

AI News Digest

11/17/2025 - 11/23/2025

WEEKLY DIGEST: OVERALL AI NEWS

Weekly AI Developments: Innovations in Design, Data Analysis, and AI Infrastructure

This week in AI news, significant advancements were made in various fields, including CAD software, natural language visualization, and AI model enhancements. MIT engineers introduced an AI agent capable of creating 3D objects from sketches, while researchers explored the implications of modified AI models for sensitive topics. Additionally, the emergence of natural language visualization is set to revolutionize data analysis, making it more accessible for users. The ongoing evolution of AI infrastructure highlights the importance of trust and data quality in scaling AI applications effectively.

KEY FINDINGS

- ✓ MIT's new AI model can operate CAD software, simplifying the design process.
- ✓ Modified DeepSeek R1 shows promise in providing uncensored responses to sensitive inquiries.
- ✓ Natural Language Visualization is changing how users interact with data, allowing for more intuitive queries.

IMPLICATIONS

- The development of AI agents for CAD could democratize 3D design, making it accessible to non-experts.
- Enhanced AI models may lead to more open discussions on sensitive topics, impacting information dissemination.
- Natural Language Visualization could significantly reduce the learning curve for data analysis tools, increasing their adoption.

Key Metrics: 41,000 videos in the VideoCAD dataset; \$1.5 trillion projected global investment in AI by 2025; Three-part approach for scaling AI: trust, data quality, IT leadership

1. New AI agent learns to use CAD to create 3D objects from sketches

Category: Applications

Departments: Computer Science, Industrial, Manufacturing & Systems Engineering, Aerospace & Mechanical Engineering

Topics: Computer Vision, CAD Co-Pilot, VideoCAD Dataset, 3D Object Generation

Impact Score: 26.0

Rubric Scores

Research Relevance: 5

Educational Value: 4

Innovation / Novelty: 5

Practical Impact: 5

Interdisciplinary Potential: 4

Ethical / Policy Implications: 3

MIT engineers have developed an AI model designed to operate Computer-Aided Design (CAD) software in a manner similar to human users, addressing the software's steep learning curve. The team created a dataset called VideoCAD, comprising over 41,000 videos that demonstrate the step-by-step construction of 3D models from 2D sketches, capturing the specific user-interface interactions involved. This AI system aims to function as a "CAD co-pilot," assisting users by suggesting next steps and automating tedious tasks, thereby enhancing productivity and accessibility for those without extensive CAD training. The research will be presented at the upcoming Conference on Neural Information Processing Systems (NeurIPS).

<https://news.mit.edu/2025/new-ai-agent-learns-use-cad-create-3d-objects-sketches-1119>

2. Quantum physicists have shrunk and "de-censored" DeepSeek R1

Category: Research

Departments: Computer Science, Political Science & Public Administration, Electrical & Computer Engineering

Topics: Generative AI, Censorship Handling, Model Compression Techniques, Quantization and Pruning

Impact Score: 25.0

Rubric Scores

Research Relevance: 5

Educational Value: 4

Innovation / Novelty: 4

Practical Impact: 5

Interdisciplinary Potential: 3

Ethical / Policy Implications: 4

Researchers at Multiverse tested a modified AI model against the original DeepSeek R1 to evaluate its ability to provide uncensored responses to sensitive questions, such as those related to Chinese political topics. The modified model demonstrated the capacity to deliver factual answers comparable to Western models, indicating a potential advancement in handling censorship. This work is part of Multiverse's initiative to create more efficient AI models that require less computational power while maintaining performance. The study highlights ongoing challenges in model compression, where techniques like quantization and pruning often result in trade-offs between size and capability.

<https://www.technologyreview.com/2025/11/19/1128119/quantum-physicists-compress-and-deconsor-deepseekr1/>

3. Natural Language Visualization and the Future of Data Analysis and Presentation

Category: Applications

Departments: Computer Science, Engineering Education & Leadership

Topics: Natural Language Processing, Natural Language Visualization, User-Friendly Data Interaction, AI-Driven Query Formulation

Impact Score: 25.0

Rubric Scores

Research Relevance: 4

Educational Value: 4

Innovation / Novelty: 5

Practical Impact: 4

Interdisciplinary Potential: 5

Ethical / Policy Implications: 3

The article discusses the emerging paradigm of Natural Language Visualization (NLV) in data analysis, which aims to simplify the interaction between users and data by allowing users to pose questions in natural language rather than navigating complex software interfaces. This approach is likened to the artistic process of Fujiko Nakaya, where the user conceptualizes the inquiry while the system handles the technical execution, including query formulation and data visualization. The article acknowledges the current limitations of AI tools in delivering accurate and reliable outputs, highlighting the need for further advancements to realize the full potential of NLV. Ultimately, it presents NLV as a transformative step towards making data analysis accessible to a broader audience beyond just data experts.

<https://towardsdatascience.com/natural-language-visualization-and-the-future-of-data-analysis-and-presentation/>

4. The cost of thinking

Category: Research

Departments: Computer Science, Electrical & Computer Engineering

Topics: Large Language Models, Reasoning Models, Reinforcement Learning, Complex Problem Solving

Impact Score: 24.0

Rubric Scores

Research Relevance: 5

Educational Value: 4

Innovation / Novelty: 5

Practical Impact: 4

Interdisciplinary Potential: 3

Ethical / Policy Implications: 3

Researchers at MIT's McGovern Institute have developed a new generation of large language models (LLMs) known as reasoning models, which demonstrate improved capabilities in solving complex problems, including math and coding tasks. The study, published in PNAS, reveals that the "cost of thinking" for these models parallels that of human cognition, as both require time to process challenging problems. The reasoning models utilize a step-by-step approach and reinforcement learning during training, which enhances their ability to arrive at correct solutions. This advancement suggests that reasoning models may exhibit a human-like approach to problem-solving, despite not being designed with this intention.

<https://news.mit.edu/2025/cost-of-thinking-1119>

5. Designing digital resilience in the agentic AI era

Category: Business

Departments: Computer Science, Electrical & Computer Engineering, Political Science & Public Administration

Topics: AI Policy & Regulation, Digital Resilience, Data Fabric Integration, Agentic AI Adaptation, Machine Data Utilization

Impact Score: 23.0

Rubric Scores

Research Relevance: 4

Educational Value: 4

Innovation / Novelty: 3

Practical Impact: 5

Interdisciplinary Potential: 4

Ethical / Policy Implications: 3

Global investment in AI is expected to reach \$1.5 trillion by 2025, yet less than half of business leaders express confidence in their organizations' ability to maintain service continuity and security during unexpected events. The article emphasizes the necessity of a data fabric—an integrated architecture that connects and governs information across all business layers—to enhance digital resilience and empower both human teams and agentic AI systems. It highlights the critical role of machine data, which includes logs and telemetry from devices and systems, in enabling agentic AI to operate effectively and adapt to risks. The current lack of comprehensive machine data integration poses significant challenges, potentially leading to errors and limiting the capabilities of agentic AI.

<https://www.technologyreview.com/2025/11/20/1127941/designing-digital-resilience-in-the-agentic-ai-era/>

6. Realizing value with AI inference at scale and in production

Category: Business

Departments: Computer Science, Electrical & Computer Engineering

Topics: Enterprise AI, Trusted Inference, Data-Centric AI, AI Factory, Data Architecture

Impact Score: 23.0

Rubric Scores

Research Relevance: 4

Educational Value: 4

Innovation / Novelty: 3

Practical Impact: 5

Interdisciplinary Potential: 3

Ethical / Policy Implications: 4

The article discusses a three-part approach essential for scaling AI effectively, emphasizing the importance of trust, data quality, and IT leadership. It highlights that trusted inference is crucial for high-stakes applications, as unreliable data can lead to decreased trust and productivity. The shift from model-centric to data-centric thinking is noted, with organizations focusing on breaking down data silos and enhancing data architecture to unlock value. This evolution is characterized by the emergence of the "AI factory," which facilitates continuous intelligence through efficient data management.

<https://www.technologyreview.com/2025/11/18/1128007/realizing-value-with-ai-inference-at-scale-and-in-production/>

7. Scaling innovation in manufacturing with AI

Category: Business
Departments: Industrial, Manufacturing & Systems Engineering, Computer Science
Topics: AI in Manufacturing, Digital Twins, Real-time Visualization, Data Integration Techniques

Impact Score: 22.0

Rubric Scores

Research Relevance: 4	Educational Value: 4	Innovation / Novelty: 4
Practical Impact: 5	Interdisciplinary Potential: 3	Ethical / Policy Implications: 2

AI-powered digital twins are transforming manufacturing by enabling comprehensive real-time visualization of entire production lines, rather than just individual machines. This integration of various data types-shop-floor telemetry, enterprise data, and immersive modeling-allows manufacturers to enhance efficiency and minimize downtime, which can reach up to 40% in high-speed industries. Current estimates indicate that approximately 50% of manufacturers are deploying AI in production, a significant increase from 35% reported in 2024. Larger manufacturers, particularly those with revenues exceeding \$10 billion, are leading this trend, with 77% already implementing AI use cases.

<https://www.technologyreview.com/2025/11/19/1128067/scaling-innovation-in-manufacturing-with-ai/>

8. Top 7 Open Source AI Coding Models You Are Missing Out On

Category: Applications
Departments: Computer Science, Electrical & Computer Engineering
Topics: Consumer AI, Open Source Coding Models, Long-Horizon Reasoning, Interactive Agent Workflows

Impact Score: 22.0

Rubric Scores

Research Relevance: 4	Educational Value: 4	Innovation / Novelty: 4
Practical Impact: 5	Interdisciplinary Potential: 3	Ethical / Policy Implications: 2

The article discusses the advantages of using open-source coding models that can be run locally, emphasizing privacy, control, and cost savings compared to cloud-based AI coding assistants. It highlights seven top-tier models, including Kimi-K2-Thinking and MiniMax-M2, which excel in coding benchmarks and offer significant improvements in long-horizon reasoning and efficiency. Kimi-K2-Thinking features a complex architecture with 1 trillion parameters and achieves high scores in various coding tasks, while MiniMax-M2 is designed for efficiency with a focus on interactive agent workflows. The article aims to inform users about viable alternatives to proprietary tools that do not require sending code to external servers.

<https://www.kdnuggets.com/top-7-open-source-ai-coding-models-you-are-missing-out-on>

9. Networking for AI: Building the foundation for real-time intelligence

Category: Applications

Departments: Electrical & Computer Engineering, Computer Science

Topics: AI Hardware, Real-Time Data Processing, AI-Ready Networking, Ultra-Low Latency Infrastructure

Impact Score: 21.0

Rubric Scores

Research Relevance: 3

Educational Value: 4

Innovation / Novelty: 4

Practical Impact: 5

Interdisciplinary Potential: 3

Ethical / Policy Implications: 2

The Ryder Cup partnered with HPE to develop a centralized operational hub that utilized a high-performance network and private-cloud environment to enhance decision-making through data visualization. This initiative demonstrated the necessity of AI-ready networking, highlighting that effective AI implementation relies on the ability to manage and process large volumes of real-time data efficiently. A survey indicated that while 45% of organizations can now perform real-time data operations, many still face challenges in integrating data collection with decision-making. The event showcased the importance of specialized network infrastructure capable of supporting AI workloads, emphasizing characteristics such as ultra-low latency and lossless throughput to ensure optimal performance.

<https://www.technologyreview.com/2025/11/18/1127997/networking-for-ai-building-the-foundation-for-real-time-intelligence/>

10. Google's new Gemini 3 "vibe-codes" responses and comes with its own agent

Category: Applications

Departments: Computer Science, Electrical & Computer Engineering

Topics: Generative AI, Multi-Step Task Management, AI Integration with Services, AI-Generated Summaries

Impact Score: 21.0

Rubric Scores

Research Relevance: 3

Educational Value: 4

Innovation / Novelty: 4

Practical Impact: 5

Interdisciplinary Potential: 3

Ethical / Policy Implications: 2

Google has announced the launch of Gemini 3, which includes the introduction of Gemini Agent, an experimental feature that facilitates multi-step task management within the app. This agent can integrate with services like Google Calendar and Gmail to perform tasks such as inbox organization and schedule management, breaking tasks into steps and requiring user approval at each stage. Additionally, Gemini 3 Pro offers subscribers enhanced AI-generated summaries based on reasoning, marking a deeper integration of Gemini with Google's existing products. The Gemini Agent feature will be available to Google AI Ultra subscribers in the US starting November 18.

<https://www.technologyreview.com/2025/11/18/1128065/googles-gemini-3/>