only powerful, but also secure, sustainable and widely accessible.

Her advice to early-career researchers reflects this broader perspective. Technical expertise is essential, but influence comes from understanding the needs of the research community and delivering solutions that enable real outcomes. Building strong foundations, embracing interdisciplinary collaboration and being willing to take early risks on emerging technologies are all key elements of long-term impact.

At its core, Dr Alam's work is about enabling discovery at scale. By designing and delivering the infrastructure that underpins modern research, she is helping to shape how science is conducted, ensuring that computational capability continues to expand in step with scientific ambition.

Nisha Agrawal

Organisation: Centre for Development of Advanced Computing (CDAC) India
Role: Joint director
Based in: Pune, India

Nisha Agrawal has spent more than two decades helping researchers unlock better performance from increasingly complex computing systems. Her work focuses on a persistent challenge in scientific computing: ensuring that large-scale simulations can fully exploit modern heterogeneous architectures.

At the Centre for Development of Advanced Computing in Pune, India, she specialises in GPU computing, parallel programming and performance optimisation across scientific applications. Her research spans molecular dynamics, density functional theory, weather forecasting and flood simulation – areas where computational efficiency directly affects the speed and scale of scientific discovery.

Agrawal has worked extensively with CUDA, OpenACC, MPI and OpenMP, optimising applications such as NAMD, WRF and ANUGA for GPU-accelerated supercomputers. Her research includes multi-GPU scaling, hardware counter analysis and memory optimisation, helping scientists run complex workloads more efficiently across CPU-GPU systems.

She also chaired the First Workshop on Accelerating Scientific Discovery in Chemistry and Material Science with HPC and AI, held in conjunction with the 32nd IEEE International Conference on High Performance Computing, Data and Analytics (HiPC 2025) in Hyderabad at the end of 2025.

Alongside her technical contributions, she has become a strong advocate for HPC education. As an NVIDIA-certified mentor, she regularly supports global OpenHackathons and has delivered more than 30 invited talks at institutions including IITs and IISERs.

Her work reflects a broader commitment to making advanced computing infrastructure more accessible to researchers across disciplines.

Martin Berzins

Organisation: University of Utah
Role: Professor, School of Computing and Scientific Computing and Imaging Institute
Based in: Salt Lake City, United States

Martin Berzins has spent his career addressing a fundamental challenge in scientific computing: how to model increasingly complex physical systems at scale. His research combines applied mathematics, computer science and engineering, with a long-standing focus on partial differential equation software for extreme-scale simulations.

He is best known for advancing the Uintah framework, helping to transform it into a task-based system capable of handling complex multiphysics simulations across large supercomputers.

That work has supported applications including combustion, materials science, fire modelling and explosive simulations.

Trained in mathematics at the University of Leeds, where he completed both his undergraduate degree and PhD in numerical analysis, Berzins built his early career at Leeds before moving to the University of Utah in 2003.

More recently, Berzins has focused on portability and performance at exascale, combining task-based runtimes with tools such as Kokkos to ensure scientific codes can run efficiently across increasingly heterogeneous systems.

His work reflects a broader goal: enabling scientists to tackle problems that would otherwise be computationally impossible.

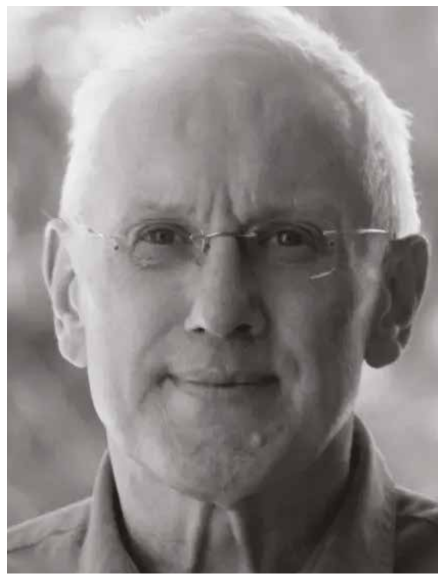

Bhushan Bonde

Organisation: University of Suffolk
Role: Visiting Professor of Computational Biology
Based in: UK

Bhushan Bonde works across scientific computing, artificial intelligence and drug discovery, applying advanced computational methods to accelerate the identification and testing of potential therapies.

With a background spanning pharmaceutical R&D, high-performance computing and machine-learning, Bonde has built a career focused on applying emerging technologies to healthcare challenges. His work combines large-scale computing infrastructure with AI-driven methods to shorten research timelines and improve decision-making in drug development.

A major area of focus has been the application of high-performance computing, deep learning and generative AI to pharmaceutical research. He has worked on computational approaches that enable faster molecular modelling, simulation and drug screening workflows, helping researchers to process larger datasets and identify promising candidates more efficiently. His work has also expanded into quantum computing, where he is exploring novel algorithms that could help solve complex biomedical problems.

His work has also expanded into quantum computing, where he is exploring novel algorithms for complex biomedical problems. During the pandemic, he received a Moonshot Covid Drug Discovery Award for applying HPC to accelerate treatment discovery efforts.

Samar Aseeri

Organisation: King Abdullah University of Science and Technology **Based in:** Saudi Arabia
Role: Computational Scientist

Samar Aseeri develops computational approaches that bridge traditional high-performance computing with next-generation quantum methods. Her work combines technical research with community-building efforts to expand access to advanced computing capabilities.

Her early training was in applied mathematics, but a pivotal moment came after completing her PhD in 2009, when she was selected for advanced supercomputing training at the IBM Thomas J Watson Research Center. There, she developed a foundation in large-scale systems that would define her future direction.

In 2010, she joined a university supercomputing laboratory as a computational scientist, becoming one of the first specialists supporting what was then her country's first academic supercomputer. This marked a decisive shift from theoretical work to applied HPC, where she combined infrastructure support with research in parallel algorithms, performance optimisation and scalable numerical methods.

Operating in an environment with limited local expertise, Aseeri developed a largely self-directed approach. Over time, her work expanded beyond technical contributions to include benchmarking, reproducibility and the broader challenge of enabling others to use advanced computing effectively.

Today, her research sits at the intersection of HPC, reproducible science and emerging hybrid quantum-classical computing.

A central theme of her work is bridging the gap between established HPC methods and next-generation quantum approaches. She leads a quantum computing reading group at KAUST, bringing together researchers from different disciplines to explore hybrid algorithms and their integration into existing scientific workflows. Alongside this, she collaborates internationally on scalable methods for hybrid computing, addressing practical challenges such as efficiency, portability and real-world applicability.

Her work also extends beyond research into ecosystem development.

Since 2017, she has led international HPC initiatives focused on connecting fragmented communities, promoting knowledge exchange and supporting capacity building across regions. These efforts reflect a recognition that access to expertise can prove as important as access to hardware in advancing scientific computing.

Technically, Aseeri's work is grounded in distributed computing and numerical simulation. She has worked extensively with MPI-based programming (including MPI4Py), scalable FFT libraries and spectral methods on HPC clusters. By applying domain decomposition techniques and optimising communication strategies, she has improved efficiency in large-scale simulations, including reducing communication overhead in parallel FFT workloads by up to 10%.

Equally important has been her focus on reproducibility. Through structured benchmarking and reproducible workflows, she has contributed to improving how performance is evaluated and compared across systems, an ongoing challenge in scientific computing.

Despite rapid advances in computing capability, she sees a persistent gap between technological potential and practical adoption. HPC systems continue to face challenges in scalability, communication bottlenecks and portability, while quantum computing introduces new layers of complexity, from algorithm design to the lack of mature tooling.

At the same time, she highlights a less technical, but equally significant, issue: the shortage of expertise needed to fully exploit these technologies.

Looking ahead, Aseeri's focus is on building bridges between classical and quantum computing, and between technology and the people who use it. Over the next few years, she aims to advance hybrid computational workflows and expand access to quantum computing through structured education and collaborative research. A key milestone will be demonstrating practical use cases where quantum-classical methods can deliver meaningful scientific outcomes.

Her advice to early-career researchers reflects her own experience: build strong fundamentals, stay adaptable and be prepared to create your own opportunities.

In a field that evolves as quickly as scientific computing, persistence and curiosity remain as important as technical skill.

At its core, Aseeri's work highlights an often-overlooked dimension of scientific computing. Progress is not only about faster systems or more advanced algorithms, it is also about building the knowledge, communities and practices that allow those technologies to be used effectively.

"Be persistent. Progress in scientific computing often comes with technical and institutional challenges, but consistent effort, curiosity and a willingness to collaborate will open doors over time," notes Aseeri.

David Bader

Organisation: New Jersey Institute of Technology
Role: Distinguished Professor
Based in: Newark, New Jersey, United States

David Bader develops computing methods that make it possible to analyse some of the world's largest and most complex datasets. His work in high-performance computing and graph analytics has helped shape how researchers study networks in fields ranging from cybersecurity to biology.

It is a career defined by pivotal moments. "Each pivot has reinforced a throughline: building the people, platforms and intellectual frameworks that let computation transform science and society," Bader explains. From early innovations in supercomputing to leadership in data science, his work has consistently focused on enabling computation at scale.

Prof Bader's research sits at the crossroads of high-performance computing, data science and large-scale network analysis. His work focuses on developing algorithms and systems capable of analysing massive, complex datasets – particularly graphs that underpin applications in cybersecurity, biology, social networks and infrastructure.

His foundational breakthrough came in 1998, when he built one of the first high-performance Linux supercomputers at the University of New Mexico. At a time when proprietary systems dominated, this work demonstrated that commodity hardware and open-source software could deliver competitive performance at significantly lower cost. The approach would go on to reshape the field, with Linux now underpinning virtually all of the world's fastest supercomputers.

Following this, Bader moved to Georgia Institute of Technology, where he founded and chaired the School of Computational Science and Engineering.

This marked a shift from individual research to institution-building, creating a new academic structure dedicated to interdisciplinary computational methods.

He later joined New Jersey Institute of Technology, where he established the Institute for Data Science. Here, his work reflects a broader convergence between high-performance computing and data-driven research, recognising that large-scale computation is now central across disciplines. Alongside this, he has remained deeply engaged with the global community, serving as Editor-in-Chief of ACM Transactions on Parallel Computing and contributing to major benchmarking initiatives such as Graph500.

Today, his research focuses on scalable graph analytics and network algorithms. Through frameworks such as Arachne and Arkouda, his group is developing interactive tools that allow analysts to work with billion-edge graphs using accessible programming environments.

These systems aim to democratise capabilities that were previously limited to specialists, enabling real-time exploration of complex networks.

His work extends into broader questions of computing's societal impact. As a recognised expert on AI and computing infrastructure, Bader contributes to public discourse, helping to translate technical developments into accessible insights for policymakers and wider audiences.

A central challenge in his field lies in the nature of graph computation itself. Unlike traditional numerical workloads, graph problems are irregular, memory-bound and difficult to optimise on modern hardware. This creates a gap between theoretical algorithmic advances and their practical application. Bridging this divide, making large-scale graph analytics both performant and accessible, remains a core focus.

Looking ahead, Bader aims to make real-time analysis of billion-scale dynamic graphs a routine capability.

His goal is to enable researchers and practitioners to interactively analyse evolving networks using intuitive tools, without requiring deep expertise in parallel computing. Achieving this would significantly expand access to advanced computational methods and accelerate insight across a range of domains.

For early-career researchers, his advice reflects the lessons of his own journey. Building something tangible, a system that works at real scale, can have far greater impact than theoretical contributions alone. At the same time, investing in people and institutions is essential. Mentoring students, fostering collaborations and shaping research environments all contribute to long-term influence.

He also emphasises the importance of communication. Translating complex ideas for broader audiences, including policymakers and the public, is increasingly critical as computing technologies shape society at large.

With recognition including the IEEE Sidney Fernbach Award and induction into the Computer History Museum, Bader's contributions have lasting impact. Yet throughout his career, the focus has remained: enabling computation to move beyond theory and into practice, where it can drive meaningful change across science and society.

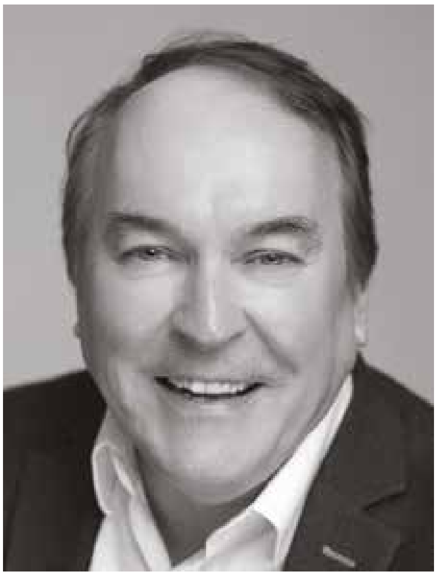

Peter Coveney

Organisation: University College London
Role: Director, Centre for Computational Science; Professor of Physical Chemistry
Based in: London, UK

As scientific computing moves towards exascale, Peter Coveney has focused on ensuring researchers can make meaningful use of that power. His work spans large-scale simulation, performance optimisation and uncertainty quantification, helping scientists extract reliable insight from complex computing environments.

Now Director of the Centre for Computational Science and Professor of Physical Chemistry at University College London, Coveney has built an influential career in computational science working across physics, chemistry and medicine. Early in his career, he worked on statistical mechanics and complex systems research alongside Ilya Prigogine before moving into industrial research at Schlumberger, where he helped pioneer large-scale simulations of complex fluids and materials.

He led the development of HemeLB, a blood-flow simulation platform capable of modelling circulation in complex vascular systems, including the human brain. His team has also developed GPU-accelerated versions that scale to large supercomputing systems.

Coveney is also a leading voice on reproducibility, validation, verification and uncertainty quantification, arguing that simulations must be trustworthy enough to support decisions in healthcare, engineering and materials science.

Lori Diachin

Organisation: Lawrence Livermore National Laboratory
Role: Principal Deputy Director of Computing
Based in: California, United States

Lori Diachin is a senior high-performance computing leader whose work has helped guide the United States' transition to exascale scientific computing.

As Project Director for the US Department of Energy's Exascale Computing Project, after previously serving as deputy director, Diachin has played a central role in coordinating one of the most ambitious computing programmes in scientific research. The project was designed to deliver a capable exascale ecosystem spanning applications, software technologies, hardware integration and user readiness.

At Lawrence Livermore National Laboratory, her leadership has focused on ensuring that exascale systems are usable for real scientific discovery, not simply theoretical performance milestones. That requires aligning application teams, software developers, facilities and hardware vendors around common goals.

Diachin's work demonstrates the importance of computational engineering leadership at the programme scale. By helping coordinate software, applications and infrastructure, she has contributed to the foundations that enable researchers to exploit the world's most powerful systems.

Annarosa Farina

Organisation: European Institute of Oncology/Centro Cardiologico Monzino
Role: Chief Information Officer (CIO)
Based in: Milan, Italy

Annarosa Farina is helping modernise digital infrastructure at two of Italy's leading medical institutions, supporting research, clinical care and data-driven innovation.

As Chief Information Officer at the European Institute of Oncology (IEO) and Centro Cardiologico Monzino (CCM), Farina leads technology strategy across organisations operating at the intersection of healthcare delivery, biomedical research and precision medicine. She oversees the systems that increasingly underpin clinical research, genomics, imaging, patient data management and digital transformation at both institutions.

Her work focuses on building infrastructure capable of supporting growing volumes of clinical and research data, particularly as oncology and cardiovascular medicine become increasingly dependent on computational analysis. This includes modernising enterprise systems, strengthening cybersecurity, improving interoperability and ensuring researchers can securely access resources for data-intensive work.

As precision medicine continues to expand, her work reflects a broader transformation across healthcare: ensuring hospitals are equipped not only to treat patients, but also to operate as highly connected, data-driven research environments.

Sunita Chandrasekaran

Organisation: University of Delaware
Role: Director of the First State AI Institute
Based in: United States

From large-scale simulations to artificial intelligence (AI)-driven models, Prof Sunita Chandrasekaran's work focuses on enabling complex scientific applications to run efficiently on modern computing systems. Her research spans compilers, programming models and performance optimisation. Her work is defined by a consistent effort to bridge disciplines, ensuring that advances in high-performance computing (HPC) and AI translate into practical, scalable impact.

Chandrasekaran began with a foundation in electrical and electronics engineering, grounding her understanding in hardware systems. During her PhD at Singapore's Nanyang Technological University, she moved into the intersection of hardware and software, focusing on mapping algorithms onto field-programmable gate arrays (FPGAs), an early step into parallel and heterogeneous computing.

A key transition followed during her post-doctoral work at the University of Houston, where her focus shifted toward HPC software, programming models and compilers. This marked a decisive move from embedded systems into large-scale computational environments.

Since joining the University of Delaware in 2015, Chandrasekaran has built the Computational Research and Programming Lab, expanding her research into compiler design, GPU programming and usability of high-performance systems.

Her work has included developing validation frameworks for directive-based programming models and leveraging machine-learning to build predictive models in paediatric disease research. Access to leadership-class supercomputers further shaped her research, enabling her to stress-test systems and develop software for real-world scientific workloads.

Her most recent pivot extends beyond technical research into leadership and governance. She co-founded the AI Center of Excellence in 2021 and later established the First State AI Institute in 2025, where she now serves as Director. In parallel, she contributes to public policy as Vice-Chair of Delaware's AI Commission, helping to shape frameworks for the responsible deployment of AI at societal level.

Currently, her work spans several ambitious initiatives. She is exploring computational challenges in high-energy laser particle accelerator simulations, which require coordinating thousands of GPUs and real-time data movement across nodes. She is also leading efforts to build localised AI models for sensitive research data, addressing challenges in storage, scalability and multi-user accessibility.

Across the university, she is driving the adoption of AI tools for research and education, while also developing training programmes to support workforce development and small-business innovation.

Her scientific computing approach integrates HPC, machine learning and advanced performance engineering.

By combining algorithmic improvements with profiling and performance analysis tools, her work ensures that software can effectively utilise increasingly complex hardware systems.

A central challenge she identifies is the widening gap between rapidly evolving hardware and the software required to exploit it. As architectures grow more heterogeneous and specialised, maintaining efficient, portable and user-friendly tools becomes increasingly difficult.

Looking ahead, Chandrasekaran aims to apply AI and HPC to high-impact domains such as healthcare and finance, while also advancing digital twin technologies for cities. These models could play a critical role in disaster response and emergency evacuation planning, particularly in vulnerable urban environments.

Her advice to early-career researchers emphasises the importance of interdisciplinary thinking.

Building effective tools requires not only technical expertise but also a clear understanding of end users and real-world applications. In her view, long-term impact comes from designing systems that are both powerful and accessible.

Beyond her research, she is actively involved in shaping future computing infrastructure, including sustainability-focused initiatives for research computing in Delaware. Her work reflects a broader vision: aligning technological progress with societal needs, ensuring that advanced computing systems remain both impactful and responsible.

Florina M. Ciorba

Organisation: University of Basel
Role: Professor for High Performance Computing
Based in: Basel, Switzerland

At the intersection of theory, systems and scientific application, Florina M. Ciorba is redefining how high-performance computing (HPC) is designed, understood and used. Her research tackles not only performance at scale but also the deeper questions of reliability, efficiency and accessibility that determine whether HPC can truly support next-generation science.

It is a perspective shaped early in her career. "The gap between elegant algorithmic theory and the complexity of real HPC systems was both the most challenging and intellectually stimulating problem to address," she notes. That conviction has guided her work ever since, driving a research agenda focused on making high-performance computing more efficient, trustworthy and accessible.

Prof Ciorba's work spans scheduling, load balancing and system-level optimisation, with a strong emphasis on enabling scientific applications at scale. Her research aims to reduce the barriers that prevent domain scientists from fully leveraging HPC, ensuring that computational advances translate into real scientific impact.

Her journey began in Zalău, Romania, where an early fascination with computing led her to study computer engineering at the University of Oradea. A formative Erasmus exchange at the National Technical University of Athens opened the door to doctoral studies, where she spent six years developing theoretical foundations for scheduling and dynamic load balancing in parallel systems.

Following her PhD, a post-doctoral position at Mississippi State University grounded her work in real scientific applications, particularly in multi-scale materials science. This experience reinforced the importance of connecting computational performance with meaningful scientific outcomes.

A further pivotal period came at the Center for Information Services and High Performance Computing (ZIH) at Technische Universität Dresden, where she gained first-hand experience of operating large-scale HPC systems. There, her research expanded to include reliability, observability and fault tolerance – laying the groundwork for her later contributions to trusted computing.

Since joining the University of Basel in 2015, Ciorba has built a research group that addresses multiple dimensions of HPC, from extreme-scale simulation to privacy-preserving computing in healthcare.

Her work reflects a belief that HPC must evolve, not only in performance, but also in usability, trust and inclusivity.

Today, her research spans several interconnected areas. One focus is improving system visibility through tools such as SIREN, which uses lightweight, privacy-preserving telemetry to identify and classify workloads on large-scale systems. Another addresses "tail waste", a largely overlooked inefficiency where incomplete jobs result in lost computation and energy. By developing predictive models and autonomous checkpointing strategies, her work aims to recover wasted resources and improve sustainability.

Ciorba is also involved in large-scale scientific collaborations, including projects that integrate astrophysical simulation frameworks to enable new multi-physics studies at unprecedented scale. At the same time, she is advancing research into algorithm selection and scheduling for heterogeneous architectures, ensuring that systems can make optimal use of available resources.

More recently, her work has expanded into healthcare, where she is developing privacy-preserving AI pipelines for hereditary cancer care. By combining high-performance computing with secure data processing techniques, this research aims to support sensitive medical applications while maintaining strict data protection standards.

Despite these advances, Ciorba identifies a central challenge in the field: the growing complexity of optimisation. Modern HPC systems must balance multiple competing objectives – performance, energy efficiency, reliability and usability – without clear frameworks for evaluating trade-offs. This is compounded by a broader reproducibility challenge, where the lack of standardised benchmarks makes it difficult to compare approaches and assess real-world impact.

Looking ahead, she aims to demonstrate a fully autonomous control loop for job management on a production supercomputer. Such a system would monitor execution in real time, predict inefficiencies and trigger corrective actions automatically, reducing both wasted computation and energy use. Achieving this would represent a significant step towards more intelligent and sustainable HPC systems.

For early-career researchers, Ciorba emphasises the importance of working at the boundary between theory and practice. Addressing real-world challenges, building inclusive communities and maintaining a commitment to openness and reproducibility are key to lasting impact.

Seeking diverse mentorship and perspectives, she adds, can also play a crucial role in shaping a successful career.

Her work extends beyond research into community building and advocacy. As a co-founder of the Swiss chapter of Women in HPC and an active participant in international initiatives, she is committed to fostering diversity and collaboration within the field. She also contributes to broader efforts in sustainable computing, highlighting the importance of reducing energy waste in large-scale systems.

At its core, Ciorba's work reflects a vision of high-performance computing that is not only powerful, but also responsible, transparent and accessible – ensuring that its benefits can be realised across an increasingly diverse range of scientific domains.

Garth Gibson

Organisation: VDURA
Role: Chief Technology and AI Officer
Based in: United States

Few researchers have influenced scientific computing infrastructure as much as Garth Gibson. Across academia, entrepreneurship and research leadership, his work has helped define how modern high-performance systems store, move and protect data at scale.

Gibson is best known as the co-inventor of RAID (Redundant Array of Independent Disks), a breakthrough that transformed how large-scale computing systems manage reliability, availability and storage performance. RAID fundamentally changed the economics of data storage by enabling systems to scale capacity and throughput while maintaining resilience. It remains a foundational technology across supercomputing, enterprise infrastructure and cloud computing.

His influence extends beyond RAID. As a professor at Carnegie Mellon University, Gibson founded the Parallel Data Laboratory, a major research group in high-performance storage systems.

His work in parallel I/O, distributed storage and scale-out file systems helped establish principles behind modern HPC storage.

He also co-founded Panasas, now VDURA, translating academic research into production systems deployed across laboratories, universities and research centres.

Today, he focuses on storage for AI-driven and data-intensive computing.

Debra Goldfarb

Organisation: Amazon Web Services
Role: Director of Products and Strategy, Advanced Computing, Simulation and Quantum Computing
Based in: United States

Debra Goldfarb has spent her career helping high-performance computing evolve from a specialist technology into critical infrastructure for science, industry and national competitiveness.

Now Director of Products and Strategy for Advanced Computing, Simulation and Quantum Computing at Amazon Web Services, Goldfarb helps shape AWS's strategy across HPC, simulation and emerging quantum technologies. Her work includes partnerships with organisations such as Harvard Data Science Initiative, Lawrence Livermore National Laboratory and RIKEN, helping researchers access scalable computing infrastructure.

Her path into computing was unconventional. Originally trained in Chinese studies, Goldfarb joined International Data Corporation as a sales representative before identifying the growing strategic importance of HPC. She later built IDC's first dedicated HPC research practice and spent 16 years advising governments, vendors and investors.

Goldfarb focuses on expanding access to advanced computing. That includes initiatives such as the AWS-Harvard collaboration on ecological tipping points, the Missing Middle Initiative to broaden access to HPC resources, contributions to the America COMPETES Act, and the Wheeling Initiative, which introduced simulation tools to underserved schools.

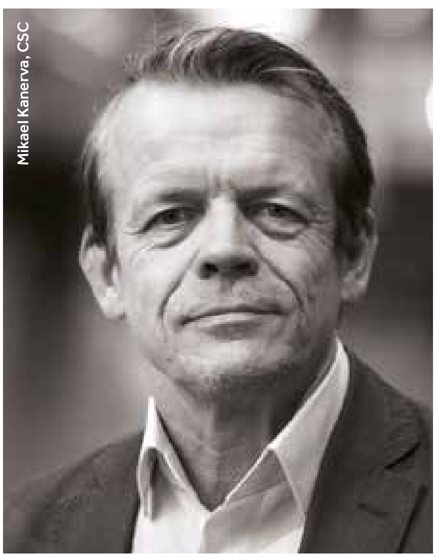

Kimmo Koski

Organisation: CSC – IT Center for Science
Role: Chief Executive Officer
Based in: Finland

Kimmo Koski has spent his career building the infrastructure that enables scientific discovery at national and international scale. From maintaining early supercomputers to leading one of Europe's most advanced research computing organisations, his career mirrors the evolution of scientific computing.

It is a journey shaped as much by opportunity as by intent. "Partly by accident I entered the area of scientific computing," he notes, recalling his early involvement with the Cray X-MP system at CSC.

At the time, the organisation was small, with just a few dozen employees, but it would progress to become the foundation for a career defined by growth, international collaboration and infrastructure leadership.

His work centres on sustaining the large-scale computing ecosystems that underpin modern research and innovation.

Today, his focus is ensuring CSC delivers world-class infrastructure and expertise across HPC, AI, data and quantum technologies.

A flagship example is the LUMI supercomputer, which is one of Europe's most powerful systems, alongside the development of the LUMI AI Factory for research and industrial innovation.

Thomas Lippert

Organisation: Jülich Supercomputing Centre
Role: Director
Based in: Germany

Thomas Lippert has spent his career building the infrastructure that enables Europe to compete at the highest level of computational research. His work has shaped large-scale supercomputing systems and Europe's strategy for technological sovereignty in HPC.

As Director of the Jülich Supercomputing Centre and a member of the board of Forschungszentrum Jülich, Lippert has played a central role in developing some of Europe's most advanced computing systems.

Under his leadership, Jülich has become one of the continent's most important hubs for exascale computing, AI infrastructure and quantum integration.

Originally trained as a physicist, Lippert built his academic career in computational physics, specialising in lattice quantum chromodynamics and large-scale simulation methods. That background shaped his understanding of how infrastructure decisions affect research outcomes.

At Jülich, he has overseen major systems including JUWELS and JUPITER, Europe's first exascale supercomputer. Lippert has also become a major voice in European digital strategy through his involvement in EuroHPC Joint Undertaking, helping to shape the continent's long-term approach to sovereign computing capability.

Simon McIntosh-Smith

Organisation: Bristol Centre for Supercomputing
Role: Director
Based in: UK

Simon McIntosh-Smith is helping to redefine how quickly national-scale computing infrastructure can be built, and how efficiently it can operate once deployed.

As Director of the Bristol Centre for Supercomputing (BriCS), McIntosh-Smith has built a career at the forefront of computer architecture and high-performance computing.

His work spans industry and academia, with a focus on systems that push performance boundaries while solving practical infrastructure challenges.

He began as a microprocessor architect at INMOS and STMicroelectronics, working on early multi-core, multi-threaded and SIMD CPU designs. He later co-designed one of the world's first fully programmable GPUs at PixelFusion before co-founding ClearSpeed Technology, where he helped develop many-core accelerators used in systems such as TSUBAME.

Subsequently, he founded the High Performance Computing Research Group at the University of Bristol before launching BriCS.

His most prominent recent achievement is Isambard-AI, a £225 million AI-focused supercomputer built rapidly on a new site and brought into production in just more than a year.

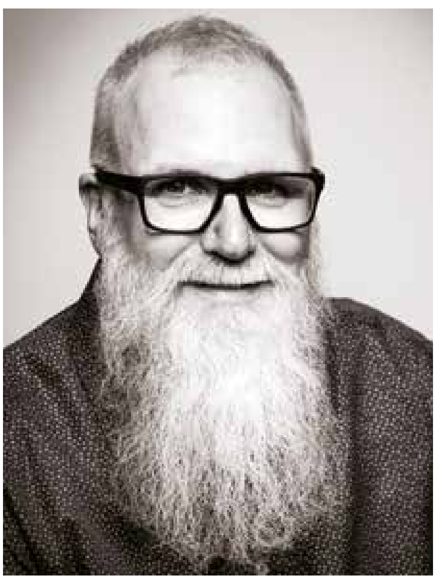

Michael Papka

Organisation: Argonne Leadership Computing Facility
Role: Director
Based in: Illinois, United States

Michael Papka has built his career around addressing one of scientific computing's most persistent challenges: ensuring researchers can effectively use the world's most powerful machines.

As Director of the Argonne Leadership Computing Facility at Argonne National Laboratory, Papka oversees some of the world's most advanced supercomputing infrastructure. His work combines large-scale computing systems, user enablement and emerging technologies.

Papka's early academic and research background spans physics, electrical engineering and computer science, culminating in a PhD in computer science from the University of Chicago. Much of his research has focused on scientific visualisation, large-scale data analysis, remote collaboration and research infrastructure. Through his work with the Electronic Visualization Laboratory and Argonne, he has contributed to tools and environments that help researchers explore complex datasets generated by large-scale simulations.

Before Argonne, Papka held senior roles at Northern Illinois University and the University of Chicago. Today, he leads operations around systems including Aurora, one of the world's first exascale supercomputers.

Sarah Neuwirth

Organisation: Johannes Gutenberg University, Mainz
Role: Professor and Director of NHR South-West HPC Center
Based in: Mainz, Germany

Modern supercomputing is no longer defined solely by raw performance, but by how efficiently systems can move and process data. Prof Sarah Neuwirth's work addresses this shift, focusing on data-centric approaches that improve both performance and energy efficiency in large-scale computing environments.

Neuwirth makes complex computing systems easier to understand. Her work focuses on one of most pressing problems of high-performance computing (HPC): how to make increasingly powerful infrastructures transparent, efficient and sustainable enough to support the next generation of science.

Her career has developed at the boundary between systems research and real-world supercomputing. During her doctoral work at Heidelberg University, she focused on parallel I/O and communication, with research stays at Oak Ridge National Laboratory proving especially formative. There, she learned how to identify gaps in leadership-class HPC software stacks and turn experimental ideas into production-ready solutions.

That experience shaped a career-long commitment to bridging research and operational computing. As deputy group lead of the Modular Supercomputing and Quantum Computing group at Goethe University Frankfurt, she moved into scientific leadership, building a new research group in close collaboration with the Jülich Supercomputing Centre and developing experimental infrastructure for systems research.

Today, as Professor at Johannes Gutenberg University Mainz and Director of the NHR South-West HPC Center, Neuwirth leads one of Germany's national high-performance computing centres. Her role combines scientific leadership, infrastructure strategy, system design, procurement and user support across a multi-institutional environment.

Her research focuses on explainable, data-centric and energy-aware HPC systems. A central concern is data movement: how information travels across increasingly complex hardware, storage and software layers. As modern workloads become more data-intensive, particularly with the rise of AI-driven science, data movement increasingly dominates both performance and energy use.

To address this, Neuwirth develops methods for cross-layer observability, reproducible benchmarking and model-driven performance analysis. Tools and frameworks such as XIO, DataCrumbs, FlexBench, JUBE-ML, and PtlBench help researchers identify bottlenecks and make system behaviour more interpretable.

Her work has demonstrated measurable impact, including reductions in I/O-related energy consumption and improvements in runtime performance.

Earlier in her career, she contributed to network-attached accelerator architectures through the DEEP project series and to load-balancing strategies for leadership-class systems at Oak Ridge. These experiences reinforced her focus on closing the gap between system architecture, software stacks and application behaviour.

For Neuwirth, the greatest challenge in scientific computing is that HPC systems have become multi-layered ecosystems whose complexity often exceeds our ability to reason about them. Existing tools offer fragmented views, while meaningful diagnosis still depends heavily on expert intuition. The result is a field that can build extraordinary machines, but does not always fully understand their efficiency, limitations or energy costs.

Her ambition over the next two-to-three years is to establish explainable, energy-aware data movement as a first-class design principle in HPC. By combining observability, workload modelling and AI-assisted analysis, she aims to make performance and energy behaviour transparent across the full software and hardware stack.

The goal is not simply faster computing, but more sustainable and reproducible computing. Integrating these capabilities into national and European HPC production environments would mark a significant step towards systems designed from the outset for performance-per-watt, scalability and trust.

Her advice to early-career researchers is to build technical depth, but never in isolation. Scientific computing, she emphasises, is fundamentally collaborative. Community engagement, shared infrastructure and work across boundaries between systems, applications, research and operations are often where the most meaningful opportunities emerge.

Neuwirth is also active in national and international HPC strategy, including EuroHPC projects, the ETP4HPC Strategic Research Agenda, the Gauss Alliance and the NHR Alliance. Her contributions to sustainable, reproducible and energy-aware HPC have been recognised with the PRACE Ada Lovelace Award for HPC and the ZONTA Science Award.

At its core, her work reflects a clear vision for supercomputing's future: systems that are not only powerful but also explainable, efficient and sustainable by design.

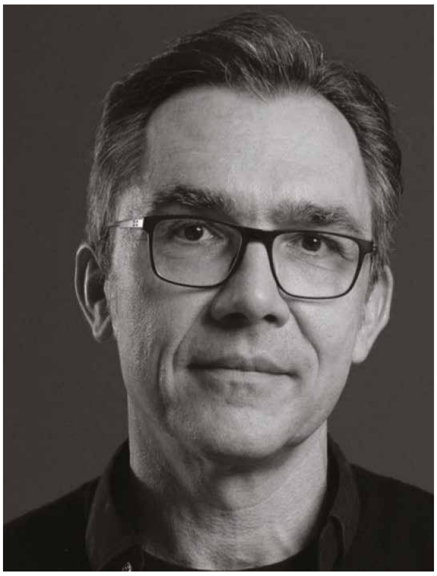

Jürgen Popp

Organisation: Leibniz Institute of Photonic Technology
Role: Scientific Director
Based in: Jena, Germany

Jürgen Popp has built his career around making complex molecular information visible, using photonics, spectroscopy and computational analysis to improve how diseases are detected and understood.

As Scientific Director of the Leibniz Institute of Photonic Technology and Professor at Friedrich Schiller University Jena, Popp is internationally recognised for advancing Raman spectroscopy, biophotonics and optical diagnostics. His work combines photonic instrumentation with computational methods that help translate complex biological signals into clinically useful insights.

Originally trained in chemistry, Popp's early research focused on vibrational spectroscopy before expanding into biomedical photonics. Over time, he became a leading figure in applying Raman-based techniques to infectious disease diagnostics, cancer detection and precision medicine.

A major focus of his recent work has been developing rapid diagnostic technologies that combine spectroscopy with machine learning and advanced data analysis.

These systems are designed to identify pathogens, classify disease states and support faster clinical decision-making.

He has also played a major leadership role in European research infrastructure through initiatives that connect photonics, AI and translational healthcare.

Alan Real

Organisation: Durham University
Role: Director of Advanced Research Computing
Based in: Durham, UK

Alan Real is focused on building the infrastructure that modern research increasingly depends on, ensuring universities can support the computational demands of AI, simulation and data-intensive science.

As Director of Advanced Research Computing at Durham University, Real leads strategy for research computing infrastructure supporting scientists across disciplines, from physics and engineering to life sciences and data analytics. His work centres on making advanced computing resources more accessible, scalable and aligned with changing academic research needs.

His career has been shaped by large-scale infrastructure delivery, operational leadership and enabling researchers to adopt complex computational tools. At Durham, he has helped expand capabilities in HPC, cloud services and research data infrastructure, areas increasingly important as AI places new demands on institutional systems.

He is also instrumental in research infrastructure collaborations, including the N8 Centre for Computationally Intensive Research and the EPSRC Tier-2 Bede supercomputing service.

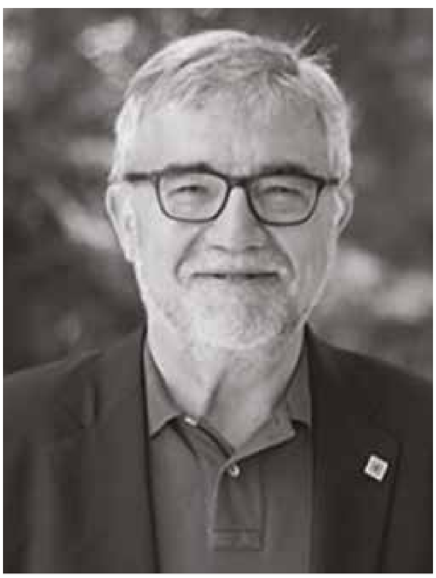

Mateo Valero

Organisation: Barcelona Supercomputing Center
Role: Director Emeritus
Based in: Barcelona, Spain

Mateo Valero has spent more than four decades shaping Europe's high-performance computing landscape, combining computer architecture research with institution-building of lasting impact.

As Director Emeritus of the Barcelona Supercomputing Center (BSC), Valero played a defining role in transforming it into one of Europe's leading supercomputing centres. It has become a major hub for processor design, exascale computing, AI infrastructure and computational science research.

Originally trained as a telecommunications engineer, Valero built his reputation through pioneering work in computer architecture, instruction scheduling, processor design and parallel computing. His research contributed to advances in vector architectures, superscalar processors and memory systems.

In 2005, he became the founding director of BSC, where he led major infrastructure projects, including the MareNostrum supercomputer family, supporting research across climate modelling, genomics, engineering, energy, and AI.

Valero has also played a important role in Europe's semiconductor ambitions through projects such as the European Processor Initiative, which aims to develop sovereign processor technologies for future HPC systems.

Amanda Randles

Organisation: Duke University
Role: Alfred Winborne Mordecai and Victoria Stover Mordecai Associate Professor
Based in: Durham, North Carolina, USA

Amanda Randles' research focuses on using high-performance computing to model the complexities of human blood flow. By combining advanced simulation with biological insight, her work provides new ways to study disease and improve understanding of cardiovascular health.

Dr Randles' career has been defined by a willingness to pivot across disciplines, moving from computing infrastructure into biology and medicine while consistently pushing the limits of what high-performance computing can achieve.

It is a trajectory marked by both adaptability and purpose. "Across all of these stages, high-performance computing has been the constant thread," she notes, reflecting on a career that spans physics, software engineering and computational biology.

Today, as a leading figure in scientific computing, Randles combines large-scale simulation, multiscale modelling and clinical insight to tackle some of medicine's most complex challenges.

Randles' research centres on developing advanced computational models of blood flow and cardiovascular systems. Her work aims to move beyond static representations of physiology, instead capturing how biological systems evolve over time. By integrating physics-based simulations with clinical data and wearable technologies, she is working to enable earlier disease detection and more proactive patient care.

Her path into this field was anything but linear. After graduating from Duke University in 2005 with a physics background, she joined IBM as a software engineer on the Blue Gene supercomputer team. Immersion in large-scale computing systems sparked a realisation: she was less interested in building machines than in using them to unlock new scientific insights. This led her to pursue a PhD in Applied Physics at Harvard University.

There, working with Efthimios Kaxiras and Hanspeter Pfister, she began developing large-scale blood flow simulations and initiated the HARVEY codebase. This marked a decisive shift into computational biology, with a focus on cardiovascular flow. As a Lawrence Fellow at Lawrence Livermore National Laboratory, with time spent at Dana-Farber Cancer Institute, her research expanded further to incorporate cancer biology and to model how cells interact within flowing systems.

Returning to Duke to establish her own lab, Randles continued to evolve her work. Her team has developed increasingly sophisticated models that incorporate cellular detail and, more recently, track physiological processes over extended periods.

This shift towards longitudinal modelling represents a significant departure from traditional snapshot-based simulations.

Today, her research focuses on building patient-specific "digital twins" of the cardiovascular system. Central to this is the Longitudinal Hemodynamic Mapping framework, which enables simulation of blood flow across thousands to millions of heartbeats. By analysing cumulative exposure and trajectories rather than isolated measurements, this approach offers a new way to assess disease risk in conditions such as heart failure and carotid disease.

Alongside this, Randles is advancing multiscale modelling techniques that bridge the gap between microscopic cellular interactions and whole-body systems. Her adaptive physics refinement approach allows simulations to capture phenomena at vastly different scales, from microns to metres, while maintaining computational efficiency. Achieving this has required innovations in load balancing, reproducibility and dynamic refinement on modern high-performance computing architectures.

Her work has consistently pushed technical boundaries. Early in her career, she demonstrated one of the first full-heartbeat simulations at cellular resolution, later scaling to whole-body vascular models and eventually simulating nearly a billion red blood cells on leadership-class systems. More recently, she has extended simulations from around 100 heartbeats to as many as 4.5 million, equivalent to several weeks of physiological time, an advance that opens entirely new avenues for studying disease progression.

Despite these achievements, Randles identifies validation as the field's greatest challenge. Many of the quantities her models seek to capture, such as detailed three-dimensional blood flow over long periods, cannot be directly measured in patients. Ensuring that simulations remain trustworthy, therefore, requires careful integration with clinical data and robust validation strategies, even when the ground truth is incomplete.

Looking ahead, her ambition is to demonstrate true longitudinal haemodynamic maps in real patients, capturing physiological changes over the course of a year. Such advances could enable entirely new biomarkers based on cumulative exposure, shifting healthcare from reactive treatment towards prediction and prevention.

Randles is also a strong advocate for early-career researchers. Taking risks, she says, is essential. "Many of the opportunities that shape your career will not feel like the safe or obvious next step, but if something excites you, it is worth pursuing." She emphasises the importance of working at the boundaries of disciplines, where new perspectives can lead to transformative breakthroughs.

Wojciech Turek

Organisation: University of Cambridge, UK
Role: Head of Research Computing Infrastructure and Platforms
Based in: United Kingdom

From building early supercomputers to developing national artificial intelligence (AI) infrastructure, Wojciech Turek's career reflects the rapid evolution of research computing. His work focuses on enabling access to advanced systems while addressing the growing complexity of data security and infrastructure design.

Turek has spent nearly two decades building the infrastructure that underpins modern scientific discovery, combining hands-on systems engineering with large-scale architectural leadership.

He began his career at the University of Cambridge in 2006 as part of the team that built Darwin, then one of the fastest supercomputers globally and the leading system in UK academia. That early experience established a practical, systems-level approach that has remained a defining feature of his career. Over time, he progressed from system administrator to architect, contributing to the development of 10 top-100 supercomputers.

Rather than a single defining transition, his career has been shaped by the expansion of his areas of responsibility. Moving into storage architecture led to a global No 1 ranking on the IO500 list in 2019, highlighting leadership in data-intensive computing. Later, the implementation of Cambridge's first ISO 27001-accredited Trusted Research Environment marked a shift toward secure computing and data governance. More recently, his involvement in co-designing the Dawn AI supercomputer, developed with industry partners, has positioned him firmly within the emerging field of national AI infrastructure.

As Head of Research Computing Infrastructure and Platforms at Cambridge, Turek now leads major national-scale initiatives. He is currently the technical lead and architect on two programmes representing significant UK government investment in AI capability.

These include the expansion of the Dawn AI supercomputer to incorporate next-generation GPU accelerators and the deployment of Zenith, a new national computational resource.

Together, these systems aim to strengthen sovereign AI and HPC capacity in the UK while operating within the constraints of an active research environment.

Alongside infrastructure development, his research focuses on one of the most pressing challenges in modern computing: enabling secure access to large-scale

systems for sensitive data analysis. Through projects such as FRIDGE and the UKRI-funded National Federated Compute Services NetworkPlus, he is developing approaches that allow Trusted Research Environments to extend securely into external HPC and AI platforms. These solutions include ephemeral secure enclaves and federated storage systems, enabling researchers to access national resources without compromising data governance or requiring complex new agreements.

The computing solutions underpinning his work span both traditional HPC and cloud-native technologies. Large-scale systems are built using GPU accelerators, high-performance storage such as Lustre, and workload management through Slurm, all orchestrated within private cloud environments using OpenStack. For secure and federated computing, he leverages Kubernetes-based platforms, Ceph storage and federated identity frameworks to enable scalable, encrypted and policy-compliant data access across institutions.

Turek identifies two major challenges shaping the field today. The first is bridging the gap between sensitive data governance and access to national AI infrastructure, a barrier that limits the potential of data-driven research in areas such as healthcare. The second is the physical constraint of infrastructure itself: modern AI systems demand unprecedented levels of power and cooling, placing increasing pressure on data centre capacity and sustainability.

Looking ahead, his goal is to establish Zenith as a leading national resource, broadening access to advanced computing for a wider research community.

A key milestone will be demonstrating federated Trusted Research Environment access at scale, enabling secure, routine use of national AI infrastructure for sensitive data applications.

His advice to early-career researchers reflects his own trajectory. Developing breadth across multiple domains, systems, storage, networking and security can be as valuable as deep specialisation. Maintaining curiosity and a willingness to engage directly with technology, even outside formal roles, creates opportunities to contribute at larger scales over time.

Through his work, Turek exemplifies the evolving role of scientific computing infrastructure: not only delivering performance, but enabling secure, collaborative and scalable research in an increasingly data-driven world.

Silvia Zorzetti

Organisation: Fermi National Accelerator Laboratory
Role: Deputy Division Head, NQI National Quantum Center – SQMS
Based in: United States

At the forefront of quantum information science, Silvia Zorzetti is developing new approaches to the design and use of quantum systems. Her work spans device innovation, system co-design and application development, connecting fundamental research with real computational challenges.

Zorzetti is working to address one of the defining challenges of quantum computing: scaling quantum systems beyond isolated devices into practical, networked technologies. Her research sits at the intersection of hardware, software and theory, where progress depends on tightly coordinated advances across the entire stack.

Her career has been shaped by a strong focus on co-design – bringing together physics, engineering and computation to push the limits of what quantum systems can achieve. Today, as Deputy Division Head of the National Quantum Initiative (NQI) National Quantum Center Division at Fermi National Accelerator Laboratory, she plays a leading role in one of the United States' flagship quantum research programmes.

Within the Superconducting Quantum Materials and Systems (SQMS) Centre, one of the US Department of Energy's national quantum information science research centres, Zorzetti leads the Architectures and Use-Cases Thrust. Her work focuses on advancing large-scale quantum devices based on the qudit (quantum dit) approach, which extends beyond conventional qubit systems by exploiting multiple energy levels within a single quantum element. This approach, enabled by superconducting cavities with exceptionally long coherence times, opens new possibilities for complex quantum simulations and applications.

A central element of her research is the development of quantum architectures that integrate hardware capabilities with computational requirements. This includes designing native gate sets tailored to specific devices, building calibration pipelines for parameter optimisation and noise characterisation, and developing compilers that translate high-level algorithms into hardware-executable instructions using modern frameworks such as LLVM and MLIR.

Alongside this, Zorzetti leads a research programme addressing a critical bottleneck in quantum computing: microwave-optical quantum transduction. This technology is essential for connecting quantum processors into larger networks. She has developed a novel device architecture that leverages long coherence times in superconducting cavities to significantly improve conversion efficiency, with projected performance far exceeding current state-of-the-art systems. This work has been recognised with a US Department of Energy Early Career Award.

Her contributions also include advances in qudit-based quantum processing units, demonstrating long coherence times and the ability to encode multiple energy levels within a single system. These developments provide a foundation for "bosonic" quantum computing approaches, which are particularly suited to problems such as quantum field theory simulations that are challenging for conventional qubit-based platforms.

Despite rapid progress, Zorzetti identifies the co-design process itself as one of the field's greatest challenges. Achieving consistent, repeatable improvements across theory, simulation, fabrication and experimental validation requires tightly integrated workflows and often bespoke tools. From multi-physics device modelling to advanced characterisation techniques such as Wigner tomography, many aspects of quantum system development remain highly specialised and resource-intensive.

She is focused on integrating artificial intelligence (AI) into the quantum computing workflow. This includes using AI to optimise device calibration, automate control systems and manage resources across increasingly complex and heterogeneous quantum platforms. Such approaches could play a key role in making quantum systems more scalable and efficient.

Her near-term goals include demonstrating high-efficiency microwave-optical quantum transduction at the single-photon level, an important step towards distributed quantum computing, and achieving a qudit-based quantum advantage for scientifically meaningful problems, such as quantum Fourier transformation in lattice field theory.

Zorzetti is deeply involved in community building and workforce development. Through summer schools, internships and lecture series, she has contributed to training the next generation of quantum scientists and engineers, engaging participants across a wide range of career stages.

Her advice to early-career researchers reflects the multidisciplinary nature of the field. Developing deep expertise is essential, but so is maintaining a broad perspective across disciplines. Curiosity, openness and a willingness to engage with unanswered questions are key to advancing in a rapidly evolving area such as quantum computing.

At its core, Zorzetti's work illustrates the importance of integration in scientific computing. By connecting advances in hardware, software and theory, she is helping to lay the foundations for scalable quantum systems capable of addressing problems beyond the reach of classical computation.

The new scientific stack

Charunethran Panchalam Govindarajan discusses how HPC, AI and quantum workflows are converging around scientific discovery

Charunethran Panchalam Govindarajan

The convergence of HPC and AI is an architectural change. AI does not reduce the importance of physics-based modelling; in many cases, it increases the value of simulation by making its outputs reusable across a larger workflow. A 2023 *Nature* review documented this shift across hypothesis generation, experimental design and simulation interpretation.[1]

Building the adaptive scientific workflow

The scientific stack that supports this work has to behave less like a queue of isolated jobs and more like an adaptive system. Researchers need to move between simulation, data processing, model development, analysis and validation without rebuilding their environment at every step. This requires processor performance, memory bandwidth, low-latency networking, orchestration, governance and AI tooling to work together. Cloud-based HPC gives teams flexibility to align infrastructure with workloads and reduce friction between stages.

AWS and the infrastructure for scientific discovery

At AWS, this direction is reflected in an advanced computing portfolio built for scientific work. For HPC users, the foundation is Amazon Elastic Compute Cloud (Amazon EC2), which includes purpose-built HPC instances for workloads that depend on throughput, memory bandwidth and efficient communication across many nodes, most recently with Amazon EC2 Hpc8a instances powered by 5th Generation AMD EPYC processors.[2]

The significance is not any single capability, but the range of infrastructure that scientific discovery now requires. CPU-based HPC instances support tightly coupled simulation. GPU and accelerated compute infrastructure support AI training, inference and visualisation. Orchestration services connect these stages so simulation can generate data, AI can explore a wider design space, and researchers can return to higher-fidelity methods with better questions.

This pattern is visible in engineering workflows on AWS. In one automotive example, AWS Batch launched computational fluid dynamics simulations that created ground truth data, Amazon SageMaker-trained drag prediction models, and GPU-based Amazon EC2 infrastructure served models through an interactive design workflow.[3]

AWS Parallel Computing Service extends this workflow approach for organisations that use Slurm. In chemical formulation and discovery, Aionics migrated to AWS Parallel Computing Service in 2025 to support computational chemistry workloads, including density functional theory simulations and GPU-accelerated molecular dynamics with machine-learned interatomic potentials.[4] This example shows the direction of the stack: not just more compute, but a managed HPC environment that helps researchers focus on scientific search rather than infrastructure maintenance.

This extends beyond engineering. The Allen Institute has used HPC and generative AI on AWS for neuroscience research, including omics pipeline runs that moved from weeks on premises to as little as one day on AWS.[5] The importance is not only an isolated performance gain. It is the ability to change the pace at which teams move from question to computation, and from computation to insight.

Quantum and the next stage of scientific computing

Quantum computing belongs in this same stack, not as a replacement for HPC or AI, but as a future accelerator for selected problems in chemistry, materials science, optimisation and quantum systems. Amazon Braket gives researchers access to quantum computers and simulators to explore quantum and hybrid algorithms alongside classical workflows. Recent work with Kvantify Qrunch on Amazon Braket demonstrates how quantum chemistry calculations can be developed using hybrid quantum algorithms on today's quantum devices.[6]

AWS is also investing in the longer-term hardware path. In February 2025, Amazon announced Ocelot, a quantum computing chip developed by the AWS Center for Quantum Computing. Ocelot uses a hardware-efficient approach to quantum error correction based on cat qubits, aiming to reduce the overhead of fault-tolerant quantum computing.[7]

Much of today's AI conversation is focused on productivity, automation and consumer applications. The deeper opportunity is to expand the rate at which researchers can explore, test and validate ideas. When HPC provides the physical and numerical foundation, AI helps navigate vast search spaces, and quantum computing begins to extend what can be modelled, the effect is not simply faster computing. It is a different cadence of discovery, one that will shape progress in energy, medicine, materials and fundamental science.

Charunethran Panchalam Govindarajan heads Product Marketing for Advanced Computing at AWS, spanning high-performance computing, AI for science and quantum computing. He holds a Master's degree in Electrical Engineering from Stanford University.

References

1.Hanchen Wang et al., "Scientific discovery in the age of artificial intelligence," Nature, 2023. https://www.nature.com/articles/s41586-023-06221-2

2. "A technical deep dive into Amazon EC2 Hpc8a performance for engineering and scientific workloads," AWS HPC Blog, March 2026. https://aws.amazon.com/blogs/hpc/a-technical-deep-dive-into-amazon-ec2-hpc8a-performance-for-engineering-and-scientific-workloads/

3. "Using machine learning to drive faster automotive design cycles," AWS HPC Blog, May 2024. https://aws.amazon.com/blogs/hpc/using-machine-learning-to-drive-faster-automotive-design-cycles/

4. "How Aionics accelerates chemical formulation and discovery with AWS Parallel Computing Service," AWS HPC Blog, November 2025. https://aws.amazon.com/blogs/hpc/how-aionics-accelerates-chemical-formulation-and-discovery-with-aws-parallel-computing-service/

5. "Powering discovery: Allen Institute advances in-silico science using AWS," AWS Case Study. https://aws.amazon.com/solutions/case-studies/allen-institute-case-study/

6. "Building advanced quantum chemistry calculations with Kvantify Qrunch on Amazon Braket," AWS Quantum Technologies Blog, February 2026. https://aws.amazon.com/blogs/quantum-computing/building-advanced-quantum-chemistry-calculations-with-kvantify-qrunch-on-amazon-braket/

7. "Amazon announces Ocelot quantum chip," Amazon Science, February 2025. https://www.amazon.science/blog/amazon-announces-ocelot-quantum-chip

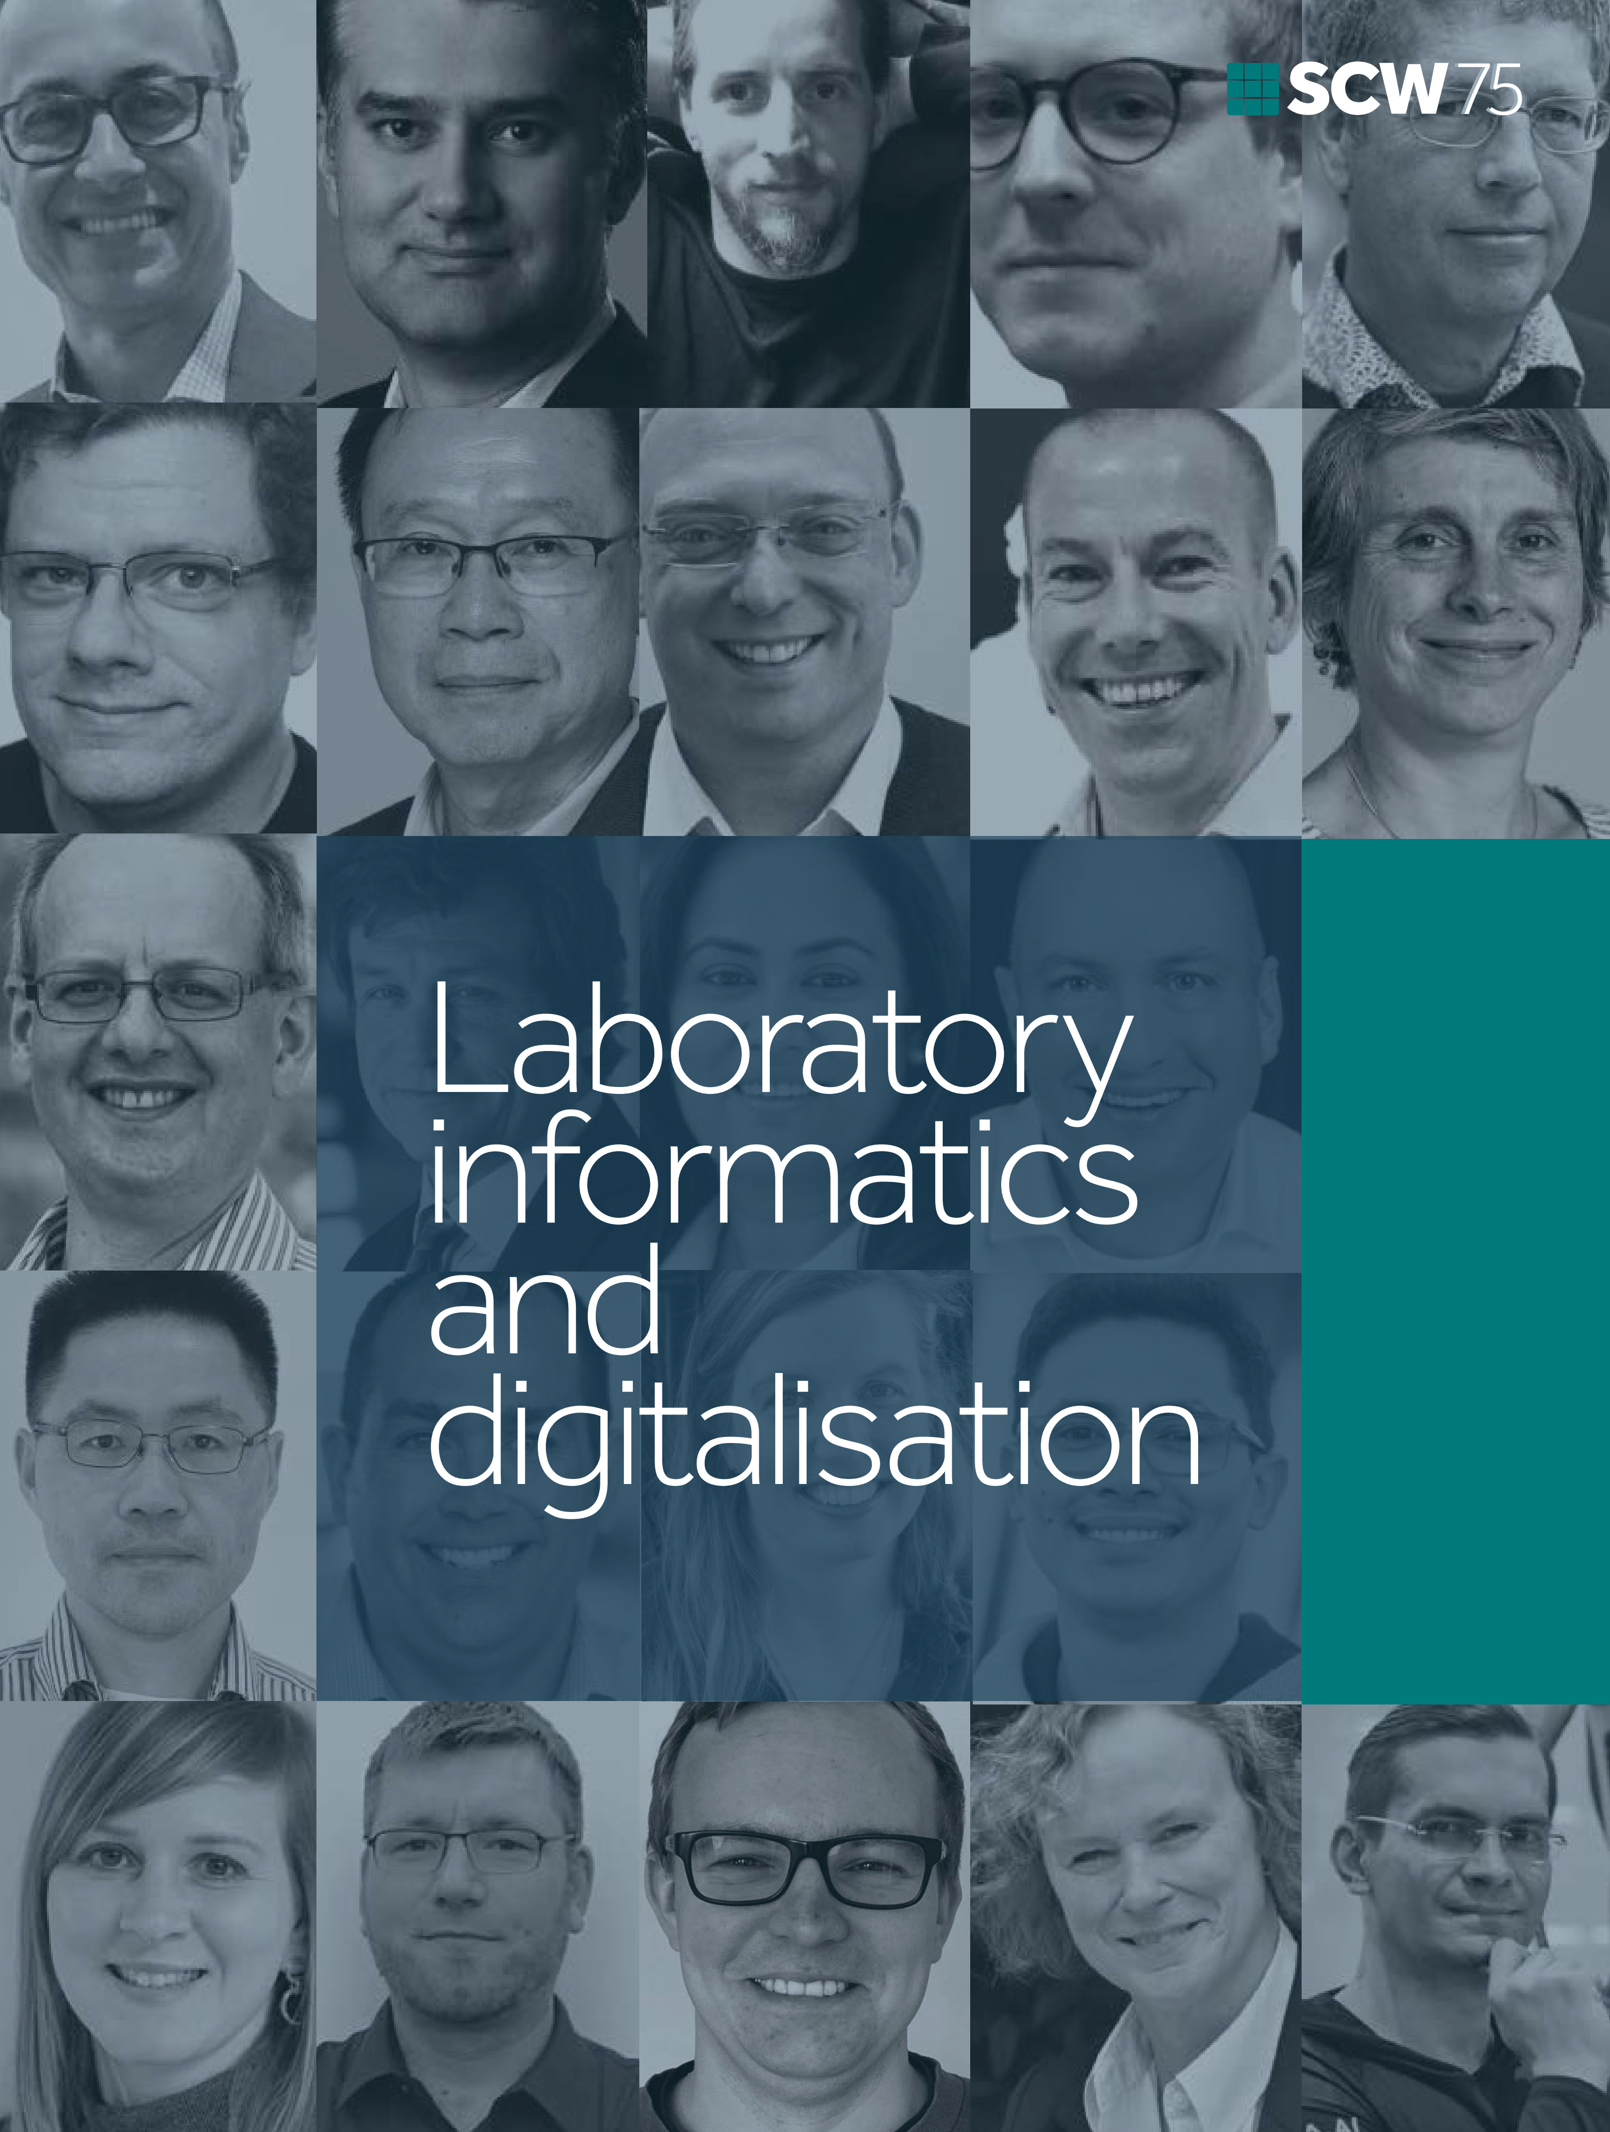

Laboratory informatics and digitalisation

Jagat Adhiya

Organisation: Bayer
Role: Vice-President, Head of IT & Digital for Pharma R&D and Cell & Gene Therapies
Based in: United States

Jagat Adhiya is engaged in work on the digital infrastructure needed to modernise pharmaceutical research and development, with a focus on connecting data, platforms and scientific workflows across complex R&D organisations.

As Vice-President and Head of IT & Digital for Pharma R&D and Cell & Gene Therapies at Bayer, his work supports the computational foundations required for AI-enabled discovery, translational research and advanced therapeutic development.

In this environment, scientific computing is not only about analytics, but about building the systems that allow data to be captured, integrated, governed and reused across the R&D lifecycle.

Before joining Bayer, Adhiya held senior roles at Merck Healthcare, including Head of R&D Data Strategy and Innovation, as well as earlier roles at EMD Millipore and Waters. That experience gives him a broad view of how digital infrastructure supports laboratory science, development pipelines and enterprise-scale research.

His career reflects a wider shift in life sciences, where IT, data strategy and computational platforms have become central to faster, more connected and more predictive drug development.

Mohit Agnihotri

Organisation: AbbVie
Role: Senior Director, Development Sciences Data & Digital Strategy
Based in: United States

Mohit Agnihotri works on the data and digital strategy that supports modern pharmaceutical development, helping embed computational thinking into the Development Sciences function within AbbVie's R&D organisation.

As Senior Director, Development Sciences Data & Digital Strategy, his work focuses on connecting scientific workflows with the data architecture, analytics platforms and digital capabilities needed to improve decision-making. Development sciences sits at a critical stage of the R&D pipeline, where experimental data, translational insight and operational execution must come together to support candidate progression and regulatory readiness.

Agnihotri's role is not only about deploying technology, but also about building the organisational capability required to use it effectively. By aligning data strategy with scientific practice, he helps create digital foundations that make R&D processes more connected.

His work reflects a broader shift in pharmaceutical development, where data platforms, workflow integration and analytics are becoming essential infrastructure for faster, more efficient and more predictive drug development.

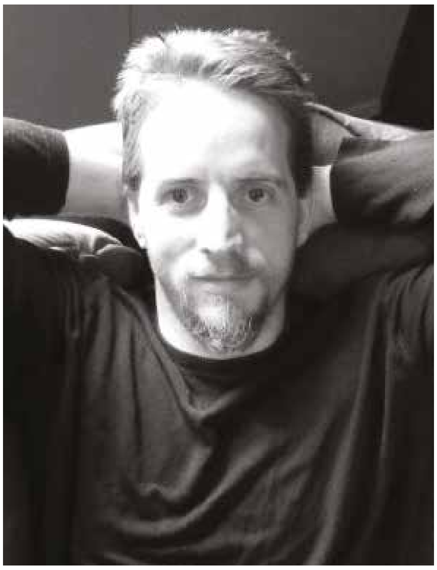

Ben Allen

Organisation: Moa Technology
Role: Head of Computational Chemistry
Based in: UK

Ben Allen works in computational chemistry for agricultural technology, applying cheminformatics and data-driven approaches to discover new modes of action.

As Head of Computational Chemistry at Moa Technology, he supports the search for novel herbicides to address difficult weeds and resistance challenges. Crop protection requires researchers to explore large chemical and biological spaces, identify promising compounds and understand how they interact with biological systems. Computational chemistry helps make that process faster and more focused by prioritising candidates before extensive experimental testing.

Allen's role brings together chemical expertise, informatics and data infrastructure to improve how scientific data is stored, managed and interrogated.

This is increasingly important as agricultural R&D becomes more dependent on high-throughput screening, complex datasets and predictive modelling.

His work reflects the growing importance of computational chemistry beyond pharmaceutical drug discovery. By applying data-driven methods to crop protection, Allen contributes to innovation in sustainable agriculture and supports the development of new approaches to weed control.

Bruno Betoni Parodi

Organisation: BASF
Role: Scientific Director
Based in: Germany

Bruno Betoni Parodi applies robotics, automation, control, machine-learning and artificial intelligence to industrial research challenges, with a focus on turning computational methods into practical tools for complex technical environments.

As Scientific Director at BASF, his work reflects the growing role of computational engineering in industrial R&D. Modelling, sensing and data-driven methods are increasingly used to improve process understanding, automate workflows and support better decision-making across chemical production, materials development and agricultural science.

His research background includes adaptive localisation, feature-map learning and machine-learning approaches for plant phenotyping, showing the breadth of his work across robotics, automation and applied AI.

These methods are important in industrial settings where systems must operate reliably, interpret complex data and adapt to changing conditions.

Parodi's career brings together academic research and industrial application.

By translating methods from robotics and machine-learning into usable engineering capabilities, he supports the development of smarter, more automated and more data-driven approaches to scientific and industrial innovation.

John Chan

Organisation: Novartis
Role: Global Head of Digital, Informatics and AI, Biomedical Research
Based in: United States

John Chan works on applying digital technologies, informatics and AI to biomedical research, helping large R&D organisations use data more effectively across discovery and development.

As Global Head of Digital, Informatics and AI, Biomedical Research at Novartis, and formerly Head of the ShinrAI Center for Artificial Intelligence and Machine-Learning, his work focuses on building the computational and organisational capabilities needed to apply AI at scale.

In biomedical research, the challenge is not simply creating models, but connecting high-quality data, scientific expertise and decision-making across complex discovery workflows.

Chan's role reflects the growing importance of AI-enabled research, where machine-learning, informatics platforms and digital infrastructure are becoming core scientific capabilities.

These tools can help researchers identify patterns, prioritise hypotheses and accelerate the translation of biological insight into potential therapies.

By leading digital and AI functions at Novartis, Chan helps shape how computational methods are embedded into discovery science and translational research.

Marek Dynowski

Organisation: Cancer Research UK Manchester Institute
Role: Head of Scientific Computing
Based in: Manchester, UK

Marek Dynowski is helping reshape how cancer researchers use large-scale computing, not by building faster machines alone, but by making advanced infrastructure easier for scientists to access and use effectively.

As Head of Scientific Computing at Cancer Research UK Manchester Institute, Dynowski works across high-performance computing, data infrastructure and biomedical research. His team supports the computational demands of modern cancer science, where genomics, imaging and biomarker research generate large datasets that require secure storage, rapid processing and advanced analysis.

A major milestone in his recent work was overseeing the design and construction of a new data centre within the Paterson Building, a major research hub that brings together scientists from multiple organisations.

The facility was built to support both day-to-day computing needs and advanced scientific workloads, while enabling secure collaboration across research teams.

He has also led efforts to make complex systems accessible through platforms such as RStudio Server, Galaxy and remote desktop tools. Alongside infrastructure development, Dynowski focuses on reproducibility, cybersecurity and researcher training.

Morgan Harris

Organisation: Pfizer
Role: MSAT CFD Group Lead
Based in: United States

Morgan Harris works on computational fluid dynamics for pharmaceutical manufacturing, using simulation to improve process understanding, scale-up and development.

As MSAT CFD Group Lead at Pfizer, his work supports the use of modelling to analyse manufacturing processes before teams commit to expensive experiments or production changes.

In pharmaceutical development, CFD can help researchers understand mixing, vessel behaviour, fluid flow and process performance, all of which are critical to robust manufacturing.

Harris's role reflects the growing importance of computational engineering in pharmaceutical process development. Simulation allows teams to compare design options, test assumptions and build confidence in process decisions while reducing trial-and-error experimentation.

In regulated environments, these models must be carefully connected to physical evidence, process knowledge and engineering judgement.

Harris's work shows how simulation is becoming part of the scientific toolkit for pharmaceutical manufacturing, supporting faster iteration, improved process confidence and more data-driven decision-making across development and technology transfer.

Nils Hoffmann

Organisation: Beiersdorf
Role: Director Future Testing and Method Development
Based in: Germany

Nils Hoffmann works on future testing and method development in consumer healthcare and personal care research and development, where product innovation increasingly depends on better predictive tools and more relevant evidence.

As Director Future Testing and Method Development at Beiersdorf, his work focuses on improving how products are evaluated, validated and brought closer to consumer needs.

In personal care, testing methods must balance scientific rigour, product performance, safety, sustainability and user experience, while reducing reliance on slow or resource-intensive approaches.

Hoffmann's background in R&D leadership and method development reflects the growing need to connect scientific testing, innovation management and digital capability. Computational and data-informed methods are becoming increasingly important in this space, helping teams interpret results faster, compare formulations and design more efficient development workflows.

His role highlights how consumer product research is becoming more predictive and method-driven. By advancing future testing approaches, Hoffmann supports faster, more reliable and more sustainable innovation in personal care science.

Lovisa Holmberg-Schiavone

Organisation: AstraZeneca
Role: Director Protein Science, Discovery Biology, Discovery Sciences, R&D
Based in: Sweden

Lovisa Holmberg-Schiavone works in protein science for drug discovery, supporting the production and characterisation of biological reagents used in early research.

As Director Protein Science, Discovery Biology, Discovery Sciences, R&D at AstraZeneca, her work supports teams that depend on high-quality protein reagents, structured workflows and reliable biological data. Protein science is increasingly central to modern discovery, providing the materials needed for assay development, structural biology, biophysics and therapeutic research.

Her role sits in an area where experimental biology and computational methods increasingly depend on one another. Protein science teams generate complex data around expression, purification, construct design, reagent quality and assay readiness.

Making that information structured, reusable and connected to discovery decisions is essential for improving speed and reproducibility.

Holmberg-Schiavone's work reflects the practical side of computationally enabled biology, ensuring that high-quality reagents, data and workflows are available to support drug discovery projects and collaborative research.

Jonathan Hirst

Organisation: University of Nottingham
Role: Professor of Computational Chemistry
Based in: Nottingham, UK

Lisa Gilligan-Lee

As the chemical sciences increasingly adopt artificial intelligence (AI), Jonathan Hirst is working to ensure those advances also support sustainability. His research combines machine-learning, molecular modelling and data-driven tools to guide more efficient and environmentally conscious chemical design.

From early work on neural networks to current efforts in AI-driven sustainable chemistry, his research reflects a consistent focus on using computation to transform how chemical science is practised.

Beginning with a degree in Chemistry at the University of Oxford, followed by a PhD in London on neural networks for modelling molecular structure and function, Hirst was working with machine learning long before it became prominent. Early post-doctoral research at Carnegie Mellon University expanded his expertise in molecular dynamics and quantum chemistry, with a focus on protein folding and spectroscopy.

After a period as an Assistant Professor at the US-based Scripps Research Institute, he returned to the UK and joined the University of Nottingham, where he has since built a wide-ranging research programme spanning computational chemistry, machine-learning and molecular design. He was promoted to Professor in 2004 and later served as Head of School, during which time he led the development of the GSK Carbon Neutral Laboratory, a major investment in sustainable chemistry infrastructure.

Today, Hirst's work focuses on integrating AI into chemical research and development, with a strong emphasis on sustainability. Through his Royal Academy of Engineering Chair in Emerging Technologies, he leads efforts to develop tools that enable chemists to combine domain expertise with advanced analytics, supporting more efficient, safer and environmentally conscious research.

Central to this vision is AI4Green, an open-source platform that supports sustainable chemical research. Built around FAIR data principles, the platform integrates data management with machine-learning tools to guide decision-making in areas such as solvent selection, synthetic route design and life cycle assessment. By combining a Python-based backend with an accessible web interface, AI4Green enables chemists to work interactively with advanced computational methods without requiring specialist programming knowledge.

The platform has gained traction across both academia and industry, reflecting growing demand for tools that align chemical innovation with sustainability goals. Features such as interactive principal component analysis for solvent selection and AI-enhanced retrosynthesis workflows demonstrate how computational approaches can directly influence practical decision-making in the laboratory.

Despite these advances, Hirst identifies a fundamental challenge in the field: data accessibility. While other disciplines have benefited from large-scale open data initiatives, chemistry still lags behind. Much valuable data remains locked within laboratory notebooks, limiting the potential of machine learning approaches. He points to the success of resources such as the Protein Data Bank as evidence of what can be achieved through open data sharing, and sees significant untapped opportunity within the chemistry community.

Looking ahead, Hirst is exploring the intersection of machine-learning and quantum computing, aiming to demonstrate practical applications in real-world chemical problems. This emerging area has the potential to further transform molecular design and drug discovery, particularly when combined with advances in computational infrastructure.

Beyond AI4Green, his research extends into drug design, polymer development and industrial collaboration.

Partnerships with organisations such as GSK and innovative small-to-medium enterprises (SMEs) have led to the development of computational tools for designing drug-like molecules, improving biopharmaceutical production and creating more sustainable materials. His work also contributes to international efforts in quantum-enabled drug discovery, collaborating with companies such as Phasecraft and QuEra Computing.

For early-career researchers, Hirst's advice is grounded in experience. Following one's own interests, engaging with the wider research community and maintaining an open mind are essential. Interdisciplinary collaboration, he notes, can be particularly powerful when driven by shared motivation and purpose.

At its core, his work reflects a broader transformation in chemistry, one in which computation, data and sustainability are becoming central to discovery. By developing tools that integrate these elements into everyday practice, Hirst is helping to shape a more open, efficient and responsible future for chemical research.

Howard Jacob

Organisation: AbbVie
Role: Vice-President and Head of Genomics and Data Integration
Based in: United States

Howard Jacob works on genomics, data integration and quantitative approaches to drug discovery, helping connect large-scale biological data with research decision-making.

As Vice-President and Head of Genomics and Data Integration at AbbVie, his work focuses on bringing genomic insight into pharmaceutical R&D. Modern drug discovery increasingly depends on the ability to integrate complex datasets from genomics, human biology, disease research and clinical studies.

Such data can help identify targets, understand mechanisms and support patient stratification, but only when they are connected, analysed and interpreted at scale.

Jacob's role reflects a central challenge in computational biology: turning large, heterogeneous datasets into knowledge that can guide research strategy and improve the probability of success in drug development.

His work also connects with the growing use of AI and machine-learning in pharmaceutical research, where integrated data can help reveal patterns, prioritise opportunities and accelerate the development of potential medicines.

Ishrat Jahan

Organisation: Regeneron
Role: Lead IT Engineer
Based in: United States

Ishrat Jahan works on the IT foundations that support digital transformation in life sciences laboratories, helping ensure scientific workflows can operate reliably in modern research environments.

As Lead IT Engineer at Regeneron, her work supports laboratory computing systems and the technology requirements of pharmaceutical research. In modern biopharma, laboratory environments increasingly depend on connected instruments, automation, data capture, compliance systems and collaboration tools. These systems must be reliable enough for day-to-day science while flexible enough to support digital transformation.

Jahan's role reflects a practical, but essential, part of computational engineering, ensuring that laboratory systems, digital tools and scientific workflows work together effectively for researchers. This requires understanding both technology infrastructure and the realities of laboratory practice.

Her work demonstrates that scientific computing depends not only on algorithms and large-scale infrastructure, but also on engineers who can connect science and IT, supporting productive partnerships between R&D teams and technology organisations.

Ross King

Organisation: University of Cambridge/ Chalmers University of Technology
Role: Professor
Based: UK and Sweden

Ross King is a leading researcher in artificial intelligence for scientific discovery, known for pioneering work on automating research itself and has worked in the field of machine learning for more than 35 years.

As a Professor at the University of Cambridge and Chalmers University of Technology, King's work focuses on using AI, robotics and computational methods to make scientific discovery more systematic and reproducible. He originated the idea of the "Robot Scientist", a computer and robotic system able to generate hypotheses, design experiments, run them physically, interpret results and repeat the cycle.

This work has helped define a major strand of computational science – closed-loop discovery, where AI and laboratory automation are integrated into the full scientific process. Rather than using computation only to analyse data after experiments, King's approach embeds computation into hypothesis generation and experimental design.

He is also involved in organising the international 'Nobel Turing Grand Challenge' to develop AI scientists: "AI systems capable of making Nobel-quality scientific discoveries highly autonomously at a level comparable, and possibly superior, to the best human scientists by 2050".

His work is becoming increasingly relevant as more laboratories adopt automation, AI planning and self-driving experimental systems to improve discovery speed, reduce manual bottlenecks and strengthen reproducibility.

Subha Madhavan

Organisation: Pfizer
Role: Head of AI/ML, Quantitative & Digital Sciences
Based in: United States

Turning complex biomedical data into actionable insight remains one of the most significant challenges in modern drug development. Subha Madhavan works at the crossroads between biomedical science and artificial intelligence (AI), where data-driven methods are reshaping how therapies are developed. Her focus is on using AI and machine-learning to accelerate clinical research and improve patient outcomes.

Her career has been defined by a consistent focus on integrating science, data and technology.

Trained in molecular biology, with a PhD followed by post-doctoral research at Johns Hopkins University, she moved early into roles that combined biomedical research with data-driven approaches. Positions at Georgetown University, the National Cancer Institute, the US Food and Drug Administration, and AstraZeneca allowed her to shape translational research and regulatory science initiatives, building a foundation at the interface of clinical science and computation.

A key transition came when she moved into enterprise AI leadership at Pfizer. There, she leads the AI/ML, Quantitative & Digital Sciences (AQDS) team, focusing on accelerating decision-making across drug development. Her work spans clinical, translational and regulatory domains, applying data-driven methods to improve efficiency, reduce timelines and enhance the quality of evidence used to guide therapeutic development.

Central to her current work is the integration of AI across the drug development lifecycle. This includes applying machine-learning and generative AI to optimise trial design, streamline workflows and generate faster insights from increasingly complex datasets. She is also driving digital health initiatives that incorporate wearable technologies and real-world data, enabling a more continuous and patient-centred view of therapeutic impact.

Technically, this work relies on a broad ecosystem of scientific computing tools. Cloud platforms such as Google Cloud, AWS and Azure provide infrastructure for large-scale data integration and model deployment, while frameworks such as TensorFlow and PyTorch support predictive and generative modelling. Knowledge graphs and semantic data integration approaches are used to connect pre-clinical, clinical and real-world evidence, enabling more comprehensive analysis of patient populations and treatment outcomes.

In addition, specialised tools support clinical and regulatory workflows, including trial emulation, protocol simulation and automated document generation. Digital health data from wearables is incorporated through signal processing and longitudinal analysis pipelines, while large language models are increasingly used to summarise evidence, automate documentation and support reasoning across complex datasets. Together, these technologies enable more efficient, data-driven decision-making throughout the entire development process.

Despite these advances, a fundamental challenge remains: data. Access to high-quality, integrated multimodal datasets spanning clinical trials, genomics, wearables and real-world evidence remains limited. Data is often fragmented across systems, stored in incompatible formats and constrained by privacy and regulatory requirements. Even when accessible, significant effort is required to standardise, harmonise and curate it before it becomes suitable for AI-driven analysis. Overcoming these barriers is essential to unlocking the full potential of AI in drug development.

Looking ahead, Madhavan's goal is to fully operationalise AI-driven decision-making across the lifecycle of therapeutic development. This includes building platforms that integrate large-scale datasets into AI-ready systems, enabling predictive modelling, optimised trial design and more patient-centred insights. A key milestone will be the deployment of generative AI and knowledge graph-based reasoning systems that can accelerate hypothesis generation, protocol development and regulatory decision support.

Her advice to early-career researchers reflects the interdisciplinary nature of this work. Deep expertise in a core domain remains essential, but it must be complemented by fluency in AI, data integration and digital technologies. Equally important is the ability to translate technical insight into strategic decisions, bridging the gap between analysis and real-world impact.

Beyond her own work, Madhavan is committed to building the next generation of scientific leaders. Through mentoring, advisory roles and engagement with start-ups, she supports emerging talent at the intersection of science, AI and translational medicine. This broader contribution reflects a view that progress in scientific computing is not only driven by technology, but by the people and ecosystems that enable it.

Ilja Kuesters

Organisation: Generate:Biomedicines
Role: Associate Director, Assay Automation & Qualification
Based in: United States

Ilja Kuesters works on laboratory automation for advanced biological research, helping turn complex assays into scalable, traceable and reliable workflows.

As Associate Director, Assay Automation & Qualification at Generate:Biomedicines, his work supports the automation infrastructure needed for data-rich therapeutic discovery. In a company applying generative biology to medicine, experimental platforms must generate high-quality data at scale while remaining robust enough to support scientific decisions.

Kuesters' work focuses on connecting assays, instruments, quality processes and data systems so research workflows can operate more efficiently.

The challenge is not simply automating individual laboratory steps, but ensuring that automated systems produce reproducible results and integrate effectively with broader discovery platforms.

His role reflects the increasing importance of computational engineering in the laboratory of the future. As biology becomes more automated and data-driven, specialists such as Kuesters help ensure that experimental systems, digital infrastructure and scientific goals are designed as connected workflows.

Zhongwu Lai

Organisation: AstraZeneca
Role: Executive Director, Bioinformatics Data Science
Based in: United States

Zhongwu Lai works in oncology bioinformatics and data science for drug development, helping translate complex molecular datasets into insight for cancer research.

As Executive Director, Bioinformatics Data Science at AstraZeneca, his work supports the use of genomics, next-generation sequencing and computational biology in oncology R&D. Cancer research generates vast and diverse datasets, from tumour genomics and transcriptomics to clinical and translational evidence. These data can help researchers understand tumour biology, identify biomarkers and support patient stratification, but only when underpinned by robust bioinformatics and data science methods.

Lai's work reflects the increasingly central role of computational biology in oncology. Bioinformatics has moved from a specialist analytical service to a strategic capability that can shape discovery, development and clinical decision-making.

By connecting genomics, data science and translational research, Lai contributes to more precise and evidence-driven oncology programmes, supporting the development of therapies informed by biological and clinical data.

Jesse Mulcahy

Organisation: Eli Lilly and Company
Role: Genetic Medicine Lab Automation Lead
Based in: United States

Jesse Mulcahy works on laboratory automation for genetic medicine, helping build scalable systems that support RNA biology, chemistry and therapeutic discovery.

As Genetic Medicine Lab Automation Lead at Eli Lilly and Company, his work supports automated workflows for one of the most rapidly developing areas of pharmaceutical research and development. Genetic medicine requires highly reproducible, data-connected laboratory systems that can handle complex biology while scaling across projects, teams and sites.

Automation can reduce manual bottlenecks, improve traceability and help researchers generate more consistent experimental data.

In genetic medicine, where workflows may span nucleic acid technologies, analytical methods and biological testing, these capabilities are essential for accelerating discovery and development.

Mulcahy's career also includes leadership in automation at Cellino, reflecting a broader focus on connecting robotics, laboratory science and data infrastructure.

His work demonstrates the growing demand for engineers who can design scalable, reliable laboratory systems closely aligned with scientific needs.

April Williams Wehner

Organisation: The Salk Institute
Role: Director, Integrative Genomics and Bioinformatics Core
Based in: United States

April Williams Wehner leads bioinformatics support for high-throughput biological research, helping scientists turn complex genomic data into interpretable biological insight. She is the Director of the Integrative Genomics and Bioinformatics Core at the Salk Institute, which she joined in April 2018.

She supports researchers working with large and diverse genomics datasets, providing the computational expertise needed to analyse, visualise and interpret results from modern high-throughput experiments.

The Salk Institute states that she obtained a PhD in Quantitative & Computational Biology and an MA in Molecular Biology from Princeton University in 2015. Her graduate work included research on transcriptional control in the malaria parasite *P. falciparum* and genome-wide studies of ageing mechanisms in *C. elegans*.

As director of a core facility, Wehner's role combines scientific computing, service leadership and collaboration. Bioinformatics cores are increasingly essential to modern biology, providing reproducible workflows, analytical standards and data interpretation for genomics-driven research.

Her work reflects how computational biology has become shared research infrastructure for the life sciences.

Iguaracy Pinheiro-de-Sousa

Organisation: EMBL-EBI
Role: Postdoctoral Fellow
Based in: Cambridge, UK

As biological data becomes increasingly complex, turning it into meaningful insight remains a major challenge. Iguaracy Pinheiro-de-Sousa is addressing this by developing computational models that capture how cells interact within tissues, helping to bridge the gap between large-scale data and experimental biology.

Dr Pinheiro-de-Sousa's work focuses on understanding how cells communicate within tissues, and how these interactions can be modelled to improve the design of organoids, laboratory-grown systems that mimic human organs. His research aims to move beyond simplified models towards more realistic representations of human tissue, particularly in the heart.

His journey began in biomedical science, with an early interest in human genetics. During his Master's in Genetics and Molecular Biology, he worked on pharmacogenetics in a hypertensive cohort from the Amazon region, examining how genetic background influences treatment response. This marked his first major pivot into cardiovascular research.

He continued along these lines during his PhD at the Heart Institute of the University of São Paulo, where he studied cardiovascular disease, endothelial cells, and blood flow. While this work provided a strong foundation in disease biology, it also highlighted the limitations of traditional experimental approaches in addressing complex, system-level questions.

A second key transition came when he joined the Novartis Next Generation Scientist programme in Basel, where he conducted RNA-seq and ChIP-seq analyses. This experience provided his first substantial engagement with computational research, enabling a shift into bioinformatics.

Now a Postdoctoral Fellow at EMBL-EBI, Pinheiro-de-Sousa combines single-cell transcriptomics with computational modelling to study tissue organisation. His work centres on cardiac organoids, where current models often lack the complexity and maturity of native tissue. By modelling cell-cell communication across developmental trajectories, his research identifies missing signalling interactions that can be tested experimentally to improve organoid design.

A key component of this work is the development of computational frameworks that translate large-scale data into actionable biological insight. By integrating transcriptomic datasets and modelling signalling networks, his team can generate hypotheses more efficiently and systematically than would be possible through manual analysis alone.


Jeff Dowling/EMBL-EBI

This enables a tighter feedback loop between computation and experiment, improving both the scalability and reproducibility of research.

Alongside this, he is developing a complementary line of research on vascular "addresses", molecular signatures that could enable targeted delivery within tissues. This work is now moving into validation, testing whether computational predictions can accurately reflect biological reality.

Despite rapid advances in data generation and modelling, Pinheiro-de-Sousa identifies a central challenge: bridging the gap between computational prediction and biological validation. While large datasets offer unprecedented opportunities, translating them into meaningful and testable hypotheses remains difficult. Effective collaboration between computational and experimental scientists is therefore essential, requiring not only technical expertise, but also clear communication and shared understanding.

Looking ahead, his goal is to establish a robust feedback loop between modelling and experimentation. By rigorously validating computational predictions, he aims to refine models and extend their application to disease contexts, including cardiomyopathies and cardiac fibrosis. Achieving this would support the development of more realistic disease organoids, with potential implications for drug discovery and therapeutic testing.

For early-career researchers, his advice reflects his own interdisciplinary path. Curiosity and openness to collaboration are key, particularly at the boundaries between fields. Some of the most valuable ideas, he notes, arise from unexpected interactions and perspectives. Building a strong foundation in one area is important, but so too is the ability to communicate and work across disciplines.

At the core of his work is a commitment to integration. Through collaborations spanning institutions and domains – from cardiovascular research in São Paulo to organoid modelling in Europe and beyond – Pinheiro-de-Sousa is helping to connect computational and experimental science. In doing so, he is contributing to a broader shift towards data-driven approaches that can better capture the complexity of human biology.

Samantha Pearman-Kanza

Organisation: University of Southampton
Role: Principal Enterprise Fellow
Based in: UK

Samantha Pearman-Kanza focuses on a deceptively simple question: how do scientists actually record and share what they do? Her work sits at the intersection of computing, chemistry and human behaviour, addressing one of the most persistent challenges in modern research, capturing data in a way that is reproducible, reusable and truly meaningful.

Her career began in computer science, with an MEng from the University of Southampton, followed by a brief period in industry. It was during this time that she realised her interest lay not just in building technology, but in applying it to solve real-world problems. This led her back into academia, where she undertook an interdisciplinary doctorate in Web Science, an area that combines computer science with sociology to understand technology as a socio-technical system shaped as much by human behaviour as by code.

That perspective has defined her research ever since. During her PhD, which bridged chemistry and computer science, she encountered a surprising reality: despite rapid advances in digital technology, much of scientific research, particularly in academia, was still recorded on paper. What initially seemed like a straightforward problem quickly revealed deep technical and cultural complexities, sparking a research focus that has now spanned more than a decade.

Today, as a Principal Enterprise Fellow at the University of Southampton, Pearman-Kanza leads multiple initiatives to transform how scientific data is recorded, managed and shared. She is Principal Investigator of the Careers and Skills for Data-Driven Research (CaSDaR) network, Pathfinder Lead on process recording for the Physical Sciences Data Infrastructure (PSDI) and a researcher within the AI in Chemistry Hub.

Her work centres on digital notetaking and process recording, an area often overlooked, but critical to scientific reproducibility.

By studying how information flows between lab notebooks, publications, theses and supplementary data, she aims to ensure that not only the results, but also the full scientific processes, are captured accurately. This includes addressing how data can be made FAIR, findable, accessible, interoperable and reusable, while also recognising that existing frameworks often fall short in guiding researchers on how to produce usable data in the first place.

A major strand of her work has been the implementation of electronic lab notebooks (ELNs). Leading a large-scale trial and roll-out within her department, she has demonstrated both the potential and the challenges of replacing traditional paper-based systems.

Adoption rates have been strong, with a significant shift from paper to digital tools and sustained user engagement following initial trials. Her work now extends to undergraduate education and to advising other institutions, supported by national and international presentations and forthcoming publications.

Beyond implementation, she is addressing the next critical challenge: interoperability. As different institutions adopt a variety of digital tools, ensuring that data can be shared and understood across systems becomes increasingly important. Her research explores how standards and frameworks can support collaboration while accommodating diverse workflows and technologies.

Another key area of her work is the application of semantic web technologies. By using ontologies, structured vocabularies that encode meaning and relationships, she is helping to make scientific data more understandable to both humans and machines. Through collaborations across the UK and Europe, she is contributing to efforts to define shared standards for chemistry data and identifying gaps where new approaches are needed.

Alongside these technical efforts, Pearman-Kanza places strong emphasis on the human dimension of data-driven research. Through CaSDaR and her broader work with PSDI, she advocates for data stewards, specialists who support researchers in managing and curating data throughout its lifecycle. She has built communities, delivered training and developed funding initiatives to show the value of this role in improving research quality and impact.

Despite progress, she identifies reproducibility and the creation of genuinely FAIR data as the most significant challenges in her field.

These issues are not purely technical; they require changes in practice, culture and incentives across the research ecosystem.

Looking ahead, her goals include expanding the adoption of electronic lab notebooks across disciplines, advancing interoperability between systems and continuing to promote best practices in data management. Through these efforts, she aims to make high-quality, reusable data a standard outcome of scientific research rather than an exception.

For early-career researchers, her advice highlights the importance of community. Building networks, sharing knowledge and learning from others are essential in a field where no single individual can master all aspects. She also emphasises the value of mentorship, reflecting on how guidance and support have shaped her career.

At its core, Pearman-Kanza's work addresses a fundamental aspect of science that is often taken for granted: how knowledge is recorded.

By combining technical innovation with an understanding of human behaviour, she is helping ensure that scientific data is not only generated, but also preserved and shared in ways that maximise its value.

Burkhard Schäfer

Organisation: Splashlake
Role: Managing Director
Based in: Darmstadt, Germany

Burkhard Schäfer is an architect of scientific ecosystems. Colleagues often describe his work as bridging the gap between data, systems and people, ensuring that scientific information is not only generated, but remains accessible, meaningful and usable over time.

It is a perspective shaped by more than two decades in laboratory informatics. Throughout his career, Schäfer has combined technical depth with practical leadership, focusing on how data flows through complex scientific environments. Today, as Managing Director of Splashlake, he works at the forefront of modern laboratory data architecture, developing systems that support interoperability, scalability and long-term data value.

His research and development efforts centre on a fundamental challenge: making scientific data easier to access and reuse. In many laboratories, data remains locked within proprietary formats or fragmented across disconnected systems, making downstream analysis – particularly with artificial intelligence (AI) – both costly and inefficient. Schäfer's approach addresses this by ensuring that data is captured once, contextualised correctly and made available in a form that can be consumed equally by researchers, automated systems and AI-driven tools.

Schäfer's career began at the intersection of science and computing, working on instrument control and scientific data standards at Los Alamos National Laboratory and later at the National Institute of Standards and Technology. These early experiences established a foundation in both technical systems and standards development, which would become a defining theme throughout his work.

He later co-founded BSSN Software, a company focused on laboratory informatics solutions, which was subsequently acquired by Merck KGaA. This transition from research to product development reinforced his focus on creating practical tools that address real-world scientific challenges.

A central element of Schäfer's work has been his commitment to standards and community building. He is widely recognised as the lead architect of the ASTM AnIML standard, a key framework for analytical data exchange, and serves on the board of the SiLA Consortium, which promotes interoperability in laboratory automation. As an SLAS Fellow, he continues to contribute to the broader laboratory informatics community, advocating for collaborative, standards-based approaches.

At Splashlake, his work focuses on developing cloud-native, service-oriented architectures for laboratory data. Rather than relying on isolated tools, his approach emphasises cohesive system design, structured data models, robust APIs and automation-ready workflows that enable seamless integration across instruments, software platforms and research domains. These architectures reduce manual data handling, accelerate the onboarding of new technologies and improve the long-term usability of scientific data.

Despite advances in computing, Schäfer identifies a persistent challenge: managing complexity at scale without constraining innovation. Modern laboratories rely on increasingly diverse instruments and software systems, each evolving independently. Without clear architectural principles and shared data semantics, this diversity can lead to fragile integrations and growing technical debt.

Balancing flexibility and robustness is therefore critical. Scientific systems must support exploratory research while maintaining traceability, compliance and trust in data.

For Schäfer, this challenge is as much cultural as it is technical, requiring alignment across organisations, disciplines and communities.

Looking ahead, he aims to advance the development of closed-loop, data-driven experimentation, laboratory environments in which data not only informs analysis, but also actively feeds back into experimental design and decision-making.

Achieving this at scale, beyond isolated pilot projects, would represent a significant step towards more adaptive and efficient scientific research.

His vision is to enable this transformation incrementally, through architectures that allow laboratories to evolve without disruptive system overhauls. In doing so, he hopes to lower barriers to adoption and make advanced data-driven approaches accessible across a wider range of scientific domains.

For early-career researchers and engineers, Schäfer emphasises the importance of both depth and breadth.

Developing expertise in a specific domain is essential, but so, too, is curiosity about adjacent fields. Influence in scientific computing, he argues, comes from the ability to translate between disciplines, connecting scientific needs with technological solutions.

Equally important is engagement with the wider community. Participation in standards organisations, collaborative initiatives and open discussions helps to shape the direction of the field. Long-term impact, he suggests, is built not through isolated achievements, but through sustained contributions to shared challenges.

At the heart of Schäfer's work is a commitment to what he describes as "digital sustainability": creating data systems that remain usable, adaptable and valuable over decades. In an era of rapidly increasing data volumes, this long-term perspective is essential.

By designing architectures that prioritise accessibility, interoperability and trust, he is helping to ensure that scientific data continues to deliver value well into the future.

Phil Spence

Organisation: HotHouse Therapeutics
Role: Head of Informatics
Based in: UK

Discovering new-to-nature chemistry presents a fundamental challenge, with few established methods or datasets to guide exploration. Traditional approaches rely heavily on trial and error, making it difficult to navigate the vast biological and chemical search space efficiently. Phil Spence is working to address this by developing computational strategies that can prioritise the most promising directions for experimental investigation.

His work sits at the intersection of chemistry, biology and computation, where artificial intelligence (AI) and informatics are opening new routes to drug discovery. By combining machine-learning, bioinformatics and knowledge-driven approaches, he focuses on turning large, complex datasets into actionable insights – identifying new molecules, understanding the genes that encode them, and accelerating the discovery of novel therapeutics.

It is a shift that began early. Trained initially as a chemist, with an MChem and PhD from University of East Anglia, Spence's career took a decisive turn during his postdoctoral research at the John Innes Centre. There, he began applying machine-learning to problems in natural product biosynthesis and bioactivity, moving from purely experimental work towards computational approaches that could more directly guide discovery.

That transition laid the foundation for his current role as co-founder and Head of Informatics at HotHouse Therapeutics. At the company, he leads the development of computational platforms designed to predict the chemistry and biology of natural products, accelerating the identification of new therapeutic candidates.

Central to this effort is BotanAI, a platform that integrates AI, bioinformatics and knowledge-driven methods to explore complex biological and chemical datasets. The goal is to identify not only novel molecules, but also the genes that encode them – effectively uncovering new pathways for drug discovery by leveraging nature's own biochemical diversity.

Spence's work reflects a broader transformation in life sciences, where computational tools are increasingly used to navigate vast search spaces that would be impractical to explore experimentally. By combining machine learning, knowledge graphs and reinforcement learning, his team can prioritise the most promising targets for laboratory validation, reducing reliance on traditional trial-and-error approaches. This strategy is particularly important in the context of natural products, where the diversity of chemical structures and biological pathways presents both opportunity and complexity. By focusing computational effort on the most informative regions of this space, Spence's work aims to make discovery more efficient and targeted.

However, he is clear about the challenges involved. Unlike more established areas of computational biology, this field lacks standardised workflows and well-curated datasets at the scale required. The ambition to generate "new-to-nature" chemistry, molecules not previously observed in biological systems, means that many of the tools and approaches must be developed from first principles.

This absence of precedent creates both risk and opportunity. On the one hand, it requires building new computational frameworks capable of integrating diverse data types and making reliable predictions in uncertain environments. But it opens the door to entirely new classes of therapeutics that could not be discovered through conventional methods.

Looking ahead, Spence's goal is to demonstrate that AI-driven discovery platforms can move beyond prediction to deliver experimentally validated drug candidates. Establishing a system that can reliably identify new genes, molecules and therapeutic opportunities, and translate these into real-world outcomes, would represent a significant milestone.

For early-career researchers, his advice reflects the interdisciplinary nature of his work. The most interesting problems, he argues, often lie at the boundaries between fields. Developing the ability to think across disciplines not only broadens perspective but also enables the identification of novel solutions that might otherwise be overlooked.

At its core, Spence's work highlights the growing role of scientific computing in reshaping drug discovery. By integrating computational intelligence with biological insight, he is helping to create new pathways from data to therapeutics – bridging the gap between what can be predicted and what can be realised in the laboratory.

Andrea Townsend-Nicholson

Organisation: University College London
Role: Vice-Dean Health, Faculty of Life Sciences
Based in: London, UK

Andrea Townsend-Nicholson applies high-performance computing to drug discovery, using simulations to predict how proteins behave and how they can be targeted therapeutically. Her approach shifts experimentation from trial-and-error to data-driven design.

It is a shift that has reshaped her own research. "We use computational methods to work out which experiments we will do, rather than trying many different things in the lab to see what works," she explains, reflecting on a transition that has seen her laboratory move from being almost entirely experimental to predominantly computational. The result is a more efficient, scalable approach to biomedical discovery, one that reduces cost, accelerates insight and even improves sustainability by minimising laboratory waste.

Prof Townsend-Nicholson's work focuses on proteins located on the surface of cells, key targets for many therapeutic drugs. These proteins are notoriously difficult to study experimentally, embedded as they are within complex cell membranes. For years, progress was limited by the lack of detailed structural data. When high-resolution atomic structures became available in the late 2000s, she seized the opportunity to integrate computational approaches into her research, marking a turning point in her career.

At University College London, her lab now uses advanced simulation techniques to understand how proteins behave and interact with potential drug molecules. By employing ensemble-based molecular dynamics simulations and quantum-informed methods, her team can explore molecular interactions at a level of detail that would be impossible to achieve by experiment alone. These approaches allow researchers to identify promising compounds computationally before validating them in the laboratory.

Her current research includes developing a novel anti-obesity therapeutic to provide alternatives for patients who do not respond to existing GLP-1-based treatments. The target protein presents unusual challenges for experimental characterisation, making computational modelling an essential tool in identifying viable drug candidates.

Beyond research, Townsend-Nicholson has become a strong advocate for expanding access to high-performance computing in the biomedical sciences. Recognising that relatively few researchers in her field engage with large-scale computing, she has embedded HPC training into undergraduate and medical education. Since 2017, she has taught about 150 students each year to use supercomputers, contributing to a growing community of computer-skilled biologists and clinicians.

Her efforts are already having a lasting impact. Thousands of students have gained hands-on experience with national computing resources, including systems such as the Snellius supercomputer and ARCHER2. For some, this exposure has been transformative, inspiring new career paths at the intersection of biology and computing.

Townsend-Nicholson is also involved in shaping the future of computational infrastructure. As one of the investigators contributing to the development of Charger, a UKRI-funded supercomputer within the UK's National Compute Resource, she is helping to ensure that advanced computing systems are accessible to a broader range of scientific disciplines. Her work extends into collaborations that integrate biomolecular simulation with quantum computing.

Despite these advances, she identifies a fundamental challenge: validation. Computational models, no matter how sophisticated, remain predictions until they are confirmed experimentally. Ensuring simulations reflect biological reality requires careful alignment between digital and laboratory approaches, and remains a concern in computational biomedicine.

Looking ahead, Townsend-Nicholson hopes to revisit experimental data generated earlier in her career. With modern computational tools, she aims to uncover the molecular mechanisms underlying observations that were previously difficult to interpret. Solving even one of these longstanding questions, she says, would represent a significant personal milestone.

Her advice to early-career researchers reflects her own experience. The most rewarding opportunities, she says, often arise unexpectedly and require a willingness to act quickly. Equally, the most exciting advances tend to occur at the boundaries between disciplines or between established fields and emerging technologies. Remaining open to these intersections is key.

Townsend-Nicholson also contributes to international efforts to advance computational biomedicine. As part of the CompBioMed consortium, she works on the multiscale challenge of building a "Virtual Human", with her focus on the molecular level, linking DNA and protein behaviour to broader physiological function. She is also active in the global high-performance computing community, co-organising workshops and student initiatives at major international conferences.

At its core, her work reflects a broader transformation in science: one in which computation and experiment are no longer separate, but deeply intertwined. By championing both the tools and the people needed to drive this change, Townsend-Nicholson is helping to shape the future of biomedical research.

Scientist-led or agent-led AI?

Rob Brown, VP and Head of the Scientific Office at Sapio Sciences, answers the question of whether AI belongs in research

Rob Brown

The question of whether AI belongs in research has been answered. It does.

The more interesting question is which operating model fits which category of work. Two distinct models are taking shape: one scientist-led, the other agent-led. They are not competing versions of the same process. They suit different kinds of research activity, place human control at different points and depend on the same underlying foundations: governed data, validated tools and a connected ecosystem.

The hurdle isn't the model's intelligence; it's the scientist's trust.

The scientist-directed model

Every scientist recognises the Design, Make, Test, Analyse (DMTA) cycle, with make and test happening in the ELN and design and analysis running in dedicated computational software. Running those tools means either learning new skills and UIs or engaging computational teams and waiting, sometimes a week, for results that should be available in minutes.

For two decades, the industry tried to solve this with custom interfaces of computational tools that scientists might actually use. They never worked. Every new interface was still another UI to learn, another context switch, another barrier between the scientist and the result.

Natural language finally removes that barrier entirely.

A single series of experiments might draw on 10 different computational packages from six vendors, each with its own UI and file format. An AI co-scientist embedded in the ELN calls the appropriate validated tool at the scientist's direction and returns the output directly into the experiment record. Scientists don't need to learn how the software works. They just say what they want.

The agent-directed model

The other approach is the agent-directed model. Here the scientist defines the research question or hypothesis; the platform orchestrates the tools, data sources and computational steps needed to test it, returning evidence or asking for scientists' input at the points where human review matters most.

For functions where scale and complexity make manual coordination impractical, such as target assessment, molecular design and lead generation, agent-led execution is the more realistic way to move from question to tested hypothesis at the pace the work demands.

What makes this model credible, rather than just fast, is transparency. Agent-led operations require high data quality, well-defined review thresholds and full audit trails across the pipeline. Without that, autonomy creates exposure instead of value. With it, the coordination burden that would otherwise consume scientific capacity gets absorbed by the system and the scientist stays accountable for the decisions that matter.

Data is the glue

Neither model works without the same underlying discipline: accurate, structured, governed data.

Scientific data accumulates across instruments, LIMS records, legacy documents and disconnected systems that were never designed to speak to one another. Inconsistent naming conventions, absent metadata and siloed departmental standards mean the same entity can mean different things depending on where in the discovery pipeline it sits. AI doesn't interrogate the data it's given; it just acts on it. So bad data doesn't slow the system down; it results in faster chaos.

The organisations making real progress treat data readiness as a prerequisite, not a parallel track. The AI readiness gap closes from the foundation up, not from the tool down.

Two companies, one ecosystem

That is the logic behind what we have built across Sapio Sciences and Sigmatic Sciences.

In the scientist-directed model, Sapio's AI co-scientist works within the ELN, helping scientists to move work forward, access validated tools and keep every step inside a governed scientific record.

In the agent-directed model, Sigmatic Sciences, a Sapio Sciences company, extends that same environment into lab-connected orchestration: a scientist's question is translated into testable hypotheses and evaluated across validated computational tools and live experimental data.

Sapio Elain and Sigmatic Scout are two points on the same journey: one making the scientist-led DMTA cycle faster and more capable, the other taking entire research cycles virtual before the wet lab is ever involved.

Science does not operate in one mode. The infrastructure supporting it should not either.

Rob Brown is Head of the Scientific Office at Sapio Sciences. He previously served as head of global presales, product marketing, and product management at Dotmatics and earlier led product marketing teams at Accelrys, SciTegic, and MSI.

Alex Zhavoronkov

Organisation: Insilico Medicine
Role: Founder and CEO
Based in: United States

Alex Zhavoronkov is mission-driven. Colleagues often point to a defining characteristic: an unwavering commitment to tackling some of the most difficult problems in medicine, particularly ageing and age-related disease.

It is a focus rooted in both personal conviction and strategic thinking. After working in the GPU industry in the early 2000s and achieving financial success in semiconductors, Zhavoronkov made a decisive shift. Rather than continue in computing hardware, he chose to invest his time and resources into biotechnology, with the explicit goal of extending healthy, productive human lifespan.

As founder and CEO of Insilico Medicine, he has built one of the most prominent companies at the intersection of artificial intelligence and drug discovery.

Under his leadership, the company has progressed 29 developmental candidates, with 12 advancing to the clinical stage, including multiple Phase I and Phase II trials. The pace and scale of this pipeline reflect a broader ambition: to fundamentally accelerate the discovery and development of new medicines.

Zhavoronkov's work centres on applying advanced computational approaches to biology and chemistry.

At the core of this effort is the Pharma. AI platform, a suite of tools including Biology42, Chemistry42 and Medicine42, designed to integrate target discovery, molecular design and clinical insight. These systems leverage machine-learning and large-scale data analysis to identify novel therapeutic targets and generate candidate compounds with unprecedented speed.

While the ultimate goal is to "solve ageing", the current focus is more immediate: addressing multiple age-related diseases by targeting pathways implicated in both ageing and disease progression. This dual approach allows the company to pursue therapies that not only treat specific conditions, but also contribute to a broader understanding of longevity.

The impact of these methods lies in their ability to compress timelines that have traditionally taken years or decades. By combining computational modelling with experimental validation, Zhavoronkov's team is working to transform drug discovery into a more efficient, data-driven process. Yet, despite advances in artificial intelligence, he is clear that "one of the field's central challenges remains experimental validation in animals and humans".

Computational predictions must ultimately be tested in biological systems before they can meaningfully influence drug development. Bridging this gap between digital prediction and real-world biology remains one of the defining challenges of modern scientific computing in healthcare.

Looking ahead, Zhavoronkov has set ambitious milestones.

Over the next two-to-three years, he aims to nominate 40-50 preclinical candidates and complete the first Phase III clinical trial for a therapeutic discovered using AI.

Achieving this would represent a major step forward, demonstrating not only the efficiency of AI-driven discovery, but also its capacity to deliver clinically validated treatments.

Beyond drug discovery, his work extends into sustainability.

Through collaborations with major energy organisations, including Aramco, Insilico Medicine is applying its molecular design capabilities to challenges such as carbon capture, hydrogen storage and clean synthetic fuels. This reflects a broader view of scientific computing as a tool not only for healthcare, but for addressing global challenges.

Zhavoronkov is also direct in his advice to early-career researchers. Motivation, he argues, must come from a clear understanding of the human impact of the work. Spending time with patients facing life-threatening illness can provide a powerful sense of purpose, for example.

At the same time, he emphasises discipline and persistence: progress in fields such as ageing and cancer research demands sustained effort and a willingness to go beyond conventional limits.

"If you are not a genius," he says, "you need to work smarter – but also harder than everyone else."

It is a philosophy that reflects both the urgency of the problems he is trying to solve and the intensity required to address them.

At its core, Zhavoronkov's work represents a convergence of AI, biology and real-world application. By pushing computational methods into the heart of drug discovery, he is helping to redefine how new therapies are created – and, potentially, how ageing, itself, is understood and treated.

Computer-aided engineering, modelling and simulation

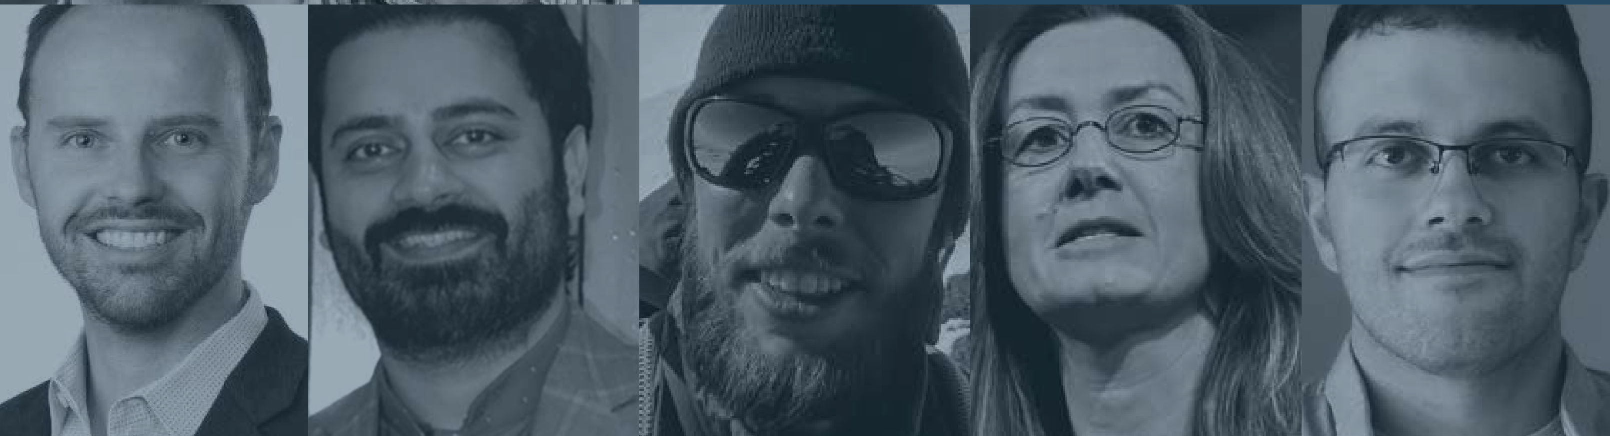

Luisa Attfield

Organisation: GSK
Role: Associate Data Scientist
Based in: UK

Luisa Attfield works on one of pharmaceutical R&D's biggest emerging challenges: turning complex datasets into decisions that can accelerate drug development.

At GSK, she applies data science and computational methods to improve efficiency across research and manufacturing workflows.

Her work reflects a growing shift in pharmaceutical development, where computational engineering is used to optimise manufacturing processes, reduce timelines and improve decision-making.

She works with large datasets generated across research, development and production environments, applying statistical modelling, machine-learning and process analytics to identify patterns that can help to improve operational performance.

In a highly regulated industry, one of the key challenges is ensuring computational tools remain robust, interpretable and compatible with existing scientific workflows.

Attfield represents a new generation of computational engineers whose work lies at the intersection of traditional engineering disciplines and modern data science.

As pharmaceutical companies continue to invest in AI, digital manufacturing and predictive analytics, professionals such as Attfield are playing an increasingly important role in translating computational innovation into practical outcomes.

Agnes Blom-Schieber

Organisation: Boeing
Role: Technical Fellow, Structures
Based in: United States

Agnes Blom-Schieber is a recognised leader in computational engineering and simulation, with a distinguished career at Boeing advancing the application of scientific computing to aerospace structures.

As a Technical Fellow in Structures, she has played a key role in integrating high-fidelity numerical methods into aircraft design and certification.

Her work has focused on scalable finite element models and multiphysics simulation frameworks, enabling engineers to predict structural behaviour under complex loading conditions with greater accuracy.

Blom-Schieber has championed the use of simulation and high-performance computing to reduce reliance on costly physical testing while maintaining rigorous safety standards.

A key achievement has been Blom-Schieber's leadership in digital twin methodologies, where she helped establish simulation-driven lifecycle management for aircraft structures.

This has improved predictive maintenance and extended service life through data-informed decision-making.

Beyond her technical contributions, she is known for fostering collaboration between structural engineers, computer scientists and data specialists, while mentoring early-career engineers and helping shape Boeing's long-term strategy in computational science.

Jamie Bramwell

Organisation: Lawrence Livermore National Laboratory
Role: Director, Center for Computational Engineering
Based in: United States

Jamie Bramwell works on computationally demanding engineering problems, using advanced simulation to help researchers understand systems that are too complex, expensive or dangerous to test physically.

As Director of the Center for Computational Engineering at Lawrence Livermore National Laboratory, Bramwell leads efforts to build next-generation engineering simulation software for modern supercomputers.

She began at the lab in 2013 as a computational engineer specialising in advanced finite element methods, initially working on ALE3D, a large-scale multiphysics code used in national security applications. There, she led development of the implicit solid mechanics physics package and helped modernise solver infrastructure.

She later joined the MFEM team, contributing to cardiac mechanics simulation tools, before launching the Smith project in 2020: a modular software development kit for agile, GPU-ready and sensitivity-enabled engineering simulations.

As engineering becomes increasingly reliant on simulation-led design, Bramwell's work reflects a broader shift toward faster, more adaptive computational tools that can keep pace with increasingly complex scientific challenges.

Christian Becker

Organisation: Sartorius Stedim Biotech GmbH
Role: Manager CAE Simulation
Based in: Germany

Christian Becker translates complex computational methods into tangible value in industrial settings, ensuring that simulation is not just a technical capability, but a core part of decision-making. Having previously worked in more traditional engineering fields, he transitioned to biomedical engineering and implemented CAE solutions at Sartorius.

Becker's work centres on the implementation of CAE simulation as a standard tool within product development. At Sartorius Stedim Biotech, he leads efforts to integrate simulation into engineering workflows, helping to accelerate development processes, reduce reliance on physical prototypes and improve efficiency.

His path into the field began with a background in mathematics, complemented by mechanical engineering.

An early role as a finite element analysis (FEA) engineer working on aeroplane doors provided his first exposure to CAE, offering a practical foundation in structural simulation and engineering analysis.

A subsequent move into technical sales at COMSOL marked an important transition. There, Becker gained exposure to a broader range of physical modelling domains and developed a deeper understanding of how simulation tools are applied across industries. This experience also sharpened his perspective on the organisational challenges of adopting computational methods, particularly the need to align technical capability with business value.

Today, his work focuses on embedding CAE simulation across Sartorius' engineering processes. Simulation is used to optimise product design, reduce development time and minimise prototyping costs. It also plays a role in customer engagement, providing insights into complex processes and supporting the scaling of biopharmaceutical systems.

Despite its advantages, Becker identifies a key challenge: driving adoption at the organisational level. Convincing senior management to invest in both simulation software and skilled personnel can be difficult, particularly in sectors such as biopharma, where laboratory-based approaches have traditionally dominated. Demonstrating the long-term value of CAE, both in cost savings and improved product performance is, therefore, essential.

For Becker, success lies in making simulation accessible and relevant to non-specialists, ensuring that it becomes an integral part of engineering culture rather than a niche capability.

Looking ahead, his focus remains on strengthening CAE's role in industrial development. By continuing to expand its use and demonstrating its value, he aims to ensure that simulation becomes a standard, trusted tool across the organisation.

His advice to early-career researchers and engineers reflects his own journey. Developing deep expertise in one area is essential, but so, too, is the willingness to broaden one's perspective. After establishing a strong foundation, he encourages moving into new fields where existing knowledge can be applied in different contexts, enabling both personal growth and wider impact.

At its core, Becker's work contributes to a larger goal.

By embedding advanced computational methods into biopharmaceutical development, he supports efforts to make drug production more efficient and cost-effective. In doing so, he helps enable better medicines to reach patients more quickly and at lower cost, a practical demonstration of how scientific computing can deliver real-world benefits.

Michael Broadhurst

Organisation: BWT Alpine Formula One Team
Role: Chief Aerodynamicist
Based in: UK

Michael Broadhurst works in one of the most simulation-intensive engineering environments, where competitive advantage is often measured in milliseconds.

As Chief Aerodynamicist at the BWT Alpine Formula One Team, he leads the computational engineering work that shapes the aerodynamic performance of Formula One cars.

Modern Formula One design relies heavily on computational fluid dynamics (CFD), virtual modelling and wind-tunnel validation to optimise vehicle performance under strict development timelines and regulatory constraints. Broadhurst's role sits at the centre of that process, overseeing aerodynamic development programmes that rely on large-scale simulation to improve speed, efficiency and race-day performance.

Before joining the BWT Alpine Formula One Team, Broadhurst built extensive experience in motorsport aerodynamics, working in environments where engineers must rapidly test thousands of design iterations while balancing simulation results against physical validation. The challenge is not simply generating more data, but identifying which simulations translate into measurable on-track gains.

As Formula One teams increasingly adopt AI tools, advanced modelling and automation, computational engineering is becoming even more central to competitive performance. Broadhurst's work reflects how digital engineering tools are reshaping elite motorsport development.

Are Magnus Bruaset

Organisation: Simula Research Laboratory
Role: Professor and Director of Software & AI
Based in: Oslo, Norway

Are Magnus Bruaset is an Adjunct Professor at the Oslo Metropolitan University and Director of Software and AI at Simula Research Laboratory, where he has established himself as a leading figure in scientific computing and high-performance software development. His career has been defined by advancing computational methods that underpin large-scale simulation and data-intensive science.

Bruaset's research has focused on numerical linear algebra, parallel algorithms, and the design of robust software frameworks for solving complex partial differential equations. He has contributed to the development of scalable computing solutions that enable efficient utilisation of modern supercomputing architectures, supporting applications in energy, climate modelling and engineering analysis.

A significant achievement has been his leadership in bridging traditional simulation with emerging artificial intelligence techniques.

By integrating machine-learning into scientific workflows, he has helped accelerate computation and improve model fidelity, particularly in areas that require rapid data assimilation and uncertainty quantification.

As Director, Bruaset has driven strategic initiatives in software engineering and AI, fostering collaboration between mathematicians, computer scientists and domain experts.

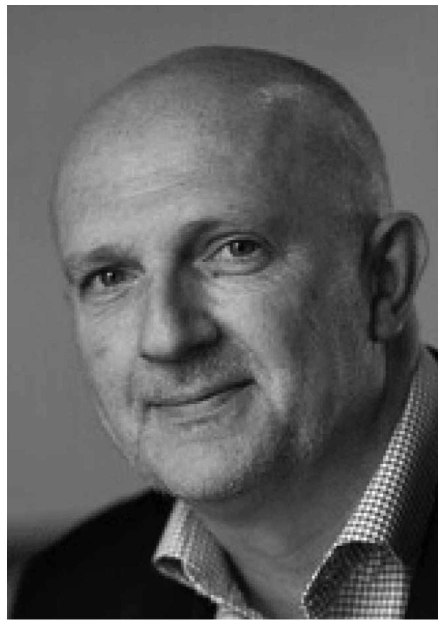

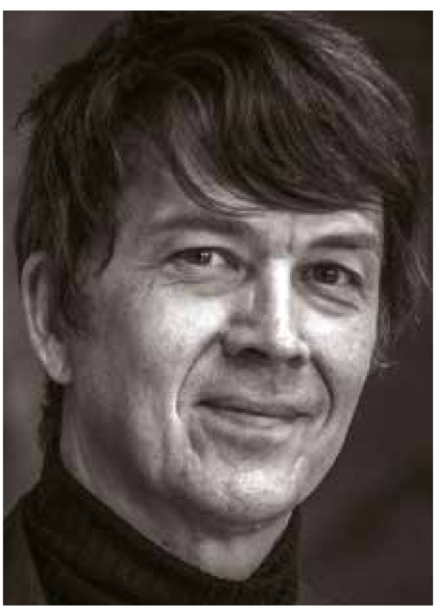

Gerbrand Ceder

Organisation: University of California, Berkeley
Role: Professor; Samsung Distinguished Chair in Nanoscience and Nanotechnology Research
Based in: California, United States

As materials science becomes increasingly computational and automated, Gerbrand Ceder has been at the centre of that transformation. His group is recognised for combining theoretical modelling, selected experiments and robotics, while contributing extensively to the Materials Project and advancing AI-driven autonomous laboratories for materials synthesis and characterisation.

As Professor of Materials Science and Engineering at University of California, Berkeley, and Samsung Distinguished Chair in Nanoscience and Nanotechnology Research, Ceder is widely recognised as a pioneer of computational materials discovery.

He played a major role in founding the Materials Project, one of the world's most widely used open materials databases, which provides computational data on thousands of compounds. It has become a major tool for scientists working on batteries, semiconductors and clean energy technologies.

Ceder has also co-founded several companies translating computational materials research into commercial applications, particularly in battery development and electrification technologies.

As artificial intelligence becomes increasingly embedded in materials discovery, his work remains central to efforts to shorten the time between theoretical prediction and real-world deployment.

Mattia De Dominicis

Organisation: Imperial Brands
Role: SVP R&D Innovation
Based in: UK

In an era where data and digital tools are redefining product development, Mattia De Dominicis is driving the transformation of R&D. His work centres on integrating AI, analytics and experimental design to deliver faster, more informed innovation.

With a career spanning multiple consumer sectors, De Dominicis has consistently used science and data to deliver commercially impactful innovation. His work combines scientific expertise, consumer insight and AI-driven approaches to accelerate innovation across complex global organisations.

His trajectory reflects both breadth and strategic evolution. With senior leadership roles at Imperial Brands, Henkel, Danone and Reckitt, De Dominicis has worked across household, consumer health, food and personal care categories. Throughout, he has focused on translating scientific knowledge and consumer understanding into "market-leading innovations with measurable commercial impact".

A defining feature of his career is global leadership. Having lived in and led R&D teams across Europe, the United States and Africa, he has developed a leadership approach rooted in multicultural collaboration, agility and continuous innovation. His work emphasises three interconnected pillars: data, people and innovation.

Now serving as SVP R&D – Innovation at Imperial Brands, De Dominicis is focused on a major industry transformation. His work aims to shift product portfolios from traditional tobacco-based products towards reduced-harm alternatives, including nicotine and eventually zero-nicotine solutions. This aligns with broader trends in health, well-being and sustainability that have shaped much of his career.

Central to this transformation is the integration of scientific computing into product development. His teams employ design of experiments and advanced data analytics to optimise formulations and processes, while increasingly incorporating AI across the innovation lifecycle. These include AI-driven synthetic personas to better understand consumer behaviour, tools that integrate and query internal knowledge across clinical and analytical datasets, and predictive systems that combine sensory data, emotional responses and biometrics.

This approach reflects a broader shift in R&D: from iterative experimentation to predictive, data-driven innovation. De Dominicis has also contributed to this

transition through thought leadership, including work on end-to-end data analytics for product development, helping to formalise how data can accelerate innovation cycles.

Despite these advances, he identifies a critical challenge in scientific computing: accurately predicting biological responses and toxicological impacts. While computational models continue to improve, they remain limited in their ability to fully capture the complexity of biological systems. Emerging technologies such as quantum computing offer potential solutions, but widespread, reliable applications remain several years away.

Looking ahead, De Dominicis aims to further integrate advanced technologies, including AI, data analytics and quantum computing, into product development, with a focus on improving wellbeing and sustainability outcomes. His goal is not only to accelerate innovation, but to ensure that it delivers meaningful societal benefit.

Beyond his core role, he is actively engaged in broader innovation and leadership communities. He participates in global forums such as London Tech Week and innovation networks, and has contributed to sustainability initiatives, including work on packaging and raw materials during his time at Reckitt and Henkel.

He has also been involved in social impact initiatives, including collaborations with organisations such as Save the Children, focused on public health challenges.

For early-career professionals, his advice reflects his leadership philosophy. Technical expertise alone is not sufficient; success also depends on emotional intelligence, empathy and the ability to lead and inspire teams in complex environments.

At its core, De Dominicis' work illustrates how scientific computing is reshaping industrial innovation.

By combining data, advanced analytics and human-centred design, he is helping to redefine how products are developed, moving from reactive processes to predictive, insight-driven systems that can deliver both commercial and societal value.

Jackie Chen

Organisation: Sandia National Laboratories
Role: Senior Scientist, Chemistry, Combustion and Materials Science Center
Based in: United States

Jackie Chen is a Senior Scientist in the Chemistry, Combustion and Materials Science Center at Sandia National Laboratories, and a leading figure in scientific computing for reactive flows. Her work has been central to advancing high-fidelity simulation of combustion processes, with wide-ranging applications in energy, propulsion and environmental modelling.

Chen is known for her contributions to direct numerical simulation (DNS) of turbulent combustion, where she has developed and applied large-scale computational models to resolve the interactions between chemical kinetics and fluid dynamics. Aided by supercomputing systems, she has enabled unprecedented insights into flame structure, ignition and pollutant formation.

A key achievement has been her role in creating scalable algorithms and data analysis techniques capable of handling the immense datasets generated by DNS. Her work has also supported the development of reduced-order models, improving the practical applicability of high-resolution simulations in engineering design.

Beyond her research, Chen has been active in fostering interdisciplinary collaboration and mentoring scientists in computational combustion.

She was elected to the National Academy of Engineering in 2018 for her work on computational simulation of turbulent reacting flows with complex chemistry, and in 2020 was named a DOE Office of Science Distinguished Scientist Fellow.

Joerg Gablonsky

Organisation: Boeing
Role: Technical Fellow
Based in: United States

Joerg Gablonsky is a Boeing Technical Fellow specialising in scientific computing, numerical optimisation and high-performance computing (HPC).

He leads the numerical optimisation group within Boeing Research & Technology's Applied Mathematics team and chairs the company's Enterprise High Performance Computing Council.

His work centres on simulation-based optimisation: developing algorithms, implementing them in software and helping engineering teams apply the right optimisation methods to demanding aerospace problems. His publications include work on the DIRECT algorithm, implicit filtering and parallel optimisation of complex computer simulations, reflecting a career grounded in both mathematical rigour and industrial application.

Gablonsky has also played a strategic role in enterprise computing as chair of Boeing's Enterprise High-Performance Computing Council, coordinating HPC resources across the company and guiding adoption of cloud and distributed computing infrastructures.

His research contributions include parallel optimisation methods and derivative-free algorithms, with publications in global optimisation, simulation-based design and computational engineering.

He has authored or co-authored around 20 scientific papers, with more than 1,000 citations.

Johannes Grau

Organisation: Volkswagen
Role: Head of Software Systems and Tools
Based in: Germany

Johannes Grau works on the software systems and engineering tools that support safety-critical automotive development. As Head of Software Systems and Tools at Volkswagen AG, his role sits within an area where simulation, requirements management and software qualification are increasingly central to vehicle engineering.

His work is linked to the toolchains used to develop and qualify complex automotive systems, including electronic power steering.

In modern vehicle programmes, engineering teams must connect requirements, simulation, testing and verification in ways that are traceable and robust enough for safety-critical development.

Grau's public comments highlight the importance of software tooling developed specifically for safety-critical projects and the use of qualified engineering tools as part of the carmaker's development foundation.

His role reflects a broader shift in computational engineering, where software systems are no longer supporting infrastructure, but essential components of product development.

Ger Janssen

Organisation: Philips
Role: AI Ethics & Compliance Lead; previously Principal Data and AI Scientist
Based in: The Netherlands

Ger Janssen works across artificial intelligence, healthcare innovation and responsible technology governance. His career reflects a shift that is increasingly important across computational engineering: moving advanced AI and modelling methods from research into regulated, trustworthy real-world use.

At Philips, Janssen has held senior roles spanning data science, computational modelling and AI governance. Previously Principal Data and AI Scientist, he worked on data-driven healthcare technologies and patient digital twin approaches, using computational models and artificial intelligence to better understand complex clinical systems and support more personalised care.

He now serves as AI Ethics & Compliance Lead, focusing on how AI-enabled healthcare technologies can be developed and deployed safely, transparently and responsibly.

This work is particularly important in medical and clinical environments, where algorithmic performance alone is not enough.

Janssen's role reflects a broader challenge facing scientific computing and computational engineering: ensuring that powerful models can be trusted by clinicians, patients, regulators and industry partners.

Holger Krag

Organisation: European Space Agency
Role: Head of Space Debris Office, ESOC
Based in: Darmstadt, Germany

Holger Krag works on one of spaceflight's most urgent engineering problems: how to keep satellites operating safely in an increasingly crowded orbital environment. As Head of ESA's Space Debris Office at the European Space Operations Centre in Darmstadt, Germany, Krag leads work that combines orbital mechanics, risk modelling, surveillance data and operational decision-making.

Krag joined ESA in 2006 as a space debris analyst, working on risk models and operational collision-avoidance systems. In 2014, he became Head of the Space Debris Office, which supports ESA's space situational awareness and space safety activities.

Krag has been instrumental in advancing computational models for space situational awareness, including high-precision orbit determination and collision risk assessment. By leveraging large-scale simulations and probabilistic modelling, he has helped develop systems capable of tracking thousands of objects in Earth orbit and forecasting potential conjunction events with operational satellites.

Krag's contributions to space flight safety have also been recognised internationally. In 2022, he received the Space Data Association's T.S. Kelso Award for his work on space debris analysis and collision avoidance.

Jarrod Murphy

Organisation: Mercedes-AMG Petronas Formula One Team
Role: Aerodynamics Director
Based in: UK

Jarrod Murphy works in one of engineering's most demanding simulation environments, where aerodynamic performance is shaped through rapid iteration, computational modelling and physical validation.

As Aerodynamics Director at the Mercedes-AMG Petronas Formula One Team, Murphy leads work at the core of motorsport development. His role involves guiding the car's aerodynamic direction, integrating computational fluid dynamics, wind-tunnel testing and track data to deliver measurable performance gains under strict regulatory and development constraints.

Murphy studied Aeronautical Engineering at Imperial College London before beginning his motorsport career in 1996 as a stress analyst at Benetton. He later moved into computational fluid dynamics, becoming Head of CFD at Renault and then Lotus before joining Mercedes in 2013. He became Chief Aerodynamicist, later Head of Aerodynamics, and was promoted to Director of Aerodynamics in 2021.

His career reflects the growing importance of computational engineering in motorsport, showing how CFD, aerodynamic modelling and engineering judgement combine in an environment where small improvements can define competitive advantage.

Steven Pierson

Organisation: Jaguar Land Rover
Role: Propulsion CAE Senior Technical Specialist
Based in: Coventry, UK

Steven Pierson works on the computational engineering behind modern propulsion systems, where simulation is essential to improving performance, reliability and efficiency before hardware reaches the test bench.

As a Propulsion CAE Senior Technical Specialist at Jaguar Land Rover, Pierson's work supports the development and validation of complex vehicle propulsion technologies. His role sits within a field where finite element analysis, thermal modelling, fluid simulation and system-level CAE are used to explore design options, reduce prototype dependency and identify issues earlier in the development cycle.

At JLR, propulsion engineering has become increasingly computational as vehicles move towards electrified, hybrid and battery-based architectures. This creates new modelling challenges, from thermal behaviour and durability to integration across mechanical, electrical and control systems. Pierson's work contributes to the development of simulation-led methods to manage this complexity and ensure that engineering decisions are grounded in robust analysis.

His public profile lists Jaguar Land Rover experience, a Coventry location and education at Cranfield University.

His work also relates to CAE integration in batteries, reflecting the growing importance of simulation in electrified propulsion development.

Pierson's career reflects the increasing role of CAE specialists in automotive engineering: not simply running models, but helping embed simulation into the development process so teams can move faster, reduce risk and build better propulsion systems.

Jordan Reynolds

Organisation: Intuitive Machines
Role: Avionics Engineer
Based in: United States

Jordan Reynolds works on the avionics and mission modelling needed to make commercial lunar exploration viable. As an Avionics Engineer at Intuitive Machines, he supports spacecraft systems that must operate reliably in some of the most unforgiving engineering environments.

Reynolds has contributed to mission simulation and analysis for lunar technologies, including Intuitive Machines' Nova-C lander and µNova Hopper. His work uses digital mission engineering tools such as Ansys Systems Tool Kit to model trajectories, lighting conditions, communications geometry and thermal behaviour across complex lunar scenarios. These simulations show how spacecraft will perform when physical testing on Earth is impossible or impractical.

A central challenge in this work is confidence. Space missions depend on systems that must survive launch, deep-space conditions, extreme temperature cycles and limited communications windows. Reynolds has described simulation as a way to model mission conditions, assess thermal and power behaviour, and make informed decisions quickly during development.

His role reflects the increasing importance of computational engineering in the commercial space sector. By combining avionics expertise with mission simulation, Reynolds helps translate spacecraft concepts into systems that can be tested, refined and flown.

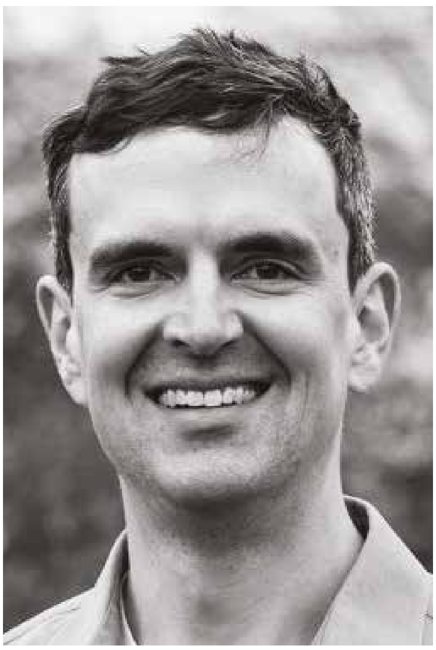

Tyler Rourke

Organisation: Perellion
Role: Mechanical Engineer
Based in: United States

Tyler Rourke works on the practical side of computational engineering, where simulation, design and fabrication come together to solve complex mechanical problems. His work reflects a hands-on approach to engineering: using digital tools not as abstract modelling exercises, but as part of a wider process that links concept development, physical manufacture and real-world performance.

As a Mechanical Engineer at Perellion, Rourke contributes to projects that combine automation, manufacturing and advanced simulation. He joined the company in 2022, bringing more than two decades of fabrication experience and a decade-plus of engineering practice to the role. That background gives him a valuable perspective in simulation-led design, where the usefulness of a model depends not only on technical accuracy, but on whether it can inform decisions that are practical to build, test and deploy.

His work includes integrating simulation into engineering design workflows. In a case study with Ansys, Rourke was part of the Perellion team using simulation to improve engineering accuracy, support iterative virtual design and optimise product development.

This work is increasingly important as mechanical engineering moves from a linear design-build-test model. Instead of waiting until late-stage prototyping to identify weaknesses, engineers use simulation to evaluate designs, understand failure modes and refine concepts before manufacture.

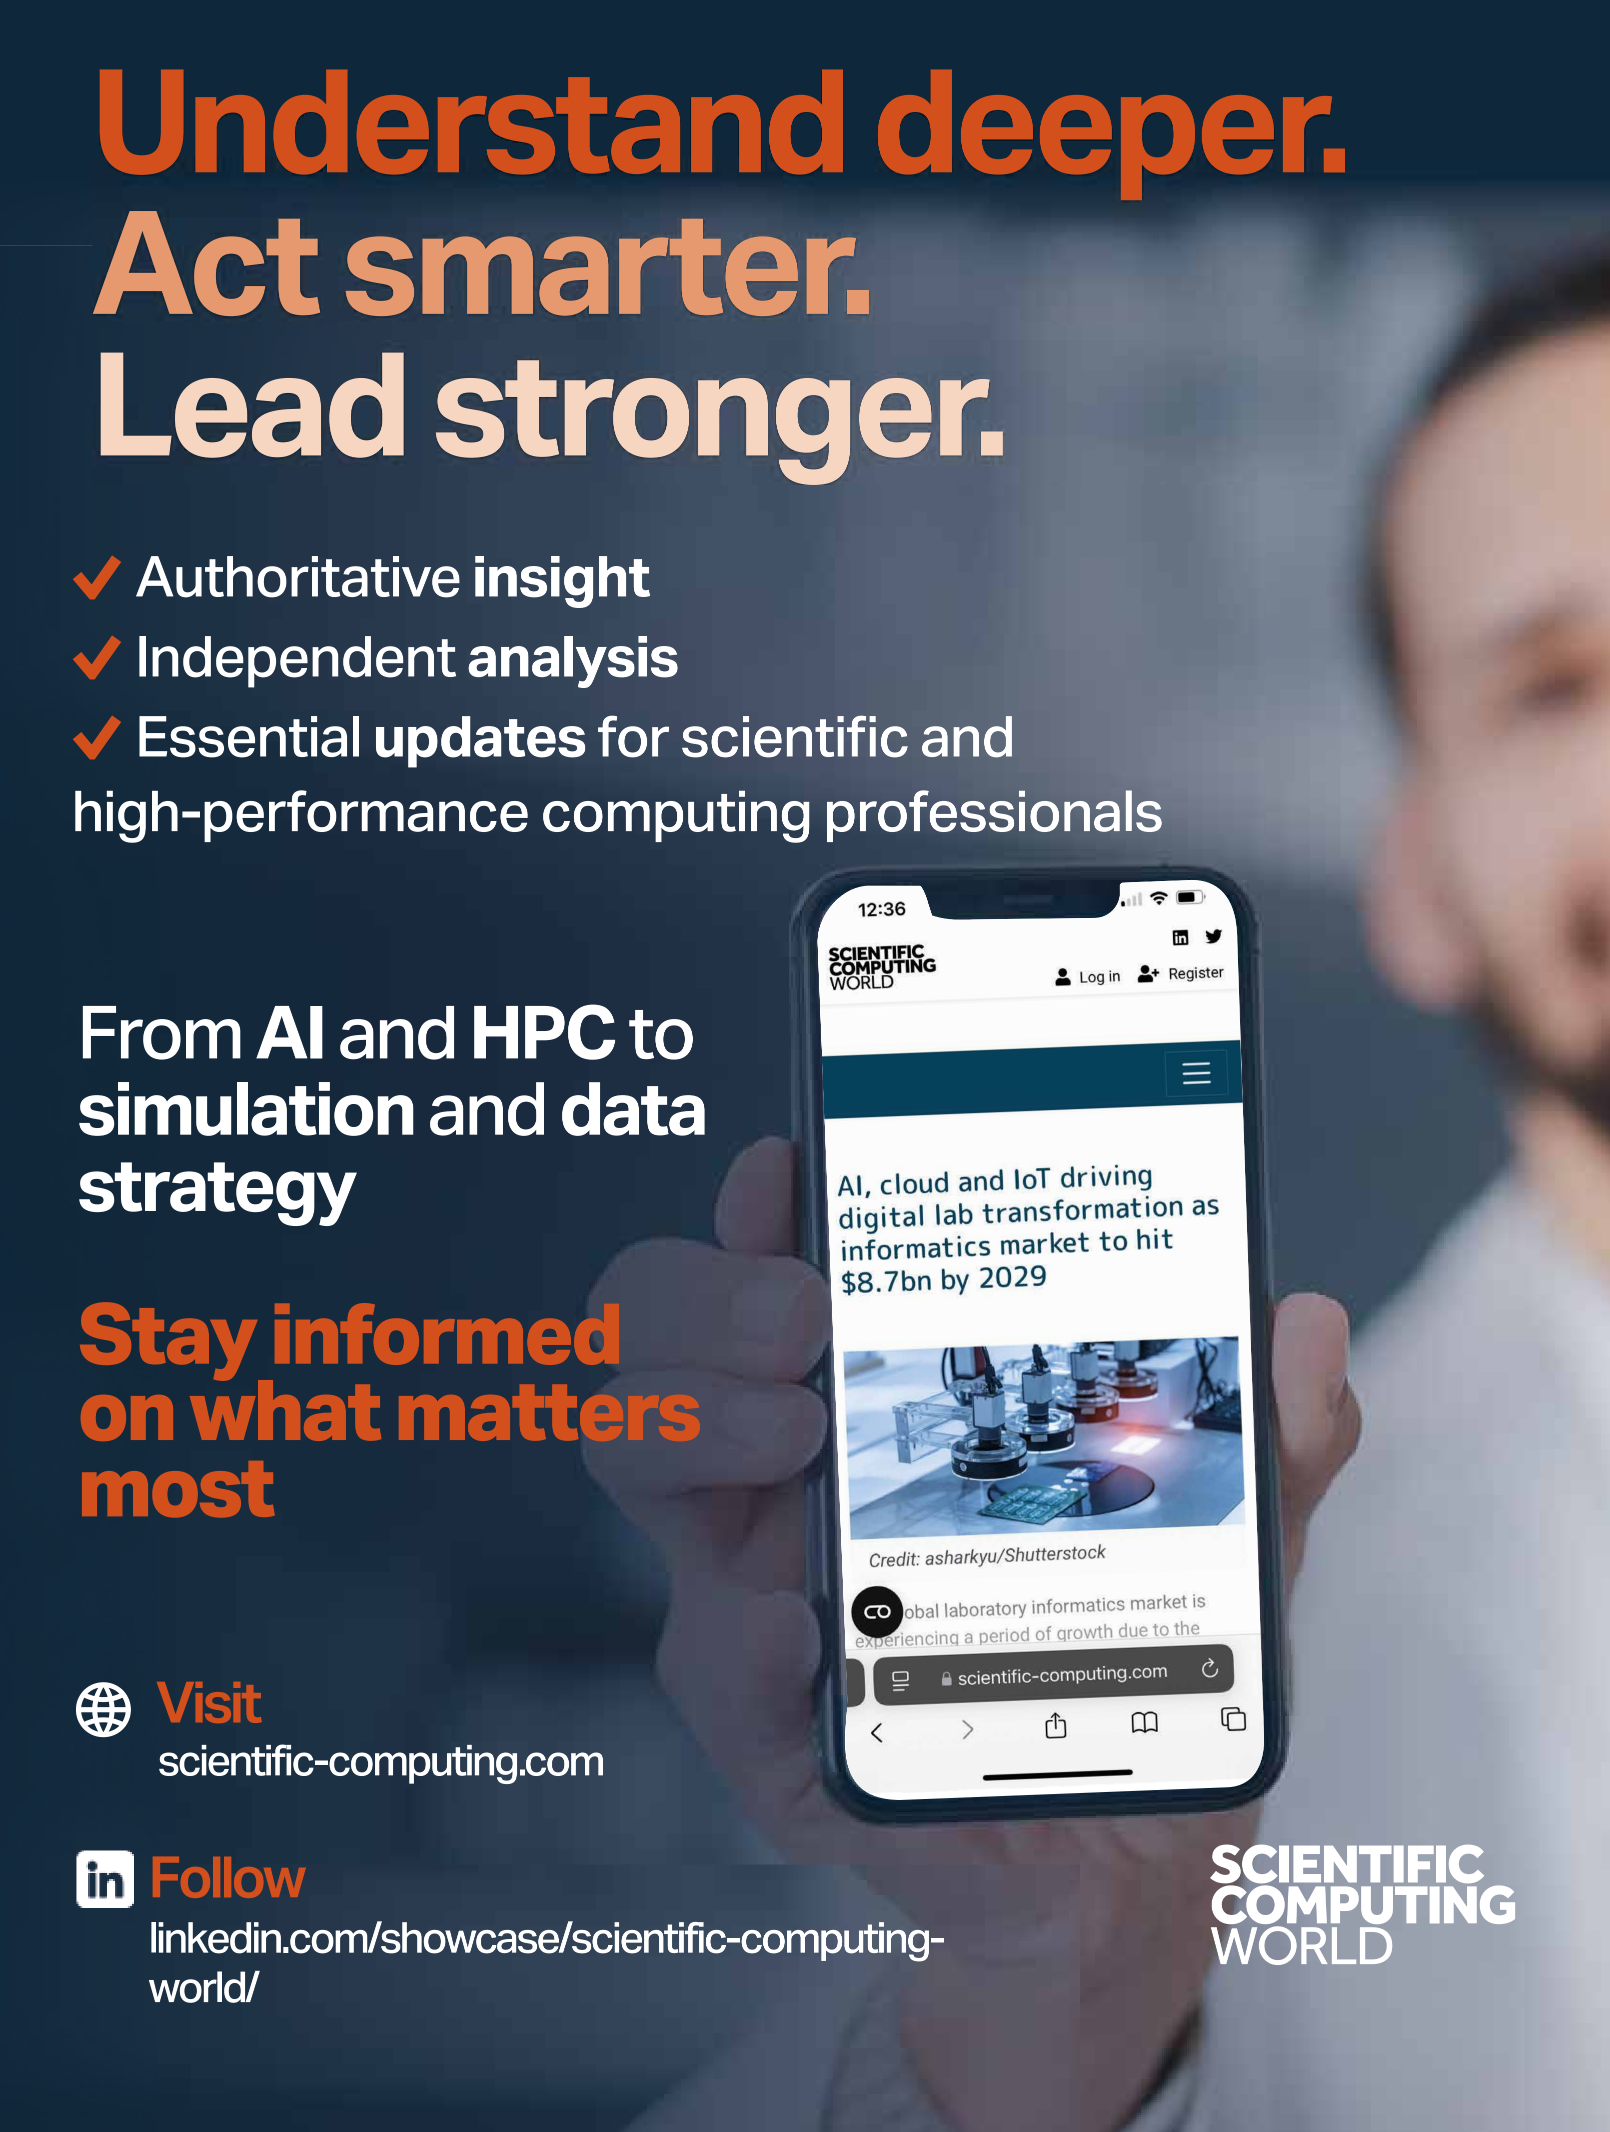
Understand deeper.
Act smarter.
Lead stronger.
Authoritative insight
Independent analysis
Essential updates for scientific and high-performance computing professionals
From AI and HPC to simulation and data strategy
Stay informed on what matters most
12:36
SCIENTIFIC COMPUTING WORLD
Log in
Register
AI, cloud and IoT driving digital lab transformation as informatics market to hit $8.7bn by 2029
Credit: asharkyu/Shutterstock
obal laboratory informatics market is experiencing a period of growth due to the
scientific-computing.com
Visit
scientific-computing.com
Follow
linkedin.com/showcase/scientific-computing-world/
SCIENTIFIC COMPUTING WORLD

Ivan Krajinović

Organisation: Rimac Technology
Role: Director of Central Engineering
Based in: Croatia

As an engineering leader, Ivan Krajinović operates at the point where simulation, testing and product development converge. His work at Rimac Technology focuses on scaling computational capabilities across teams, enabling faster iteration and more efficient engineering processes.

Krajinović is driving the integration of high-performance computing into everyday engineering workflows. By embedding GPU-accelerated simulation into development processes, he is helping to shift automotive R&D toward faster, more data-driven decision-making.

His career is defined by scale and integration. With more than a decade of experience across industry and academia, Krajinović has consistently focused on aligning simulation, testing and product development into cohesive engineering workflows. Today, as Director of Central Engineering at Rimac, he leads multidisciplinary teams at the forefront of electric vehicle and high-performance powertrain development.

His work centres on embedding advanced computational engineering into product lifecycles. Under his leadership, teams spanning simulations, testing, functional safety and cybersecurity collaborate to deliver complex systems to global automotive partners including Bugatti, Porsche, BMW and Hyundai. These collaborations reflect the growing importance of simulation-driven design in modern vehicle engineering.

Krajinović's career began with a strong foundation in engineering and research, followed by a progression into leadership roles where he combined technical expertise with team development. In a previous role, he scaled a simulation-focused engineering team from five to 35 specialists, introducing advanced numerical methods and establishing structured development processes. This experience shaped his approach to building high-performing, knowledge-driven teams.

At Rimac, he has contributed to the development of landmark electric hypercars, including the Rimac Concept One and the Rimac Nevera, vehicles widely recognised for pushing the boundaries of performance and electrification. These projects required tight integration between simulation, prototyping and testing, reinforcing the importance of computational methods in accelerating innovation.

Currently, Krajinović is focused on integrating GPU-based high-performance computing into everyday engineering workflows. The goal is to significantly increase simulation speed while reducing dependency on costly software licences. By enabling faster iteration cycles, this approach supports more efficient product development and enhances competitiveness in a rapidly-evolving industry.

Central to this effort is the development of a seamless feedback loop between modelling, simulation, prototyping and testing. While computational methods have advanced rapidly, Krajinović identifies verification as a key challenge. Ensuring that simulation results accurately reflect physical behaviour requires rigorous testing and validation, particularly in safety-critical applications such as automotive systems.

Looking ahead, his objective is to refine this feedback loop to the point where simulation and physical testing operate as a tightly coupled system. Achieving this would allow engineering teams to iterate more rapidly, reduce development costs and bring complex products to market with greater confidence.

Beyond technical innovation, Krajinović places strong emphasis on people and culture. He is known for fostering environments that encourage continuous learning, initiative and ownership, as well as for mentoring engineers and promoting STEM engagement.

His leadership approach reflects a belief that long-term success in scientific computing depends as much on developing talent as on advancing technology.

His advice to early-career engineers is characteristically direct: "Read, think, build, test." It is a philosophy rooted in practical experience, emphasising both intellectual curiosity and hands-on problem solving.

At its core, Krajinović's work represents the convergence of simulation, computation and industrial application. By integrating advanced HPC methods, alongside emerging approaches such as physics-informed neural networks and AI, into engineering practice, he is helping to define the next generation of automotive development.

Alistair Revell

Organisation: The University of Manchester
Role: Professor of Computational Engineering & Flow Physics, Director, Head of Department
Based in: Manchester, UK

From aerospace engineering to healthcare and elite sport, Alistair Revell's career has been defined by applying high-fidelity computational methods to problems where accuracy and impact matter most.

It is a defining feature of a career shaped by both technical depth and practical impact. "The thread has always been bringing high-fidelity simulation out of the lab and into real decisions," he explains, reflecting on work that spans aerospace engineering, bio-inspired design and clinical modelling.

Prof Revell's research sits at the intersection of computational engineering, turbulence modelling and fluid–structure interaction. His work focuses on developing advanced simulation frameworks that can operate across scales and levels of fidelity, enabling engineers and clinicians to better understand complex flow systems.

His path into the field began in industry, when he worked as a CAD technician before training as an aerospace engineer specialising in computational fluid dynamics. Early experiences in France, as an ERASMUS student in Poitiers and later as a Marie Curie Fellow at IMFT Toulouse, were followed by a role in Paris contributing to the development of code_saturne, EDF's open-source CFD code. This immersion in industrial-scale simulation proved formative, shaping his long-term focus on practical, scalable computational tools.

After a period at the Stanford University Centre for Turbulence Research, Revell joined the UK's University of Manchester in 2007, where he established a research group that has since grown into a leading hub for computational engineering. Early collaborations with aerospace and defence partners, including Airbus and Rolls-Royce, focused on external aerodynamics and turbulence modelling, later expanding into bio-inspired engineering and surface transport applications.

A significant pivot came in 2013, following an exchange at Massachusetts Institute of Technology, where he began applying computational fluid dynamics to cardiovascular flows. What started as work on congenital heart disease has evolved into AI-driven prediction of cardiac events, culminating in him co-leading computational efforts within a major British Heart Foundation Centre for Research Excellence.

Alongside scientific work, Revell has taken on substantial leadership responsibilities. From Deputy Head of School to founding Head of Department for

Mechanical & Aerospace Engineering, he has overseen rapid organisational growth, recruiting new academic staff and shaping institutional strategy. He also directs Manchester's Modelling & Simulation Centre, a partnership with EDF.

Today, his research spans three interconnected areas.

The first is the development of multi-fidelity simulation frameworks for exascale computing, coupling finite-volume and lattice Boltzmann methods to enable system-level simulations, such as modelling entire nuclear reactor cooling circuits.

The second focuses on integrating machine-learning into turbulence modelling, using techniques such as diffusion models and physics-informed neural networks to reconstruct complex flows efficiently.

The third applies these advances in real-world contexts, from patient-specific cardiovascular simulations to performance optimisation in elite sport.

Central to this work is a commitment to open-source tools and scalable computing. Revell's group makes extensive use of platforms such as code_saturne and OpenFOAM, alongside in-house solvers, deploying them across high-performance computing systems ranging from local clusters to national facilities such as ARCHER2. His contributions include developing coupled multi-fidelity libraries, advancing GPU-accelerated simulation techniques, and creating interactive CFD platforms used in industry and education.

Despite rapid progress, Revell identifies a key challenge: translating advances in computational power into tools that are accessible, reliable and usable in practice. While exascale systems enable unprecedented simulation detail, workflows remain complex and data-intensive, limiting their adoption by engineers and clinicians. Addressing this requires improved automation, robust uncertainty quantification and reduced-order models that can support decision-making without requiring specialist expertise.

He also highlights the need for critical engagement with emerging technologies such as artificial intelligence. As capabilities expand, ensuring that engineers retain a strong grounding in fundamental principles is essential. Initiatives such as Manchester's MaSC portal, a dedicated digital makerspace for students, aim to provide hands-on experience with computational tools in a controlled, exploratory environment.

Looking ahead, Revell hopes to accelerate the adoption of GPU-based, high-fidelity simulation in industry, enabling a better balance between accuracy and computational cost.

He sees particular potential for impact in nuclear engineering, clinical applications and elite sports performance.

For early-career researchers his advice is clear: focus on problems that matter beyond your discipline. "The most impactful computational work I have done started with a conversation with a clinician, an industrialist, or a coach – not with a journal paper. Build genuine partnerships where you learn their language and constraints, not just apply your methods to their data."

With a strong record of publications, extensive research funding and the supervision of more than 25 doctoral students, Revell considers his greatest contribution to be the people he has trained. In a field driven by complexity and collaboration, that investment in future talent may prove his most enduring legacy.

Ryan Slabaugh

Organisation: Boeing
Role: Executive Director, Boeing Emulation, Simulation, Test and Analysis Framework
Based in: United States

Ryan Slabaugh works on the engineering systems that help aerospace teams connect virtual design with physical reality. His role focuses on bringing emulation, simulation, testing and analysis into a more integrated framework for the development of complex aircraft and systems.

As Executive Director for Boeing's Emulation, Simulation, Test and Analysis Framework, Slabaugh leads work that supports one of aerospace engineering's most important shifts: using computational methods not only to explore designs, but to validate systems, reduce risk and improve confidence before physical testing is complete.

In an industry where safety, reliability and certification are critical, simulation must be accurate, traceable and closely linked to real-world test evidence.

His work reflects the increasing importance of computational engineering across aerospace. Emulation and simulation environments allow engineers to model aircraft systems, test behaviours under different conditions and analyse performance earlier in the development cycle. When effectively connected to physical test data, these tools can shorten feedback loops, improve decision-making and reduce reliance on costly late-stage redesign.

Matt Smolka

Organisation: Lotus
Role: Principal Engineer - Advanced Simulation at Group Lotus
Based in: UK

Matt Smolka works on advanced simulation in automotive engineering, where computational models are used to improve safety, performance and development efficiency before physical prototypes are constructed.

As Principal Engineer – Advanced Simulation at Lotus, Smolka's work focuses on CAE workflows for vehicle development, particularly crash and safety simulation. Public material from Lotus Engineering shows his role in advanced simulation.

This work highlights the use of LS-DYNA and Oasys tools across crash, NVH, durability and CFD workflows, with simulation used throughout design iteration and reporting processes.

His career has included roles in crash and safety engineering, CAE and advanced simulation, with public profiles listing experience at Group Lotus, TECOSIM Group, Penso and Tata Technologies.

These roles reflect a career built around applying simulation to demanding automotive engineering problems, from structural optimisation to crashworthiness and safety analysis.

Smolka has also contributed to industry discussions on artificial intelligence in engineering, appearing as a panellist discussing the use of AI in simulation and engineering workflows.

This reflects the growing role of data-driven methods alongside established CAE tools in automotive development.

His work demonstrates how computational engineering is becoming central to vehicle design: not simply validating designs after the fact, but helping engineers explore options, reduce physical testing and arrive at faster, more informed decisions.

Muhammad Usman

Organisation: Ghulam Ishaq Khan Institute of Engineering Sciences and Technology
Role: Associate Professor, Head of the Quantum Engineering Lab
Based in: Pakistan

Dr Muhammad Usman's work focuses on pushing the limits of semiconductor optoelectronics, where advances in materials and device design underpin the next generation of photonic technologies.

Operating at the forefront of III–V semiconductor research, Usman combines physics-based modelling with emerging AI methods to improve the performance of LEDs, laser diodes high electron mobility transistors (HEMTs), and microLED display technologies. His research spans device engineering, numerical simulation and AI-driven optimisation, with a focus on improving efficiency across demanding spectral ranges.

Dr Usman is an Associate Professor at the Ghulam Ishaq Khan Institute of Engineering Sciences and Technology (GIKI), Pakistan, where he also heads the Quantum Engineering Lab and the Photolithography Lab. He was named a 2024 Top Scholar by ScholarGPS and received the 2025 Physica B Award for Condensed Matter Physics.

Usman's academic journey began with a PhD in Semiconductor Photonics at Hanyang University, where he carried out pioneering work on GaN-based LEDs. He later joined the Ghulam Ishaq Khan Institute of Engineering Sciences and Technology, where he now serves as Associate Professor and leads the Quantum Engineering Lab.

There, he has developed a strong research ecosystem, producing more than 100 peer-reviewed publications and mentoring a new generation of researchers, many of whom have gone on to pursue fully-funded doctoral studies internationally.

His research focuses on next-generation optoelectronic and quantum semiconductor devices, particularly in challenging spectral regions, such as the green and deep-ultraviolet. These regimes are notoriously difficult due to instability in device physics, including efficiency loss, carrier leakage and material defects. By developing quantum-engineered active regions, graded barrier structures and advanced transport mechanisms, his work aims to overcome these long-standing limitations.

Central to this effort is the use of advanced scientific computing. Usman employs state-of-the-art semiconductor simulation platforms, including APSYS and SiLENSe®, to model complex physical processes such as carrier transport, recombination dynamics and thermal effects. These tools allow for precise virtual prototyping of device architectures before experimental validation, significantly accelerating the design cycle.

More recently, he has integrated artificial intelligence (AI) into his research workflows. By combining machine-learning and generative modelling with traditional physics-based simulations, he is able to explore large design spaces, optimise device parameters and address challenges related to data scarcity. This hybrid approach has already had impact in areas such as deep-ultraviolet laser diode design, where AI-assisted methods have enabled faster and more reliable identification of optimal structures.

Despite these advances, Usman identifies a fundamental challenge in the field: the gap between simulation and experimental reality. In complex semiconductor systems, uncertainties in material parameters and limitations in physical models can lead to discrepancies between predicted and observed behaviour. This challenge is particularly acute in extreme operating conditions, such as high-field or quantum-confined environments.

To address this, his work focuses on developing integrated frameworks that combine physics-based modelling with machine-learning corrections. By improving parameter accuracy and reducing computational cost, these approaches aim to make simulations more predictive and reproducible – an essential step towards accelerating device innovation.

Looking ahead, Usman's goal is to establish a validated, physics-informed machine-learning framework that aligns simulation with experimental outcomes. Achieving this would represent a significant milestone, enabling faster discovery and optimisation of high-efficiency optoelectronic devices across a range of applications.

Beyond his research, he has played an important role in advancing scientific capacity, contributing to professional development initiatives, organising international conferences and mentoring emerging researchers. His work reflects a commitment not only to technical excellence, but also to building a strong, collaborative research community.

His advice to early-career researchers is simple, but powerful: remain focused on your passion and avoid becoming overly constrained by metrics. In a field driven by complexity and innovation, persistence and curiosity remain key to meaningful impact.

At its core, Dr Usman's work demonstrates how scientific computing is transforming semiconductor engineering. By combining advanced simulation with AI-driven optimisation, he is helping to overcome long-standing barriers in device performance and paving the way for the next generation of optoelectronic technologies.

Martin Wearing

Organisation: European Space Agency
Role: Digital Twin Earth Scientist, ESA Science Hub
Based in: Italy

Martin Wearing works on digital replicas of the Earth system, using Earth observation data, numerical modelling and artificial intelligence to improve understanding of environmental change.

As a Digital Twin Earth Scientist at the European Space Agency's Centre for Earth Observation, ESRIN, Wearing helps coordinate the development of computational replicas of elements of the Earth system.

This work brings together satellite observations, modelling workflows and AI-enabled analysis to create tools that can represent environmental change with greater detail and responsiveness.

His background in numerical ice-sheet modelling and satellite Earth observation gives him a strong foundation in the computational challenges of climate science.

In particular, work on ice-sheet dynamics and sea-level rise requires models that can integrate sparse observations, physical processes and long-term uncertainty.

Wearing's work demonstrates how digital twins can combine simulation, satellite data and artificial intelligence to support a better understanding of complex Earth systems.

These approaches can help researchers explore future scenarios, test assumptions and provide evidence for climate resilience and environmental decision-making.

Karen Willcox

Organisation: University of Texas
Role: Director, Oden Institute for Computational Engineering and Sciences
Based in: United States

Karen Willcox is a leading figure in computational engineering, known for developing methods that make large-scale simulation faster, more reliable and more useful for decision-making.

As Director of the Oden Institute for Computational Engineering and Sciences at the University of Texas at Austin, she leads one of the world's major centres for computational science and engineering. Her work connects applied mathematics, engineering simulation and data-driven modelling, with a focus on making complex computational models practical for real-world design and analysis.

Willcox's research focuses on reduced-order modelling, multifidelity methods, scientific machine-learning and simulation-based engineering. These approaches allow researchers to approximate complex systems more efficiently, explore design options, quantify uncertainty and make decisions without relying solely on expensive full-scale simulations.

Her work has had particular influence in aerospace, energy and complex engineered systems, where simulation must balance accuracy, cost and speed. By developing methods that make computational models more efficient and trustworthy, Willcox has helped shape how engineers use simulation to guide high-consequence decisions.

John Ziadat

Organisation: Stoke Space Technologies
Role: Engineering Lead, Stoke Analysis
Based in: United States

John Ziadat works on the engineering analysis needed to develop reusable launch systems, using simulation to understand how spaceflight hardware behaves under extreme loads.

As Engineering Lead for Stoke Analysis at Stoke Space Technologies, Ziadat supports the development of reusable rocket systems that must withstand intense thermal, pressure, inertial and aerodynamic forces.

His work focuses on applying computational analysis to evaluate how components perform during launch, flight, re-entry and repeated use.

His background in mechanical engineering and computational solid mechanics, including numerical simulation of ballistic impact, aligns closely with the demands of launch vehicle development. These systems require engineers to understand not only whether the hardware survives a single mission, but also how it behaves across multiple operating cycles and under high-stress conditions.

Ziadat's work reflects the growing role of computational engineering in the commercial space sector. Simulation allows teams to assess design options, reduce risk and build confidence before expensive physical testing or flight, helping accelerate development while maintaining reliability.