TROY TECHNICAL

MATERIALS INFORMATICS WEEKLY REPORT

November 14 - November 20, 2026 | Issue 2026-W46

Closed-loop learning from data quality to experimentally validated material performance

12	10
ANALYZED ARTICLES	PRIORITY THEMES

Issued 2026-11-20 | English edition | fixed approved cover artwork



Materials Informatics: the boundary moved from isolated claims to qualification evidence

12 English articles were reviewed independently from the Japanese edition. Ten themes are retained for management decisions.

SIGNAL 01 | DATA CAPTURE

Wireless Optical Chips: Real-time Reaction Monitoring for 2026

Researchers have developed 0.5mm wireless optoelectronic integrated circuits capable of monitoring reaction vessel temperatures with 0.1°C precision every half second.

SIGNAL 02 | MODELING

NVIDIA DAQIRI: Real-Time GPU Data Acquisition for Scientific Instruments

NVIDIA DAQIRI enables the direct transfer of high-bandwidth sensor data into GPU memory, bypassing disk writes to allow real-time AI processing.

SIGNAL 03 | CANDIDATE DESIGN

NVIDIA Nemotron for Science: AI Agents for Scientific Discovery

NVIDIA has introduced Nemotron for Science, combining the Nemotron 3.5 Super VL open model with specialized chemical datasets and post-training recipes to accelerate scientific discovery.

SIGNAL 04 | EXPERIMENT

Magnetic Loss Analysis: TUS AI Approach for High-Frequency Materials

Researchers at the Tokyo University of Science developed a physics-interpretable AI method to identify microscopic structures causing energy loss in high-frequency magnetic materials.

DECISION

Focus this week's management attention on closed-loop learning from data quality to experimentally validated material performance. Move only when the linked evidence can be reproduced under your own operating conditions.

THREE MANAGEMENT MOVES

1. Buyers: select high-value decisions where data lineage and experimental feedback can be measured.
2. Suppliers: provide machine-readable composition, process, test conditions, uncertainty and failure data.
3. Executives: track dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EVIDENCE DISCIPLINE

FACT states what the collected English source reports. BUSINESS READ and ACTION are editorial interpretations. UNKNOWN lists information that the source file does not establish. A linked article is not treated as independent confirmation of every claim.

Numbers and evidence anchors that require conditions

0.5mm
A02

SOURCE-BOUND CONDITION
Researchers have developed 0.5mm wireless optoelectronic integrated circuits capable of monitoring reaction vessel temperatures with 0.1°C precision every half second.

500 G
A12

SOURCE-BOUND CONDITION
A case study with Lawrence Berkeley National Laboratory demonstrates the ability to process over 500 Gbps of electron microscopy data for immediate experimental adjustment.

2075 cyc les
A01

SOURCE-BOUND CONDITION
This architecture achieves a discharge capacity of ~60,000 mAh g-1 and a stable lifespan of 2075 cycles.

89.33%
A04

SOURCE-BOUND CONDITION
It achieves F1 scores of 89.33% for property extraction and 71.35% for composition extraction, processing data 6-496× faster than contemporary frontier LLMs.

22x
A06

SOURCE-BOUND CONDITION
For projected OMat24 workloads, Daviator-enabled VASP can achieve up to 22x faster time-to-solution and 75% lower costs compared to CPU execution.

HOW TO USE THESE NUMBERS
Each value is shown with the condition stated in the collected English article. Do not compare values across different test methods, device sizes, operating temperatures or commercial stages without normalizing those conditions.

Where the value chain moved

Each stage combines one leading theme and one supporting theme. The report treats handoffs as evidence transfers, not decorative arrows.

1

Data capture

P01/P06

Researchers have developed 0.5mm wireless optoelectronic integrated circuits capable of monitoring reaction vessel.

2

Modeling

P02/P07

NVIDIA DAQIRI enables the direct transfer of high-bandwidth sensor data into GPU memory, bypassing disk writes to allow.

3

Candidate design

P03/P08

NVIDIA has introduced Nemotron for Science, combining the Nemotron 3.5 Super VL open model with specialized chemical datasets.

4

Experiment

P04/P09

Researchers at the Tokyo University of Science developed a physics-interpretable AI method to identify microscopic structures.

5

Scale & governance

P05/P10

An AI-integrated framework leveraging machine learning and density functional theory has identified Ni-Co dual-atom catalysts.

EVIDENCE REQUIRED AT EACH HANDOFF

Data capture to Modeling	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.
Modeling to Candidate design	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.
Candidate design to Experiment	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.
Experiment to Scale & governance	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

COMMON CONTROL POINT

No theme moves to a commercial commitment until a named owner has reproduced the relevant evidence and documented dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

Priority map: maturity and next decision locus



P01	Wireless Optical Chips: Real-time Reaction Monitoring for 2026 System / adoption | cue: integrated
P02	NVIDIA DAQIRI: Real-Time GPU Data Acquisition for Scientific Instruments Business / rules | cue: Acquisition
P03	NVIDIA Nemotron for Science: AI Agents for Scientific Discovery System / adoption | cue: deployable
P04	Magnetic Loss Analysis: TUS AI Approach for High-Frequency Materials Tech / process | cue: no system/business cue
P05	Li-O2 Batteries: AI-Driven Ni-Co Dual-Atom Catalyst for LiOH Pathway Business / rules | cue: capacity
P06	MatSKRAFT: Large-scale materials knowledge extraction from tables System / adoption | cue: system
P07	MACE-MP-0: ML Model for Lithium Battery SEI Reaction Simulation Tech / process | cue: no system/business cue
P08	Daviator: NVIDIA's GPU acceleration for plane-wave DFT in VASP Tech / process | cue: no system/business cue
P09	NVIDIA cuPotential: Accelerating Molecular Dynamics Simulations Tech / process | cue: no system/business cue
P10	AlphaDiffract: Deep Learning for PXRD Lattice Determination 2026 System / adoption | cue: system

REPRODUCIBLE POSITIONING RULE

Horizontal position uses five discrete stages: Research, Prototype, Joint evaluation, Product adoption and Scale. The vertical bands classify the next decision locus: Tech / process, System / adoption, or Business / rules (including supply systems). Article-title cues take precedence; the overview is used only when the title is silent. These are categories, not impact ranks. Bubble diameter reflects the editorial score.

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P01 Wireless Optical Chips: Real-time Reaction Monitoring for 2026

A02 | Data capture | Joint evaluation

2026-11-20

FACT

Researchers have developed 0.5mm wireless optoelectronic integrated circuits capable of monitoring reaction vessel temperatures with 0.1°C precision every half second. These commercially available chips provide high-density data for self-driving labs and industrial safety monitoring by converting heat outputs into reaction rates.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P01.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 20/25 | MATURITY 3/5 | DECISION LOCUS: SYSTEM / ADOPTION

P02 NVIDIA DAQIRI: Real-Time GPU Data Acquisition for Scientific Instruments

A12 | Modeling | Product adoption

2026-11-19

FACT

NVIDIA DAQIRI enables the direct transfer of high-bandwidth sensor data into GPU memory, bypassing disk writes to allow real-time AI processing. A case study with Lawrence Berkeley National Laboratory demonstrates the ability to process over 500 Gbps of electron microscopy data for immediate experimental adjustment.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P02.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P03 NVIDIA Nemotron for Science: AI Agents for Scientific Discovery

A05 | Candidate design | Pilot

2026-11-18

FACT

NVIDIA has introduced Nemotron for Science, combining the Nemotron 3.5 Super VL open model with specialized chemical datasets and post-training recipes to accelerate scientific discovery. The framework enables researchers to build domain-specific agents for tasks such as protein design and chemical synthesis, deployable from desktops to supercomputing centers.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P03.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 20/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

P04 Magnetic Loss Analysis: TUS AI Approach for High-Frequency Materials

A10 | Experiment | Pilot

2026-11-17

FACT

Researchers at the Tokyo University of Science developed a physics-interpretable AI method to identify microscopic structures causing energy loss in high-frequency magnetic materials. By combining supercomputing with topological machine learning, the team pinpointed specific domain-wall features that contribute most to energy barriers. This framework provides a pathway to designing more efficient power conversion components for electric vehicles and data centers.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P04.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: TECH / PROCESS

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P05 Li-O2 Batteries: AI-Driven Ni-Co Dual-Atom Catalyst for LiOH Pathway

A01 | Scale & governance | Product adoption

2026-11-18

FACT

An AI-integrated framework leveraging machine learning and density functional theory has identified Ni-Co dual-atom catalysts that shift lithium-oxygen battery reactions from a two-electron to a four-electron LiOH pathway. This architecture achieves a discharge capacity of ~60,000 mAh g-1 and a stable lifespan of 2075 cycles. The result demonstrates a data-guided strategy to overcome traditional kinetic bottlenecks in high-energy-density energy storage.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P05.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

P06 MatSKRAFT: Large-scale materials knowledge extraction from tables

A04 | Data capture | Research

2026-11-16

FACT

MatSKRAFT is a computational framework that automatically extracts materials science knowledge from scientific tables using constraint-driven GNNs. It achieves F1 scores of 89.33% for property extraction and 71.35% for composition extraction, processing data 6-496× faster than contemporary frontier LLMs. This system has enabled the construction of a database with 509,281 entries from over 45,500 research publications.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P06.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P07 MACE-MP-0: ML Model for Lithium Battery SEI Reaction Simulation

A11 | Modeling | Pilot

2026-11-16

FACT

Researchers have fine-tuned the MACE-MP-0 machine learning foundation model to simulate the formation and degradation of the solid-electrolyte interphase (SEI) in lithium metal batteries. The fine-tuned R4 potential reduced force root mean square error to 52 millielectronvolts per angstrom for electrolyte structures, enabling nanosecond-scale simulations of interface chemical dynamics.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P07.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: TECH / PROCESS

P08 Daviator: NVIDIA's GPU acceleration for plane-wave DFT in VASP

A06 | Candidate design | Research

2026-11-18

FACT

NVIDIA has introduced Daviator, a GPU kernel library that eliminates performance bottlenecks in plane-wave DFT by replacing sequential k-point processing with batched operations. For projected OMat24 workloads, Daviator-enabled VASP can achieve up to 22x faster time-to-solution and 75% lower costs compared to CPU execution. This optimization is critical for generating the massive DFT datasets required to train next-generation machine learning interatomic potentials (MLIP).

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P08.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 1/5 | DECISION LOCUS: TECH / PROCESS

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P09 NVIDIA cuPotential: Accelerating Molecular Dynamics Simulations

A07 | Experiment | Pilot

2026-11-16

FACT

NVIDIA cuPotential provides modular C API primitives that achieve end-to-end speedups of 2x to over 3x for systems with 24,000 to 400,000 atoms. By optimizing long-range interaction evaluation under periodic boundary conditions, the library significantly reduces energy drift and improves dispersion energy accuracy.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P09.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: TECH / PROCESS

P10 AlphaDiffract: Deep Learning for PXRD Lattice Determination 2026

A03 | Scale & governance | Research

2026-11-15

FACT

AlphaDiffract is a deep learning framework that jointly predicts crystal systems, space groups, and lattice parameters directly from powder X-ray diffraction (PXRD) data. Trained on a dataset of over 31 million simulated patterns, it achieves 81.7% crystal system accuracy and 66.2% space group accuracy on the RRUFF dataset. This unified model enables rapid, high-throughput materials discovery by reducing the need for iterative, expert-led crystallographic analysis.

WHY IT MATTERS

The decision relevance is closed-loop learning from data quality to experimentally validated material performance. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, provide machine-readable composition, process, test conditions, uncertainty and failure data.

ACTION

Within 30 days, select high-value decisions where data lineage and experimental feedback can be measured; record the evidence and owner against P10.

UNKNOWN

Still unproven or incompletely stated: dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.

EDITORIAL SCORE 19/25 | MATURITY 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

Playbook and 0-5 year roadmap

STAKEHOLDER	NOW | 0-30 DAYS	NEXT | 1-12 MONTHS	LATER | 1-5 YEARS
Operating companies	select high-value decisions where data lineage and experimental feedback can be measured	Run production-like qualification with explicit acceptance and stop criteria.	Build a portfolio and architecture that can absorb technology and supplier changes.
Materials, equipment and technology suppliers	provide machine-readable composition, process, test conditions, uncertainty and failure data	Create shared reliability, process-window and failure-analysis data with customers.	Standardize interfaces, traceability, lifecycle support and recovery services.
Trading, investment and legal teams	Map dataset bias, reproducibility, IP, model transfer, experiment cost and adoption.	Contract for capacity, quality, change notice, liability, alternatives and exit.	Diversify regional, technical and customer concentration risk.

ROADMAP

NOW | 0-30 DAYS
Evidence ledger
Owner: Strategy + quality
Deliverable: claim, source, condition, owner

NOW | 0-30 DAYS
Risk screen
Owner: Procurement + legal
Deliverable: unknowns, alternatives, stop

NEXT | 1-12 MONTHS
Joint qualification
Owner: R&D + supplier
Deliverable: process window and failure data

NEXT | 1-12 MONTHS
Commercial gate
Owner: Business + finance
Deliverable: unit economics and capacity

LATER | 1-5 YEARS
Platform strategy
Owner: Executive team
Deliverable: portfolio and interface roadmap

LATER | 1-5 YEARS
Service layer
Owner: Operations + channel
Deliverable: monitoring, maintenance and recovery

GO / NO-GO GATES

GATE	OWNER	REQUIRED EVIDENCE	NO-GO CONDITION
G1 Evidence	**Quality**	Source, test conditions and reproducibility	No traceable evidence
G2 Integration	**R&D**	Interface, process window and failure analysis	Cannot meet system conditions
G3 Commercial	**Business**	Capacity, total cost, contract and alternatives	Economics or supply cannot be supported
G4 Lifecycle	**Operations**	Monitoring, maintenance, change notice and exit	No accountable operating plan

ANALYZED ARTICLES | ITEMS 01-12

Every title and URL is displayed in full. URLs wrap without shortening and remain clickable. A verified-reference label means the translated HTML did not retain its original source URL.

ID	FULL ARTICLE TITLE	FULL SOURCE / VERIFIED REFERENCE URL - CLICKABLE
A01	Li-O2 Batteries: AI-Driven Ni-Co Dual-Atom Catalyst for LiOH Pathway	https://doi.org/10.1038/s41467-026-79335-6
A02	Wireless Optical Chips: Real-time Reaction Monitoring for 2026	https://www.chemistryworld.com/news/sand-grain-sized-wireless-chip-can-monitor-reactions-to-provide-data-for-self-driving-labs/4024549.article
A03	AlphaDiffract: Deep Learning for PXRD Lattice Determination 2026	https://doi.org/10.1038/s41524-026-02388-0
A04	MatSKRAFT: Large-scale materials knowledge extraction from tables	https://doi.org/10.1038/s41524-026-02359-5
A05	NVIDIA Nemotron for Science: AI Agents for Scientific Discovery	https://blogs.nvidia.co.jp/blog/nvidia-introduces-nemotron-for-science-to-accelerate-discoveries-with-agents/
A06	Daviator: NVIDIA's GPU acceleration for plane-wave DFT in VASP	https://developer.nvidia.com/blog/tackling-gpu-bottlenecks-in-plane-wave-dft-with-k-point-batching-and-mixed-precision/
A07	NVIDIA cuPotential: Accelerating Molecular Dynamics Simulations	https://developer.nvidia.com/blog/accelerate-molecular-dynamics-simulations-with-accurate-flexible-primitives-using-nvidia-cupotential/
A08	All-Natural Plastic Substitutes: Robotics and ML Workflow 2026	https://doi.org/10.1038/s41596-026-01449-4
A09	MatOS: Open-Source Materials Science Agent for R&D Automation	http://eu.36kr.com/en/p/4073503237329796
A10	Magnetic Loss Analysis: TUS AI Approach for High-Frequency Materials	https://www.miragenews.com/how-ai-identifies-magnetic-structures-behind-1778948/
A11	MACE-MP-0: ML Model for Lithium Battery SEI Reaction Simulation	https://bioengineer.org/machine-learning-model-simulates-solid-electrolyte-interphase-reactions-in-lithium-batteries/
A12	NVIDIA DAQIRI: Real-Time GPU Data Acquisition for Scientific Instruments	https://developer.nvidia.com/blog/enabling-real-time-ai-for-scientific-instruments-with-gpu-native-data-acquisition/

MaterialsInformatics — Selected Articles

Date: 2026-11-21

Articles: 12

Table of Contents

#01 Li-O2 Batteries: AI-Driven Ni-Co Dual-Atom Catalyst for LiOH Pathway

#02 Wireless Optical Chips: Real-time Reaction Monitoring for 2026

#03 AlphaDiffract: Deep Learning for PXRD Lattice Determination 2026

#04 MatSKRAFT: Large-scale materials knowledge extraction from tables

#05 NVIDIA Nemotron for Science: AI Agents for Scientific Discovery

#06 Daviator: NVIDIA's GPU acceleration for plane-wave DFT in VASP

#07 NVIDIA cuPotential: Accelerating Molecular Dynamics Simulations

#08 All-Natural Plastic Substitutes: Robotics and ML Workflow 2026

#09 MatOS: Open-Source Materials Science Agent for R&D Automation

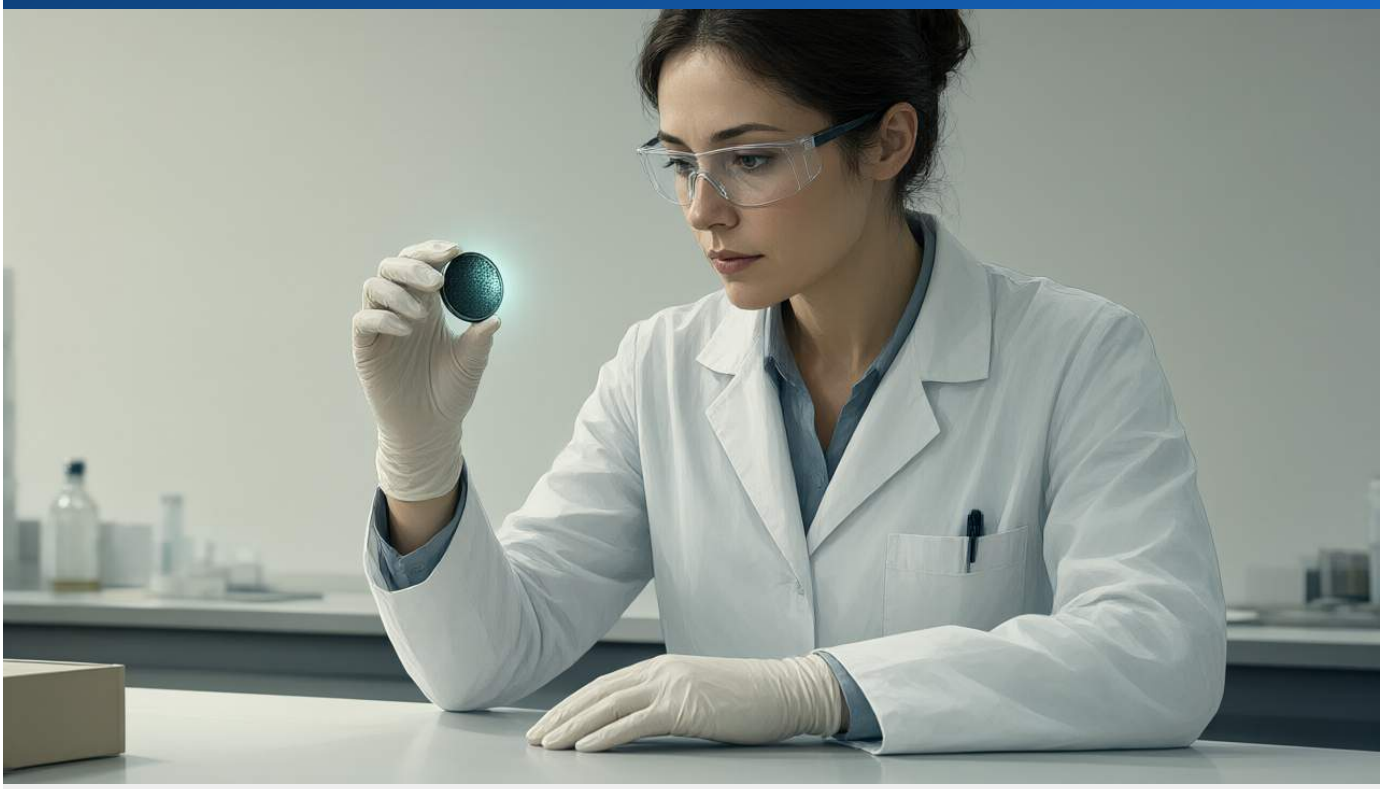
#10 Magnetic Loss Analysis: TUS AI Approach for High-Frequency Materials

#11 MACE-MP-0: ML Model for Lithium Battery SEI Reaction Simulation

#12 NVIDIA DAQIRI: Real-Time GPU Data Acquisition for Scientific Instruments

#01 Li-O2 Batteries: AI-Driven Ni-Co Dual-Atom Catalyst for LiOH Pathway

Published November 18, 2026 Nature Communications



OVERVIEW

An AI-integrated framework leveraging machine learning and density functional theory has identified Ni-Co dual-atom catalysts that shift lithium-oxygen battery reactions from a two-electron to a four-electron LiOH pathway. This architecture achieves a discharge capacity of ~60,000 mAh g−1 and a stable lifespan of 2075 cycles. The result demonstrates a data-guided strategy to overcome traditional kinetic bottlenecks in high-energy-density energy storage.

IN DEPTH

Key Findings

The integration of interpretable machine learning (ML) and density functional theory (DFT) has led to the discovery of Ni-Co dual-atom catalysts (DACs) that enable a controllable transition to a four-electron (4e−/O2) reaction pathway in lithium-oxygen batteries. This shift results in the formation of lithium hydroxide (LiOH) rather than the conventional lithium peroxide (Li2O2), yielding a discharge capacity of ~60,000 mAh g−1. Furthermore, the system demonstrates a stable lifespan of 2075 cycles with reduced voltage decay.

Technical Details

The AI framework navigated a chemical space of thousands of active sites by analyzing high-dimensional, non-linear correlations between physicochemical descriptors—such as valence electron configuration and electronegativity—and electrochemical performance. Theoretical analysis indicates that the synergistic electronic coupling between the Ni and Co dual-sites reconfigures the local d-band center, lowering the thermodynamic barriers for oxygen redox and optimizing proton-coupled electron transfer (PCET) kinetics.

Beyond its catalytic role at the positive electrode, the NiCo-DAC architecture serves as a multifunctional interfacial protective layer. This layer stabilizes the lithium metal negative electrode by suppressing the growth of dendrites.

Background & Context

Lithium-oxygen batteries are candidates for future electrification due to a high theoretical specific energy of ca. 3500 Wh kg−1. However, the conventional two-electron pathway (2Li + O2 + 2e− ↔ Li2O2) produces insulating and insoluble lithium peroxide, which passivates the electrode interface and causes severe overpotentials and poor cycle life. While shifting to a four-electron LiOH pathway offers better chemical stability and higher capacity, traditional single-metal active sites are limited by linear scaling relations that hinder the precise balance of protonated intermediates.

Significance & Outlook

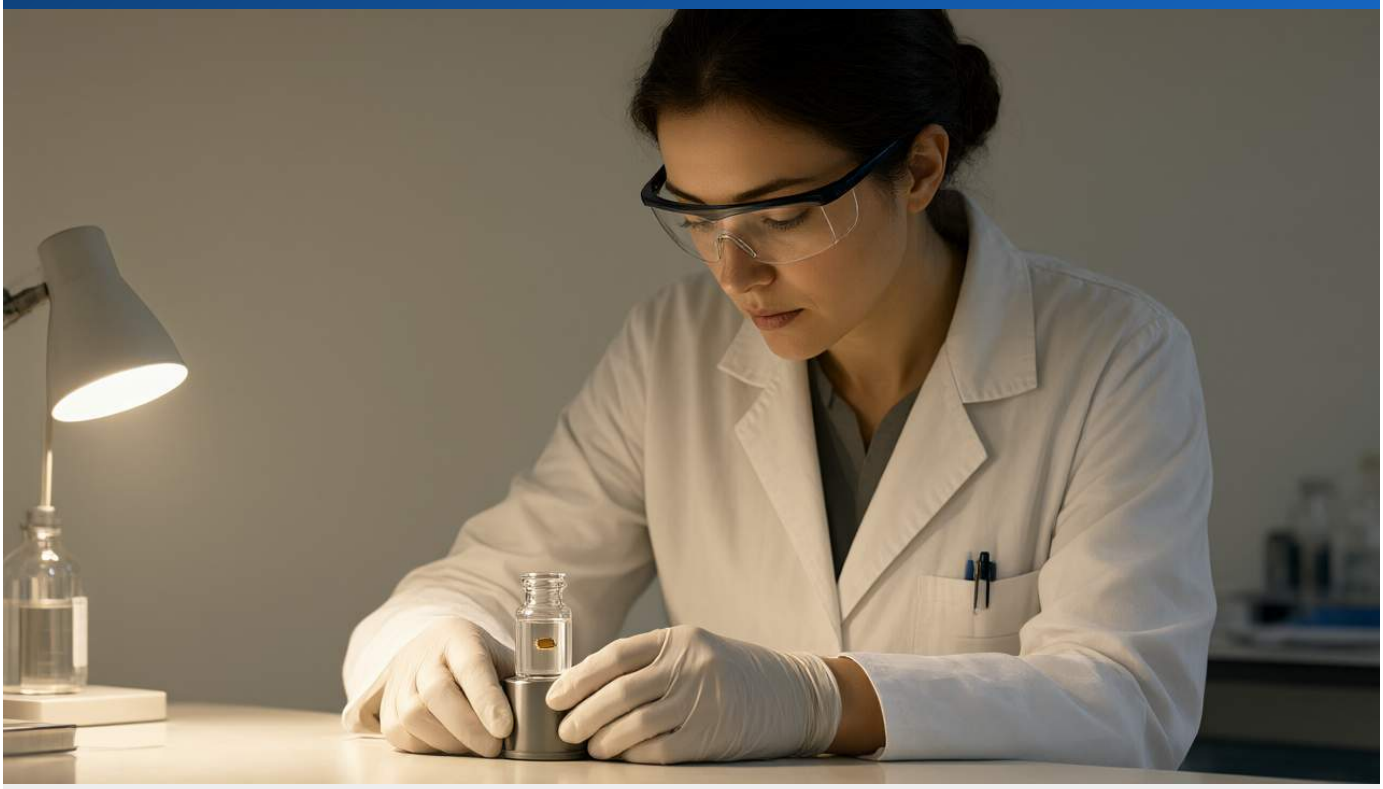
By expanding the design space from single atoms to multi-metallic configurations, this research breaks traditional catalytic bottlenecks. The use of a data-driven protocol replaces inefficient trial-and-error exploration, bridging the gap between atomic-scale theory and practical application for high-energy-density batteries.

Source: https://doi.org/10.1038/s41467-026-79335-6

Collected: November 21, 2026 | Automated Research System (AI)

#02 Wireless Optical Chips: Real-time Reaction Monitoring for 2026

Published November 20, 2026 Chemistry World USA



OVERVIEW

Researchers have developed 0.5mm wireless optoelectronic integrated circuits capable of monitoring reaction vessel temperatures with 0.1°C precision every half second. These commercially available chips provide high-density data for self-driving labs and industrial safety monitoring by converting heat outputs into reaction rates.

IN DEPTH

Key Findings

US researchers, including Song Lin of Stanford University and Paul McEuen of Cornell University, have developed submillimetre wireless optical chips that monitor reaction temperatures in real time. These 0.5mm optoelectronic integrated circuits can be placed inside or attached to the walls of test tubes, allowing for the simultaneous monitoring of up to 24 tubes. The system accumulates over 100,000 data points per hour, providing a high-resolution stream of information that can be converted into heat outputs and reaction rates.

Technical Details

The chips operate using a wireless optical mechanism:

- **Power Source:** An external LED emits light that is absorbed by a semiconductor photovoltaic on the chip and converted into electrical energy.
- **Signal Generation:** This energy pumps a laser that emits nanosecond optical pulses.
- **Temperature Sensing:** The repetition rate of these pulses depends on the temperature.
- **Data Acquisition:** An external photodetector detects the pulse rate, enabling temperature measurements within 0.1°C every half second.

Background and Context

Efficiently monitoring experimental reactions is traditionally difficult. Existing methods like NMR spectroscopy often require sample extraction for ex situ analysis, while in situ spectroscopy typically provides cumulative concentration measurements that are hard to interpret. While calorimetric measurements offer the most direct data on reaction rates via heat production or absorption, high-end calorimeters can cost over a hundred thousand dollars. The new chips were tested on various reactions, including Buchwald–Hartwig amination, Suzuki–Miyaura coupling, and the strongly exothermic cobalt-catalysed hydrolytic kinetic resolution of terminal epoxides.

Significance and Outlook

The technology is designed to integrate with self-driving labs, where robots perform parallel experiments but struggle to read results to inform subsequent steps. Donna Blackmond of the Scripps Research Institute notes that this system provides the high-quality experimental information necessary for sophisticated AI-driven discovery models. Additionally, the chips have significant implications for industrial process chemistry, specifically in safety monitoring to prevent reaction runaway in drug manufacturing.

Source: https://www.chemistryworld.com/news/sand-grain-sized-wireless-chip-can-monitor-reactions-to-provide-data-for-self-driving-labs/4024549.article

Collected: November 21, 2026 | Automated Research System (AI)

#03 AlphaDiffract: Deep Learning for PXRD Lattice Determination 2026

Published November 15, 2026 npj Computational Materials USA



OVERVIEW

AlphaDiffract is a deep learning framework that jointly predicts crystal systems, space groups, and lattice parameters directly from powder X-ray diffraction (PXRD) data. Trained on a dataset of over 31 million simulated patterns, it achieves 81.7% crystal system accuracy and 66.2% space group accuracy on the RRUFF dataset. This unified model enables rapid, high-throughput materials discovery by reducing the need for iterative, expert-led crystallographic analysis.

IN DEPTH

Key Findings

AlphaDiffract provides a single-shot method for the simultaneous prediction of crystal system, space group, and all six lattice parameters from experimental PXRD patterns. When validated against the RRUFF dataset, the model demonstrated strong transferability from simulated data to experimental measurements, achieving an accuracy of 81.7% for crystal systems and 66.2% for space groups.

Technical Details

The framework utilizes a 1D adaptation of the ConvNeXt architecture, a modern convolutional neural network that incorporates design principles from transformers. This architecture is coupled with dedicated prediction heads for each specific crystallographic property. To train the model, the researchers generated a physics-based dataset of over 31 million simulated diffraction patterns by augmenting 312,267 curated structures sourced from the Materials Project and ICSD databases.

Background & Context

Identifying materials and understanding their structures via PXRD is a fundamental challenge in materials science. Accurate determination of the crystal lattice—including symmetry and lattice parameters—is a prerequisite for full structure solution. Traditional analysis methods are typically iterative and rely heavily on expert input. While previous deep learning attempts have shown promise, a robust method capable of comprehensive lattice determination from experimental data has remained a primary goal for the field.

Significance & Outlook

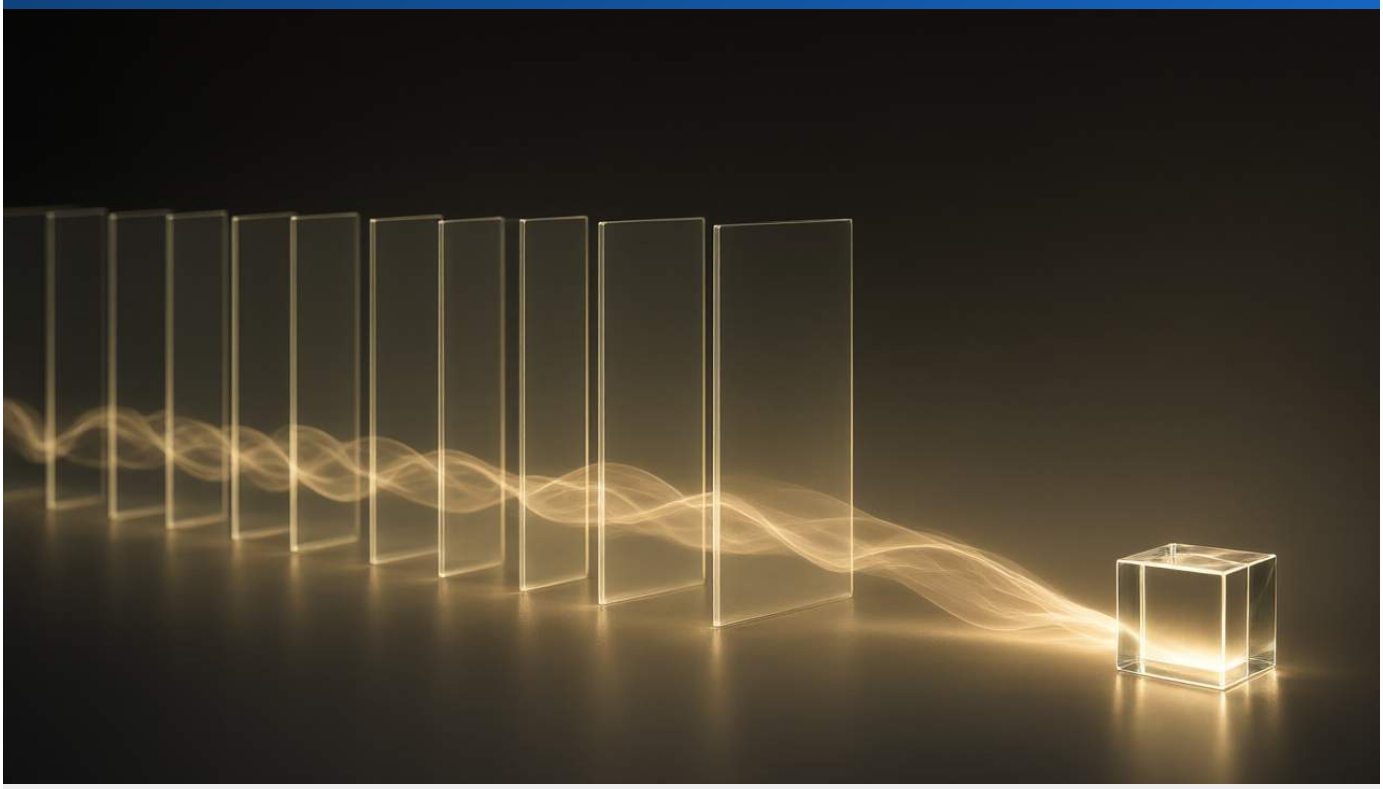
By automating the determination of lattice symmetry and parameters, AlphaDiffract represents a significant advancement in the use of deep learning for materials science. The ability to perform rapid, unified predictions directly from PXRD data facilitates high-throughput materials discovery and characterization. This research was supported by the U.S. DOE Office of Science-Basic Energy Sciences and utilized resources from the Advanced Photon Source at Argonne National Laboratory.

Source: https://doi.org/10.1038/s41524-026-02388-0

Collected: November 21, 2026 | Automated Research System (AI)

#04 MatSKRAFT: Large-scale materials knowledge extraction from tables

Published November 16, 2026 npj Computational Materials India



OVERVIEW

MatSKRAFT is a computational framework that automatically extracts materials science knowledge from scientific tables using constraint-driven GNNs. It achieves F1 scores of 89.33% for property extraction and 71.35% for composition extraction, processing data 6-496× faster than contemporary frontier LLMs. This system has enabled the construction of a database with 509,281 entries from over 45,500 research publications.

IN DEPTH

Key Findings

MatSKRAFT significantly outperforms contemporary frontier large language models in both accuracy and efficiency. The framework achieves an F1 score of 89.33% for property extraction and 71.35% for composition extraction. In terms of speed, MatSKRAFT processes data 6-496× faster than the fastest and slowest compared models, respectively, while maintaining modest hardware requirements.

Technical Details

The framework addresses the challenge of extracting data from semi-structured tabular formats in scientific literature. MatSKRAFT transforms tables into graph-based representations, which are then processed by constraint-driven Graph Neural Networks (GNNs). This architecture allows the model to encode scientific principles directly into its structure, ensuring that the extracted knowledge adheres to established scientific laws.

Background & Context

Much of the experimental data in materials science remains trapped in tables, resisting systematic extraction and analysis. To demonstrate the scalability of MatSKRAFT, the researchers applied the framework to 66,267 tables sourced from more than 45,500 research publications. This process resulted in a comprehensive database containing 509,281 entries, which includes 104,000 compositions and a wide array of material properties. This dataset expands coverage beyond the scope of major existing databases.

Significance & Outlook

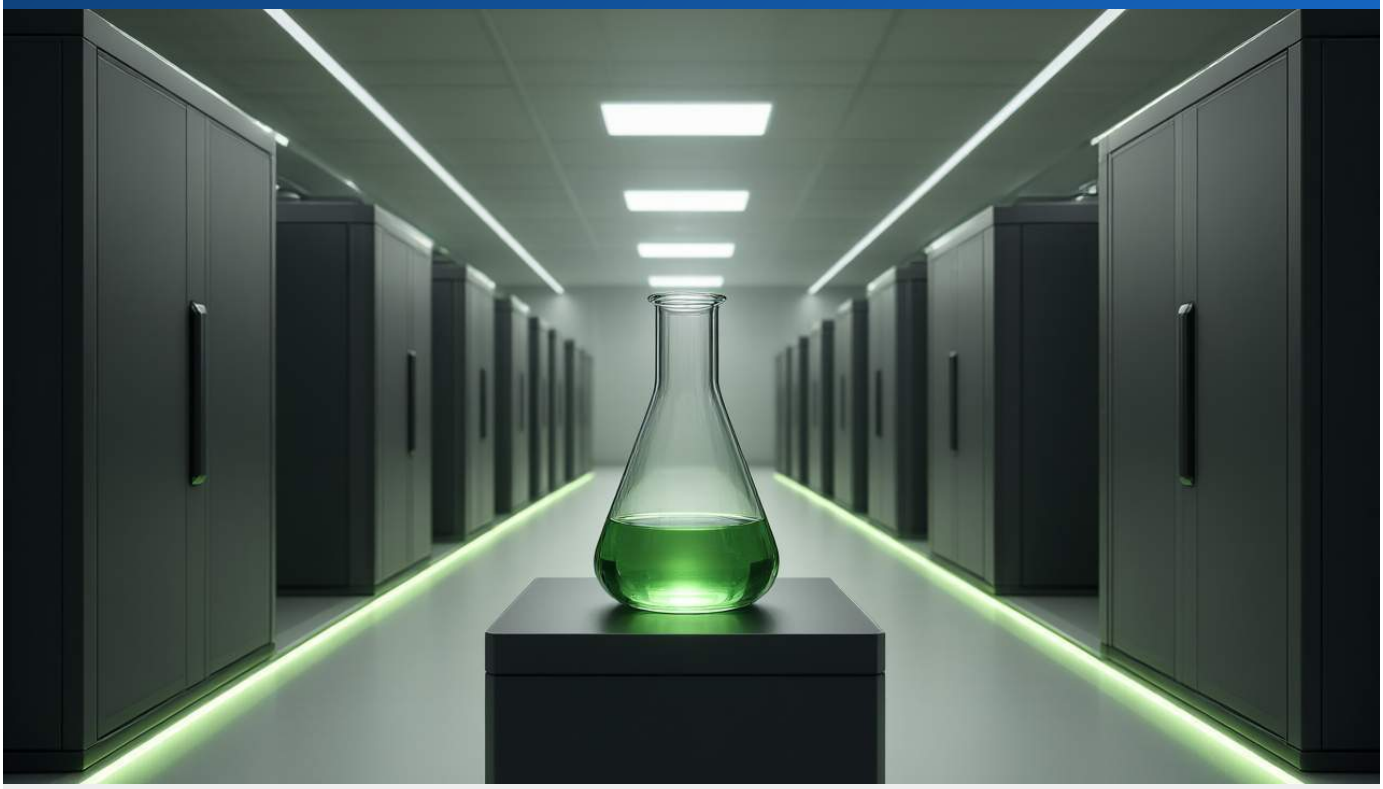
The ability to systematically extract data at this scale reveals previously overlooked materials with distinct property combinations. By integrating vast amounts of literature into a structured format, MatSKRAFT enables the data-driven discovery of composition-property relationships. This capability serves as a cornerstone for future materials and scientific discovery, allowing researchers to synthesize knowledge across extensive literature more effectively than previous manual or LLM-based methods.

Source: https://doi.org/10.1038/s41524-026-02359-5

Collected: November 21, 2026 | Automated Research System (AI)

#05 NVIDIA Nemotron for Science: AI Agents for Scientific Discovery

Published November 18, 2026　NVIDIA 日本版ブログ　USA



OVERVIEW

NVIDIA has introduced Nemotron for Science, combining the Nemotron 3.5 Super VL open model with specialized chemical datasets and post-training recipes to accelerate scientific discovery. The framework enables researchers to build domain-specific agents for tasks such as protein design and chemical synthesis, deployable from desktops to supercomputing centers.

IN DEPTH

Key Findings

NVIDIA Nemotron for Science provides a comprehensive toolkit for developing scientific agents capable of reasoning through complex problems and utilizing professional research tools. A central component is the Nemotron 3.5 Super VL open model, which researchers can adapt via post-training for specialized tasks. Early adopters include life sciences companies Lila Sciences, Terray Therapeutics, and Therna, as well as the Argonne National Laboratory and the Jülich Supercomputing Centre.

Technical Details

The ecosystem includes LitmusBench v0.2, an open dataset containing approximately 50,000 questions focused on chemical reasoning. Based on molecular properties and experimental drug discovery data, this dataset improves model predictions regarding molecular behavior, such as solubility and human body processing. To ensure scalability and security, the system leverages NVIDIA DGX Spark and data center infrastructure, allowing deployment across various scales. The Nemotron 3.5 Super VL model is specifically optimized for local execution on 128 GB NVIDIA DGX Spark systems and is available via Hugging Face.

Implementation and Performance

- **Terray Therapeutics:** Post-trained the model to extract reactants, products, and conditions from scientific diagrams, extracting approximately 76% of annotated reactions in the RxnScribe test.
- **Lila Sciences:** Improved the accuracy of synthesizing quantum dots to 61% (nearly double the previous rate), surpassing GPT-6 Astra's 54% while reducing inference costs by 10 to 20 times.
- **Therna:** Enhanced the prediction accuracy of protein expression ranges by nearly double compared to frontier models.
- **KTH Royal Institute of Technology:** Built an agent that autonomously repaired 81% of failing GROMACS molecular dynamics simulation settings.

Significance and Outlook

By integrating Nemotron into multi-agent systems like StructBioReasoner at Argonne National Laboratory or the RIKYU platform at RIKEN, researchers can accelerate the design of protein-based therapeutics and new materials. The framework allows for the conversion of published chemical discoveries into structured data for synthesis planning, potentially opening new paths in cancer treatment, national security, and mineral research.

Source: https://blogs.nvidia.co.jp/blog/nvidia-introduces-nemotron-for-science-to-accelerate-discoveries-with-agents/

Collected: November 21, 2026 | Automated Research System (AI)

#06 Daviator: NVIDIA's GPU acceleration for plane-wave DFT in VASP

Published November 18, 2026 NVIDIA Technical Blog USA



OVERVIEW

NVIDIA has introduced Daviator, a GPU kernel library that eliminates performance bottlenecks in plane-wave DFT by replacing sequential k-point processing with batched operations. For projected OMat24 workloads, Daviator-enabled VASP can achieve up to 22x faster time-to-solution and 75% lower costs compared to CPU execution. This optimization is critical for generating the massive DFT datasets required to train next-generation machine learning interatomic potentials (MLIP).

IN DEPTH

Key Findings

Daviator significantly improves GPU utilization for small-system calculations in plane-wave DFT, which previously suffered from high kernel launch overhead. Benchmarks indicate that depending on hardware and configuration, the library enables up to 22x faster time-to-solution and a reduction in cost by up to 75% compared to CPU-based execution for projected OMat24 dataset workloads.

Technical Details

The library addresses the bottleneck of sequential k-point processing by implementing batched GPU operations. A central component is the batched Davidson eigensolver, which replaces sequential loops with batched FFTs, GEMMs, and custom kernels to maximize SM occupancy. Additionally, a mixed precision execution layer allows for the use of FP32 arithmetic and storage for specific Davidson solver components while maintaining FP64 for numerically sensitive steps. This mixed precision approach provides roughly 2x end-to-end acceleration on GPUs with high FP32:FP64 ratios, such as the RTX PRO 6000, which features a 64:1 ratio.

Daviator is integrated into the NVIDIA ALCHEMI VASP microservice. Users can enable these features in the INCAR file by setting LDAVIATOR=True and LDAVIATOR_MIXED_PRECISION=1, without needing to modify existing VASP inputs.

Background and Context

Training MLIPs requires millions of DFT calculations for energies, forces, and stresses. While VASP's OpenACC implementation benefits large systems, small unit cells with dense k-point grids often leave GPU performance unused. For example, a single-atom iron system with 752 k-points previously ran 8x slower on GPU hardware than on CPU. Furthermore, while traditional DFT relies on FP64, modern GPU hardware development has shifted toward lower-precision data types for AI workloads, creating a performance gap for FP64-only HPC codes.

Significance and Outlook

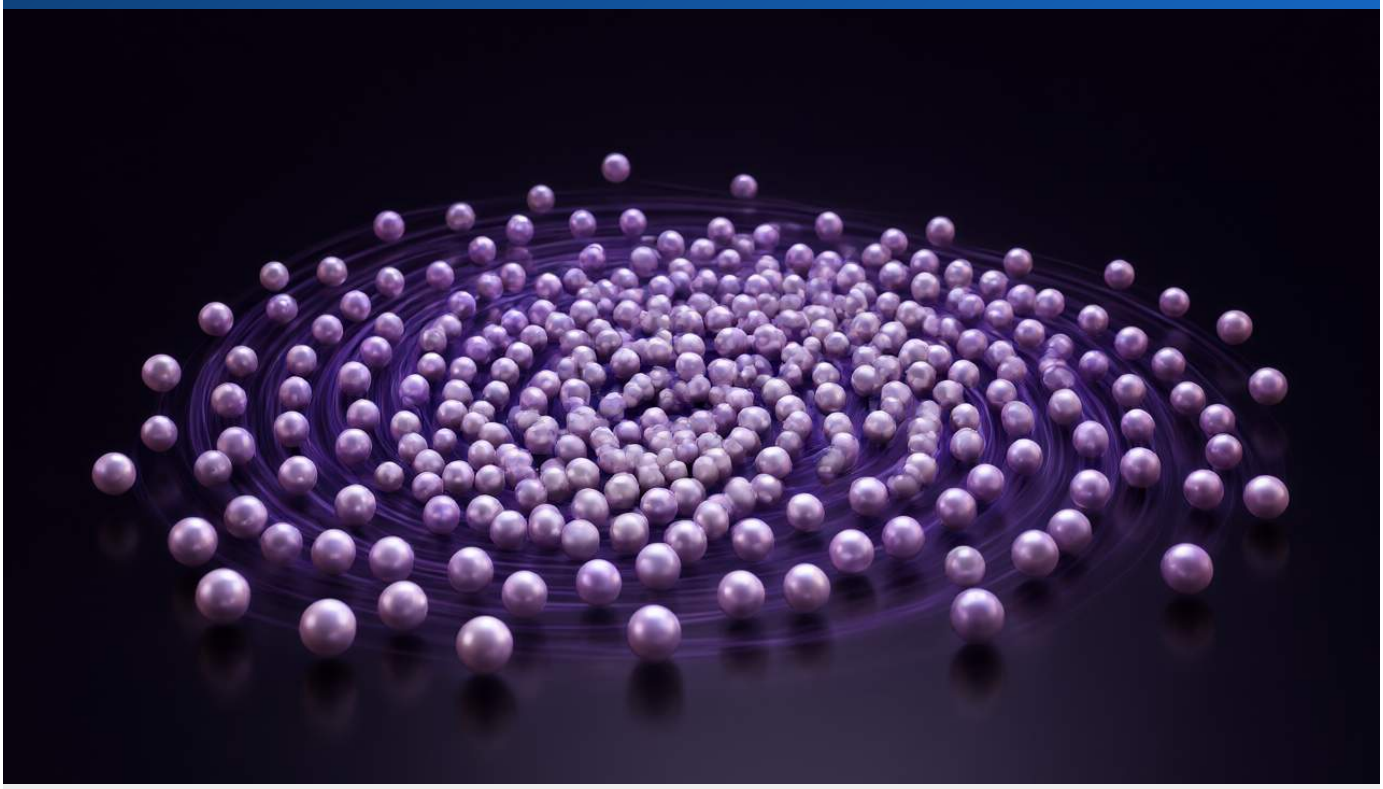
By saturating GPU capacity through k-point and spin-batching, Daviator reduces the time and cost associated with generating ground-truth data for MLIPs. This allows researchers to more efficiently produce the scale of data required for the Open Materials 2024 (OMat24) inorganic materials dataset.

Source: https://developer.nvidia.com/blog/tackling-gpu-bottlenecks-in-plane-wave-dft-with-k-point-batching-and-mixed-precision/

Collected: November 21, 2026 | Automated Research System (AI)

#07 NVIDIA cuPotential: Accelerating Molecular Dynamics Simulations

Published November 16, 2026 NVIDIA Technical Blog USA



OVERVIEW

NVIDIA cuPotential provides modular C API primitives that achieve end-to-end speedups of 2x to over 3x for systems with 24,000 to 400,000 atoms. By optimizing long-range interaction evaluation under periodic boundary conditions, the library significantly reduces energy drift and improves dispersion energy accuracy.

IN DEPTH

Key Findings

NVIDIA cuPotential delivers performance gains of 2x to over 3x across GPU architectures ranging from T4 to RTX PRO 6000 Blackwell for systems containing 24,000 to 400,000 atoms under NPT conditions. Energy drift measurements over 10 ns indicate that cuPotential reduces drift by several hundred fold compared to default PME settings. Additionally, the library improves dispersion energy accuracy by several thousand fold while maintaining force errors comparable to existing methods at a lower computational cost.

Technical Details

The library provides modular C API primitives, including direct space and reciprocal space utilities such as charge spreading functions and cutoff-based interaction evaluators. Developed in collaboration with the Flatiron Institute and D. E. Shaw Research, cuPotential utilizes new C2 continuous splitting and spreading functions. These allow the system to achieve target accuracy with reduced cutoffs and lower grid resolutions than standard u-series and Ewald methods.

For ease of adoption, cuPotential includes a plugin for OpenMM. This plugin serves as a drop-in replacement for the Langevin integrator, Monte Carlo barostat, and nonbonded force, requiring minimal changes to existing user code. The runtime is compatible with Linux x86-64 and requires NVIDIA drivers compatible with CUDA 13.1.

Background and Context

Molecular dynamics (MD) is essential for materials science, biology, chemistry, and physics. While MD codes were early adopters of GPU acceleration due to the parallel nature of atom-pair interactions, the evaluation of long-range interactions—specifically dispersion and electrostatics under periodic boundary conditions—remains a primary engineering and computational bottleneck. cuPotential addresses these challenges by providing a flexible stack of components that developers can integrate into their workloads.

Significance and Outlook

By offering high-accuracy kernels that minimize energy drift and computational overhead, cuPotential allows researchers to accelerate simulations without compromising the integrity of the results. The availability of the library via GitHub provides developers with the necessary benchmarks, examples, and integration code to implement these optimizations in various MD engines.

Source: https://developer.nvidia.com/blog/accelerate-molecular-dynamics-simulations-with-accurate-flexible-primitives-using-nvidia-cupotential/

Collected: November 21, 2026 | Automated Research System (AI)

#08 All-Natural Plastic Substitutes: Robotics and ML Workflow 2026

Published November 16, 2026 Nature Protocols



OVERVIEW

A new integrated workflow combines automated robotics and machine learning to accelerate the discovery of biodegradable, all-natural plastic substitutes. The system reduces the discovery timeline to approximately 2 months by automating the screening and inverse design of nanocomposites with programmable functions.

IN DEPTH

Key Findings

The described workflow enables the accelerated discovery of eco-friendly materials using components from the generally recognized as safe database. By integrating liquid-handling robotics with machine learning, the process automates the inverse design of products that meet user-specified property targets, such as fire resistance or transparency.

Technical Details

The procedure utilizes a multi-stage approach to define and optimize the material design space:

- **Initial Screening:** A pipetting robot prepares a library of 286 nanocomposites. A support vector machine classifier uses the basic physical properties of these products to establish the decision boundary between feasible and infeasible film compositions within a four-component system.
- **Optimization:** Through 14 active learning loops and data augmentation—which involves adding small random changes to characterization data within observed experimental error ranges—the robot fabricates 135 all-natural nanocomposites.
- **Predictive Modeling:** An artificial neural network is constructed to predict physicochemical properties based on composition and identify compositions likely to meet specific user requirements.
- **Validation and Interpretation:** The workflow incorporates Shapley Additive Explanations for model interpretation and molecular dynamics simulations for validation.

Background & Context

The accumulation of nondegradable petrochemical plastic waste has increased the demand for biodegradable substitutes derived from natural resources. However, traditional iterative optimization is inefficient, making it difficult to find all-natural composite materials that match the mechanical, thermal, and optical performance of commercial plastics.

Significance & Outlook

This hybrid approach completes the discovery process in approximately 2 months, significantly reducing the time required for traditional screening experiments. Because the workflow automates the inverse design of materials to satisfy specific property targets, it can be expanded to incorporate new building blocks and applied to other functional-material systems.

Source: https://doi.org/10.1038/s41596-026-01449-4

Collected: November 21, 2026 | Automated Research System (AI)

#09 MatOS: Open-Source Materials Science Agent for R&D Automation

Published November 17, 2026 36Kr（量子位） China



OVERVIEW

MatOS is an open-source scientific research agent system designed to automate materials R&D by unifying data, tools, and workflows into organized assets. It achieves a 97.96% success rate on the MatTools benchmark and can filter 97 solid electrolyte candidates from over 30 million structures.

IN DEPTH

Key Findings

MatOS demonstrates significant performance gains in autonomous materials research. On the ResearchClawBench benchmark, it scores 33.6% higher on average than Codex CLI using the same base model, with an 80.8% increase in materials-specific tasks. In the MatTools benchmark, MatOS achieved a full problem success rate of 97.96%, surpassing both Codex CLI and Claude Code. Practical applications include the development of the DuoCompNet alloy melting point prediction model, which reduced mean absolute error to 3.18°C, a 13.6% improvement over the second-best model.

Technical Details

Developed by the AI Materials Research Institute of the Beijing Zhongguancun Academy and the Zhongguancun Institute of Artificial Intelligence, MatOS utilizes three core architectural designs:

- **Hierarchical Scientific Asset Graph:** Organizes 3,135 datasets, 386 skills, 104 workflows, and 51 expert roles, integrating 48 materials computing software packages. Assets are connected via typed relationships such as dependency and verification.
- **Stateful Adaptive Workflow:** Executes tasks through stateful processes that preserve evidence and allow for on-demand adjustments during long-cycle tasks.
- **Evidence-Driven Asset Evolution:** Converts reusable experience into verified asset revisions, allowing the system to evolve its capabilities during usage.

Background & Context

Materials R&D typically faces challenges such as the need to screen tens of millions of candidate structures or iterative searches across hundreds of alloy formulations. Traditional agents often require full workflow reconstruction when new databases or software are integrated. MatOS addresses these pain points by treating data, knowledge, and tools as "scientific assets," enabling researchers and the agent itself to build new capabilities on top of existing ones without starting from scratch.

Significance & Outlook

By transitioning from simple question-answering to independent research execution, MatOS enables the continuous delivery of high-quality scientific assets. The system's ability to store and reuse practical experiences, such as sampling strategies and simulation durations, provides a scalable framework for autonomous materials discovery. The project's code is available on GitHub, and its technical report is published on arXiv.

Source: http://eu.36kr.com/en/p/4073503237329796

Collected: November 21, 2026 | Automated Research System (AI)

#10 Magnetic Loss Analysis: TUS AI Approach for High-Frequency Materials

Published November 17, 2026 Mirage News（Tokyo University of Science） Japan



OVERVIEW

Researchers at the Tokyo University of Science developed a physics-interpretable AI method to identify microscopic structures causing energy loss in high-frequency magnetic materials. By combining supercomputing with topological machine learning, the team pinpointed specific domain-wall features that contribute most to energy barriers. This framework provides a pathway to designing more efficient power conversion components for electric vehicles and data centers.

IN DEPTH

Key Findings

The study identifies that energy loss in soft magnetic materials is not uniformly distributed across the network but is concentrated around specific microscopic structures. Using inverse mapping, the research team discovered that high-contribution features are located around selected fine wall segments and propagating reversed-domain tips. The analysis further revealed that magnetic domains become increasingly subdivided as excitation frequency increases, with the energy barrier near the coercive region rising in tandem with frequency.

Technical Details

To achieve these results, Professor Masato Kotsugi and his team utilized the Fugaku supercomputer to perform large-scale micromagnetic simulations of an idealized iron-based nanocrystalline soft magnet at frequencies of 100 kHz, 1 MHz, and 10 MHz. The methodology integrated several advanced computational techniques:

- **Persistent Homology:** A mathematical technique used to quantify the shapes and connectivity of evolving magnetic domain-wall networks.
- **Principal Component Analysis (PCA):** Used to compress image-based descriptors into a two-dimensional representation.
- **Inverse Mapping:** Linked the PCA representation to an image-derived exchange-energy measure to construct an energy landscape and identify structures contributing to the energy barrier.

The researchers emphasized that the identified barrier represents an exchange-derived structural energy cost rather than a direct measurement of dissipated heat.

Background & Context

Modern power conversion systems, including those in renewable energy and data centers, rely on transformers and inductors. In soft magnetic materials, magnetization reversal occurs via the movement of magnetic domain walls. While conventional models estimate overall losses, they typically fail to identify which specific microscopic structures create the most resistance during reversal.

Significance & Outlook

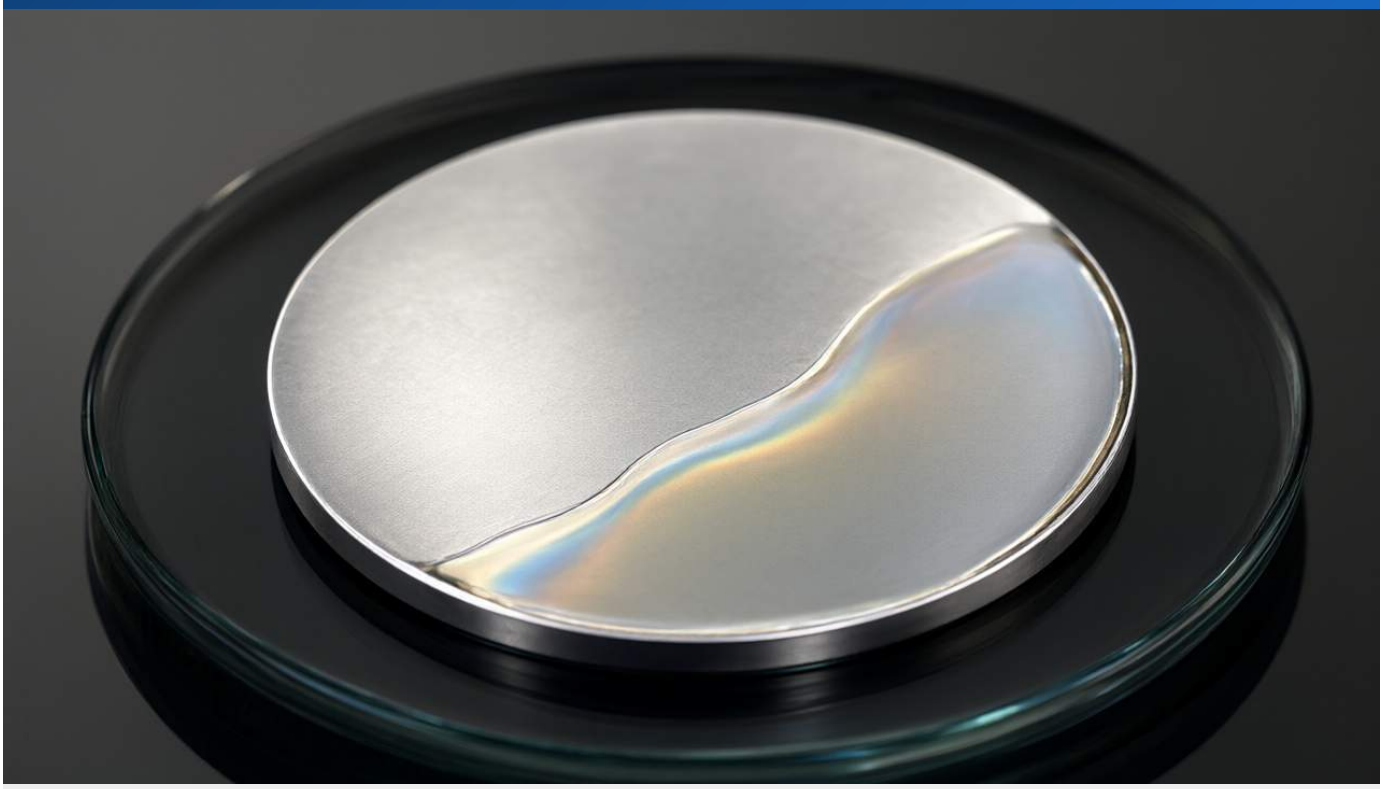
This physics-interpretable AI approach allows researchers to "see the invisible" by explaining results through physical principles rather than simple prediction. By understanding how material composition and processing influence domain-wall behavior, this framework can guide the development of low-loss magnetic materials. Such advancements are expected to improve power conversion efficiency, reduce heat generation, and contribute to CO_2 emission reduction.

Source: https://www.miragenews.com/how-ai-identifies-magnetic-structures-behind-1778948/

Collected: November 21, 2026 | Automated Research System (AI)

#11 MACE-MP-0: ML Model for Lithium Battery SEI Reaction Simulation

Published November 16, 2026 BIOENGINEER.ORG



OVERVIEW

Researchers have fine-tuned the MACE-MP-0 machine learning foundation model to simulate the formation and degradation of the solid-electrolyte interphase (SEI) in lithium metal batteries. The fine-tuned R4 potential reduced force root mean square error to 52 millielectronvolts per angstrom for electrolyte structures, enabling nanosecond-scale simulations of interface chemical dynamics. This approach overcomes the computational bottlenecks of traditional quantum mechanical methods to help design more stable battery interphase layers.

IN DEPTH

Key Findings

The research team successfully adapted the MACE-MP-0 equivariant machine learning interatomic potential to achieve quantum-level predictive precision for lithium metal anode interfaces. The fine-tuned MACE-MP-0-R4 potential demonstrated a significant increase in accuracy over the baseline model, achieving a force root mean square error of 52 millielectronvolts per angstrom on hold-out electrolyte structures, compared to 0.251 electronvolts per angstrom for the out-of-the-box baseline. Across electrolyte, adsorbate, and interface domains, the fine-tuned models maintained force errors between 38 and 67 millielectronvolts per angstrom.

Technical Details

The team employed an iterative fine-tuning workflow consisting of four revisions (R1 through R4) to expand the physical configuration space:

- **R1:** Bulk liquid electrolyte, using MACE-MP-0 with Grimme D3 dispersion corrections and ONETEP density functional theory.
- **R2:** Electrolyte molecules adsorbed onto lithium metal surfaces.
- **R3:** Pristine lithium slab interfaces in contact with bulk electrolyte.
- **R4:** Model protective inorganic coatings between the lithium slab and liquid electrolyte.

The training set included 517 configurations sampled between 150 Kelvin and 350 Kelvin. The model was fine-tuned for 50 epochs using multi-head replay with 131331 structures from the Materials Project repository. Reference calculations used the ONETEP program with a kinetic energy cut-off of 1000 electronvolts and a localization radius of 9.0 Bohr radii.

Background and Context

Lithium metal anodes offer a high theoretical specific capacity of 3860 milliampere-hours per gram and a low electrochemical potential of minus 3.04 volts. However, lithium dendrite growth and electrolyte degradation can lead to internal short circuits and thermal runaway. While ab initio density functional theory provides high accuracy, it is limited to small systems and durations of 1 to 10 picoseconds. Traditional empirical force fields often fail to capture bond formation and rupture during complex interface reactions.

Significance and Outlook

By utilizing the LAMMPS molecular dynamics package, the researchers performed nanosecond-scale simulations on lithium metal slabs with LP57 liquid electrolyte (1 molar lithium hexafluorophosphate in a 3:7 volume ratio of ethylene carbonate and ethyl methyl carbonate) at temperatures from 300 to 450 Kelvin. This capability allows for the design of stable, electronically insulating, and ionically conducting interphase layers to suppress parasitic chemical decomposition.

Source: https://bioengineer.org/machine-learning-model-simulates-solid-electrolyte-interphase-reactions-in-lithium-batteries/

Collected: November 21, 2026 | Automated Research System (AI)

#12 NVIDIA DAQIRI: Real-Time GPU Data Acquisition for Scientific Instruments

Published November 19, 2026 NVIDIA Technical Blog USA



OVERVIEW

NVIDIA DAQIRI enables the direct transfer of high-bandwidth sensor data into GPU memory, bypassing disk writes to allow real-time AI processing. A case study with Lawrence Berkeley National Laboratory demonstrates the ability to process over 500 Gbps of electron microscopy data for immediate experimental adjustment.

IN DEPTH

Key Findings

NVIDIA DAQIRI is a software-defined data acquisition library within the NVIDIA Holoscan platform that bridges high-bandwidth sensors with GPU-accelerated workflows. The library allows developers to perform real-time frame assembly, correction, and AI inference during acquisition. In a practical application at Lawrence Berkeley National Laboratory's Molecular Foundry, the system handles data streams exceeding 500 Gbps—over 500x faster than conventional electron energy-loss spectroscopy (EELS) streams—enabling researchers to adjust acquisition parameters and regions of interest in real time.

Technical Details

The latest DAQIRI release introduces several hardware-level optimizations to reduce latency and CPU overhead:

- **Hardware Packet Reordering:** Supported NVIDIA ConnectX-7 or newer NICs can now offload packet reordering. The NIC parses sequence fields and places payloads directly into ordered memory positions, eliminating the need for a software reorder copy step.
- **Throughput Improvements:** An updated ibverbs engine nearly doubles per-core throughput when handling small packet sizes.
- **Low-Latency Polling:** A direct polling mode has demonstrated application send-to-receive latency of 2-4 µs.

Background & Context

Traditional scientific instruments are often limited by the speed at which data can be moved from readout electronics to processing units. In EELS, capturing both strong and faint signals typically requires two separate acquisition processes and subsequent alignment. The high-dynamic-range camera developed at the Molecular Foundry simplifies this by capturing both signals simultaneously. To manage the resulting 500 Gbps stream, FPGA readout electronics transmit data over Ethernet, which DAQIRI then moves directly into GPU memory for processing.

Significance & Outlook

By enabling submicron spatial resolution from 5-micrometer sensor pixels and providing immediate feedback, this GPU-native approach allows for more flexible instrument design. This architecture is applicable to any Ethernet-capable instrument, allowing makers to integrate real-time AI to guide experiments and optimize operating conditions without the bottleneck of writing data to disk first.

Source: https://developer.nvidia.com/blog/enabling-real-time-ai-for-scientific-instruments-with-gpu-native-data-acquisition/

Collected: November 21, 2026 | Automated Research System (AI)