

# The Autonomy Frontier: Engineering System 2 Reasoning to Solve the 99.9% Wall

The transition of autonomous vehicle (AV) technology from experimental prototypes to global urban infrastructure has reached a critical inflection point as of March 12, 2026. While the industry successfully navigated the initial phases of Level 4 autonomy through intensive imitation learning and high-definition mapping, it has encountered a statistical plateau known as the "99.9% Wall." This wall represents the "Long Tail" of edge cases—unpredictable, low-frequency, and high-complexity scenarios that traditional "System 1" reactive architectures cannot resolve. The current state of the art represents an architectural migration toward "System 2" deliberative reasoning, where world models and foundation robotics enable vehicles to move beyond pattern matching toward a fundamental understanding of physical causality and social intent. This report investigates the technical, geopolitical, and ethical dimensions of this frontier, serving as a comprehensive strategic analysis for the next era of autonomous deployment.

## Defining the 2026 Edge Case: The Anatomy of the Long Tail

The edge case crisis of 2026 is no longer defined by sensor noise or basic classification errors, which were largely solved by the mid-2020s through multimodal sensor fusion. Instead, the current crisis involves "Active Edge Cases" that emerge from the interaction between autonomous agents and complex human social systems. In major hubs such as San Francisco, Phoenix, Beijing, and Shenzhen, the density of these interactions has exposed the limitations of rule-based heuristics.

### Categorization of Active Edge Cases in Global Hubs

As autonomous fleets scale, the specific nature of the edge cases they encounter has become localized based on urban design and social behavior. The following table categorizes the primary active edge cases across dominant global testing grounds as of early 2026.

Urban Hub	Dominant Edge Case Category	Primary 2026 Obstacle	Technical Bottleneck

San Francisco	Social Friction/Hacking	Intentional pedestrian blocking and "safety trapping" of AVs.	Non-cooperative game theory/Intent prediction.
Beijing	Unstructured Density	Swarms of non-standard 2026 delivery droids and micro-mobility devices.	High-resolution voxel occupancy in crowded scenes.
Shenzhen	Multi-Agent Interference	High-frequency signal interference in 5G-Advanced urban canyons.	V2X reliability and fall-through logic.
Phoenix	Environmental Anomalies	Extreme heat-induced sensor degradation and high-velocity dust occlusions.	Dynamic sensor calibration and physics-aware world models.

In San Francisco, the phenomenon of "social hacking" has emerged as a significant operational hurdle. Groups of pedestrians have identified the safety-critical constraints of AV software—which prioritizes collision avoidance at all costs—and intentionally exploit these triggers to block vehicle movement for social media content or protests.<sup>1</sup> Because traditional System 1 architectures lack social reasoning, they cannot distinguish between a legitimate safety hazard and an intentional social obstruction, leading to "permanent gridlock" states that require human intervention.

### Multi-Agent Friction and 2026 Micro-Mobility

The urban landscape of 2026 has been further complicated by the explosion of new micro-mobility devices, including wheel-leg robots that climb stairs and high-speed electric unicycles capable of 40 mph speeds.<sup>2</sup> These devices do not conform to the kinematic models used for traditional bicycles or pedestrians. When an AV encounters a 2026-era wheel-leg droid on a sidewalk, it must reason about the droid's ability to suddenly enter the roadway in a

non-linear fashion. This "Multi-Agent Friction" is a primary component of the Long Tail, where the sheer variance in agent types exceeds the diversity of any static training dataset.<sup>3</sup>

## The Physics of Visual Occlusion and Temporal Reasoning

Visual occlusion remains a fundamental challenge in solving the Long Tail. Traditional systems often experienced "micro-hesitations"—abrupt braking when a detected object was momentarily obscured by a static object like a mailbox or a bus stop.<sup>5</sup> System 2 reasoning addresses this through "Temporal Transformers" and "Object Permanence." By maintaining a continuous temporal buffer, the vehicle's architecture creates a persistent memory of the environment. If a car drives behind a large truck, the system "knows" the car exists and predicts its likely exit trajectory based on a learned world model of physical momentum, rather than re-detecting it as a "new" object once it reappears.<sup>5</sup>

## The Technical Solution Stack: Beyond Reactive Sensors

To breach the 99.9% wall, the industry is transitioning from "mimicking human drivers" to "understanding the physical world." This transition is powered by a new stack of AI architectures: World Models, Joint-Embedding Predictive Architectures (JEPA), and Vision-Language-Action (VLA) foundation models.

### World Models and the JEPA Paradigm

The most significant shift in early 2026 is the emergence of General World Models as the core of the autonomous brain. Led by Yann LeCun at AMI Labs (Advanced Machine Intelligence), this paradigm argues that Large Language Models (LLMs) are insufficient for physical intelligence because they lack a mental model of physics.<sup>6</sup> AMI Labs, which raised €1.03 billion in March 2026, focuses on Joint-Embedding Predictive Architectures (JEPA) to enable cars to "dream" in latent space.<sup>8</sup>

JEPA differs fundamentally from previous generative AI. Instead of predicting the next pixel in a video—which is computationally expensive and filled with irrelevant noise—JEPA predicts the next *representation* of the world state.<sup>7</sup> This allows the AI to simulate millions of "what-if" scenarios, such as near-miss collisions or unusual weather conditions, entirely in latent space before they occur on the road.<sup>7</sup>

### Tesla FSD v13 and v14: Voxelization and Architectural Efficiency

Tesla's transition to FSD v13 and v14 represents the production-scale application of these world model concepts. Released in early 2026, v13 achieved a 95% reduction in disengagements by moving away from simple imitation toward a "physics-aware" world model.<sup>5</sup> A key technical

pillar is Occupational Networks 3.0, which uses high-resolution voxelization.

Feature	Tesla FSD v12 (2024-25)	Tesla FSD v13/v14 (2026)	Technical Impact
Core Logic	Imitation Learning	World Model (System 2)	Move from mimicry to physical reasoning.
Temporal Context	Single frames/Short history	Long-horizon temporal buffers	Solves object permanence and occlusion. <sup>5</sup>
Resolution	Standard voxel grids	8x High-Res Voxelization	Navigation of "unstructured" objects (debris). <sup>5</sup>
Hardware Targets	Unified HW3/AI4 software	Bifurcated (v14 vs v14 Lite)	Optimization for compute-constrained chips. <sup>11</sup>

FSD v14 further refined this by introducing "V14 Lite" for legacy Hardware 3 (HW3) vehicles, while the full v14 runs on AI4 (Hardware 4) with 5-megapixel cameras and five times the perception bandwidth.<sup>5</sup> The architectural divergence highlights the "Compute Barrier": System 2 reasoning requires significantly more FLOPs to process high-resolution voxels and maintain 10-second predictive windows.<sup>5</sup>

## Foundation Models for Action: The Strategic Brain

While the "Tactical Controller" handles steering and braking, foundation models like Gemini Robotics 1.5 and GPT-5.4-Robotics serve as the "Strategic Brain." These are Vision-Language-Action (VLA) models that integrate reasoning with motor control.<sup>13</sup> Google's Gemini Robotics-ER (Embodied Reasoning) 1.5, for example, specializes in understanding physical spaces and making logical decisions, such as finding a safe trajectory to bypass a blocked road.<sup>13</sup>

These models allow a robot or AV to "think before acting." When a vehicle encounters a complex construction site with non-standard signage, the foundation model can reason about the visual scene—identifying which cones are relevant and which can be ignored—and then pass high-level insights to the tactical controller.<sup>13</sup> This enables a degree of "Generality" where the system can handle novel objects it has never seen in its training data.<sup>15</sup>

## V2X and 6G: The Collective Intelligence Layer

In March 2026, the 6G roadmap has become a critical enabler for collective intelligence in AV fleets. Announced at MWC Barcelona, 6G is designed as an "AI-native" system that distributes compute across the network edge.<sup>18</sup> This allows for "Vehicle-to-Everything" (V2X) communication that goes beyond simple position sharing.

Under the 6G paradigm, if one AV encounters a unique edge case (e.g., a burst water main), it doesn't just brake; it uploads a "Reasoning Trace" and a latent world model update to the local edge node. This node then pushes the update to all other vehicles in the vicinity, effectively "vaccinating" the entire fleet against that specific edge case in real-time.<sup>20</sup> This represents a shift from individual intelligence to a distributed digital nervous system.

## The Global Expansion and Geopolitical Friction: Autonomy Silos

The expansion of self-driving vehicles in 2026 is no longer a purely technological race; it is a battle for data sovereignty and regional AI independence. The world has effectively split into two "Autonomy Silos," led by the United States and China, each with distinct philosophical and regulatory frameworks.

### The Western Approach vs. The Eastern Approach

The divergence between Western and Eastern autonomous strategies has widened by March 2026, driven by different risk tolerances and urban infrastructures.

Aspect	Western Approach (Waymo, Zoox)	Eastern Approach (Xiaomi, Baidu)
Philosophy	"Safety First" through deep geo-fencing.	"Scale First" through aggressive urban deployment.
Sensor Stack	Lidar-heavy, high redundancy.	Vision-centric, heavy AI post-processing. <sup>22</sup>
Mapping	High-Definition (HD) Map dependent.	"Mapless" or standard navigation-based.
Integration	Standalone	Integrated ecosystem

	mobility-as-a-service.	(Phone-Car-Home). <sup>23</sup>
--	------------------------	---------------------------------

Xiaomi's entry into the space with the SU7 and its testing of humanoid robots on the production line exemplifies the "Eastern Approach": a vertically integrated AI strategy where the lessons learned in factory robotics are directly transferred to the vehicle's driving brain.<sup>23</sup>

## Data Sovereignty and the US-China Data War

The most significant geopolitical hurdle in 2026 is the "Data Sovereignty" crisis. Both the U.S. and China have moved to restrict the flow of autonomous data across borders, citing national security concerns. In January 2025, the U.S. Bureau of Industry and Security (BIS) finalized rules effectively banning Chinese hardware and software in connected vehicles.<sup>24</sup>

The BIS rule identifies two critical systems for regulation: Vehicle Connectivity Systems (VCS) and Automated Driving Systems (ADS).<sup>24</sup> The rule prohibits connected vehicle manufacturers from importing vehicles that incorporate covered software developed by entities subject to the jurisdiction of China or Russia.<sup>24</sup> Conversely, China has tightened its data localization rules, requiring all data collected by foreign EVs (like Tesla) to be stored and processed within Chinese borders.<sup>26</sup> This creates two separate intelligence environments, where models trained on San Francisco streets are legally and technically "blind" to the nuances of Beijing traffic.

## The Liability and Ethics Layer: Auditing Machine Intent

As AVs move into the "Delegated Intelligence" era, the focus has shifted from "Did the car crash?" to "Why did the car decide to do that?" This is the era of Agentic Liability and the Trolley Problem 2.0.

### The Reasoning Trace and Agentic Liability

In 2026, the industry has introduced the "Reasoning Trace" as a standard for post-accident auditing. When a vehicle encounters a "No-Win" scenario—for example, a child running into the street where any avoidance maneuver results in a collision—the AI must generate an auditable causal chain.<sup>4</sup>

This trace, powered by models like Gemini Robotics, converts the high-dimensional latent decisions of the neural network into natural language explanations: *"Detected pedestrian on collision path; braking force maximized; identified leftmost lane as path with lowest predicted impact energy"*.<sup>14</sup> This transparency is crucial for "Agentic Liability," where the enterprise assumes responsibility for the autonomous decisions of its digital agents.<sup>29</sup> The OSWorld-V benchmark has become a key metric here, testing how well agents can reason through multi-step workflows while maintaining safety boundaries.<sup>29</sup>

## The Insurance Metamorphosis

Insurance models are undergoing a fundamental shift from "Driver Liability" to "Product/Algorithm Liability." In early 2026, major insurers like Lemonade began offering substantial discounts (up to 50%) for owners using supervised FSD, as the data increasingly shows that System 2 world models are statistically safer than human drivers.<sup>31</sup> However, this forces a consolidation of the insurance industry, where the "Product Liability" for millions of cars shifts onto the balance sheets of the AV manufacturers themselves.

## Industrial Humanoids vs. AVs: The Benchmark of Physical Intelligence

A compelling metric for the state of autonomy in 2026 is the comparison between Xiaomi's humanoid robots in factories and Level 4 AVs in the wild.

### Xiaomi's 90.2% Success Rate

Xiaomi recently reported that its humanoid robots achieved a 90.2% success rate in autonomous fastening tasks on an EV assembly line.<sup>23</sup> These robots use a Vision-Language-Action (VLA) model that fuses visual data with fingertip tactile sensors.<sup>22</sup> While 90.2% is impressive for a factory task with a 76-second cycle time, it highlights the gap between "Industrial Autonomy" and "Road Autonomy".<sup>23</sup>

Metric	Xiaomi Industrial Humanoid	Level 4 AV in Rain/Snow
Success Rate	90.2% (Task completion)	~75-80% (Reliable operation)
Cycle Time	76 seconds (Fixed)	Milliseconds (Real-time)
Environmental Control	Semi-structured	Unstructured/Chaos
Perception Modality	Multi-modal (Tactile + Vision)	Multi-modal (Lidar + Vision)

The success of humanoids in the factory is driven by "Embodied Reasoning"—the ability to use tactile sensors to overcome visual occlusions during delicate assembly.<sup>23</sup> AVs are now attempting to replicate this through "Environmental Tactile" sensing, where world models

interpret the "feel" of road friction and tire slip as a surrogate for tactile data, allowing for better performance in rain and snow where vision is degraded.

## Technical Requirements and Infrastructure: The Gigawatt Scale

The transition to System 2 reasoning is constrained by one final factor: raw compute power. In 2026, we have entered the age of "Gigawatt-Scale" training.

### Compute Factories and the 2026 Benchmarks

Elon Musk's xAI brought the first 1-gigawatt training cluster, "Colossus 2," online in early 2026.<sup>32</sup> By April 2026, this is expected to reach 1.5 gigawatts of power, housing over one million H100 GPU equivalents.<sup>32</sup> This scale is necessary because System 2 world models require "Post-Training Scaling" and "Test-Time Scaling"—where the model "thinks" longer to generate better outcomes—consuming up to 100 times more compute than standard inference.<sup>12</sup>

### OSWorld-V and Agentic Logic

The OSWorld-V benchmark, where GPT-5.4 achieved a 75% success rate in autonomous navigation, has become the "Gold Standard" for agentic logic.<sup>2</sup> This benchmark proves that AI can now orchestrate complex, multi-step tasks across heterogeneous applications with minimal supervision.<sup>29</sup> For AVs, this logic is being adapted into "Road-Use Agents" that can navigate software-defined vehicle layers, manage V2X communications, and execute complex detours without human prompts.<sup>34</sup>

## Synthesis: The First Principles of Physical Navigation

The resolution of the Edge Case Crisis lies in the synthesis of world knowledge and physical action. By March 2026, the industry has realized that solving the final 0.1% requires a vehicle to act not just as a sensor platform, but as a "Thinking Agent."

The "First Principles" of 2026 autonomy dictate that:

1. **Physics Trumps Pattern Matching:** High-resolution voxelization and latent world models (JEPA) allow vehicles to navigate "unstructured" debris and novel objects by understanding their physical occupancy rather than their label.<sup>5</sup>
2. **Temporal Context is Essential:** The solution to visual occlusion is not more cameras, but better memory. Temporal transformers that maintain object permanence across a 10-second buffer are the only way to solve the child-behind-the-car problem.<sup>5</sup>
3. **Geopolitical Alignment is the Final Frontier:** Technical success is meaningless without regulatory and data alignment. The creation of "Autonomy Silos" threatens to slow the path to AGI in robotics by fragmenting the global dataset.<sup>24</sup>

As we look toward the remainder of 2026, the winners of the autonomy race will not be those who drive the most miles, but those who build the most robust world models—systems that can "dream" of the edge cases of tomorrow and reason their way through them today. The transition from System 1 to System 2 is not just a software update; it is a fundamental re-engineering of the machine's relationship with the physical world.

## Works cited

1. AI - 2026's New Operating System: Autonomous AI - Tech Channels, accessed on March 12, 2026, <https://www.tech-channels.com/breaking-news/ai-2026s-new-operating-system-autonomous-ai>
2. Tech for Tomorrow's World - Spreaker, accessed on March 12, 2026, <https://www.spreaker.com/podcast/tech-for-tomorrow-s-world--6551255>
3. AI Agents and Agentic AI: The Future of Autonomous Intelligence in 2026 - Fast VPS Hosting, accessed on March 12, 2026, <https://www.logicweb.com/ai-agents-and-agentic-ai-the-future-of-autonomous-intelligence-in-2026/>
4. AI Automation in 2026: The Rise of Autonomous Systems at Scale - AI World Journal, accessed on March 12, 2026, <https://aiworldjournal.com/ai-automation-in-2026-the-rise-of-autonomous-systems-at-scale/>
5. Tesla FSD v13: Navigating the Architectural Divide Between HW3 ..., accessed on March 12, 2026, <https://www.teslaaccessories.com/blogs/news/tesla-fsd-v13-navigating-the-architectural-divide-between-hw3-and-ai4>
6. SBVA Invests €30 Million in Yann LeCun-Founded AMI to Pioneer the Era of World Models, accessed on March 12, 2026, <https://www.taiwannews.com.tw/en/news/6318181>
7. World Models Race 2026 | Introl Blog, accessed on March 12, 2026, <https://introl.com/blog/world-models-race-agi-2026>
8. Yann LeCun just raised \$1bn to prove the AI industry has got it wrong, accessed on March 12, 2026, <https://thenextweb.com/news/yann-lecun-ami-labs-world-models-billion>
9. Turing Winner LeCun's New 'World Model' AI Lab Raises \$1B In Europe's Largest Seed Round Ever - Crunchbase News, accessed on March 12, 2026, <https://news.crunchbase.com/venture/world-model-ai-lab-ami-raises-europes-largest-seed-round/>
10. World Models Revolution: Yann LeCun's AMI Labs Secures \$1.03 Billion for Groundbreaking AI, accessed on March 12, 2026, <https://www.mexc.com/news/892749>
11. Tesla FSD V14 Lite Coming Summer 2026: What AI3 Owners Need to Know - basenor, accessed on March 12, 2026, <https://www.basenor.com/blogs/news/tesla-update-v14-lite>
12. Why AI's next phase will likely demand more computational power, not less -

- Deloitte, accessed on March 12, 2026,  
<https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2026/compute-power-ai.html>
13. Gemini Robotics — Google DeepMind, accessed on March 12, 2026,  
<https://deepmind.google/models/gemini-robotics/>
  14. Gemini Robotics 1.5: Pushing the Frontier of Generalist Robots with Advanced Embodied Reasoning, Thinking, and Motion Transfer - Googleapis.com, accessed on March 12, 2026,  
<https://storage.googleapis.com/deepmind-media/gemini-robotics/Gemini-Robotics-1-5-Tech-Report.pdf>
  15. Gemini Robotics brings AI into the physical world - Google DeepMind, accessed on March 12, 2026,  
<https://deepmind.google/blog/gemini-robotics-brings-ai-into-the-physical-world/>
  16. Gemini Robotics: A new era of AI-Powered Robots Plain Concepts, accessed on March 12, 2026, <https://www.plainconcepts.com/gemini-robotics/>
  17. How we built the new family of Gemini Robotics models - Google Blog, accessed on March 12, 2026,  
<https://blog.google/products-and-platforms/products/gemini/how-we-built-gemini-robotics/>
  18. Qualcomm and Other Industry Leaders Commit to 6G Trajectory Towards Commercialization Starting from 2029 Onwards, accessed on March 12, 2026,  
<https://www.qualcomm.com/news/releases/2026/03/qualcomm-and-other-industry-leaders-commit-to-6g-trajectory-towa>
  19. The Future of 6G Training: How Qualcomm's Global Coalition Is Preparing the Industry for 2029 - 5GWorldPro, accessed on March 12, 2026,  
<https://5gworldpro.com/blog/2026/03/02/6g-training-qualcomm/>
  20. Towards 6G C-V2X Networks: A Comprehensive Survey on Mobility Management, Multi-RAT Coexistence, and Machine Learning (3M) Framework for C-ITS - MDPI, accessed on March 12, 2026, <https://www.mdpi.com/2079-9292/15/5/1042>
  21. Call for Workshop Papers | IEEE Wireless Communications and Networking Conference (WCNC) 2026, accessed on March 12, 2026,  
<https://wcnc2026.ieee-wcnc.org/call-workshop-papers>
  22. Humanoid Robotics Market in 2026 Transformative Trends and Technological Advancements - Reddit, accessed on March 12, 2026,  
[https://www.reddit.com/r/robotics/comments/1qxe7rq/humanoid\\_robotics\\_market\\_in\\_2026\\_transformative/](https://www.reddit.com/r/robotics/comments/1qxe7rq/humanoid_robotics_market_in_2026_transformative/)
  23. Xiaomi tests humanoid robots on EV factory assembly line, accessed on March 12, 2026,  
<https://roboticsandautomationnews.com/2026/03/10/xiaomi-tests-humanoid-robots-on-electric-production-line-in-automotive-factory/99418/>
  24. Securing the Information and Communications Technology and Services Supply Chain: Connected Vehicles - Federal Register, accessed on March 12, 2026,  
<https://www.federalregister.gov/documents/2025/01/16/2025-00592/securing-the-information-and-communications-technology-and-services-supply-chain-connected-vehicles>

25. Another Misstep in U.S.-China Tech Security Policy - Lawfare, accessed on March 12, 2026,  
<https://www.lawfaremedia.org/article/another-misstep-in-u.s.-china-tech-security-policy>
26. Data Sovereignty 101 for Mobile Apps in 2026 - Dogtown Media, accessed on March 12, 2026,  
<https://www.dogtownmedia.com/data-sovereignty-101-for-mobile-apps-navigating-the-2026-regulations-on-where-your-mobile-data-lives/>
27. Data Sovereignty Laws Around The World | by Robert Broeckelmann | Feb, 2026 | Medium, accessed on March 12, 2026,  
<https://medium.com/@robert.broeckelmann/data-sovereignty-laws-around-the-world-22b3a2dbc2ad>
28. Arxiv今日论文| 2026-03-10 - 闲记算法, accessed on March 12, 2026,  
[http://lonepatient.top/2026/03/10/arxiv\\_papers\\_2026-03-10](http://lonepatient.top/2026/03/10/arxiv_papers_2026-03-10)
29. The Agentic Liability: Governing the "Fire and Forget" Office ... - Reddit, accessed on March 12, 2026,  
[https://www.reddit.com/user/enoumen/comments/1rqzc6i/the\\_agentic\\_liability\\_governing\\_the\\_fire\\_and/](https://www.reddit.com/user/enoumen/comments/1rqzc6i/the_agentic_liability_governing_the_fire_and/)
30. (PDF) Agentic Reward Modeling: Verifying GUI Agent via Online Proactive Interaction, accessed on March 12, 2026,  
[https://www.researchgate.net/publication/400369362\\_Agentic\\_Reward\\_Modeling\\_Verifying\\_GUI\\_Agent\\_via\\_Online\\_Proactive\\_Interaction](https://www.researchgate.net/publication/400369362_Agentic_Reward_Modeling_Verifying_GUI_Agent_via_Online_Proactive_Interaction)
31. Tesla update 2026.2.3 goes wide release, explore new features and release notes, accessed on March 12, 2026,  
<https://www.teslaoracle.com/2026/02/12/tesla-update-2026-2-3-goes-wide-release-explore-new-features-and-release-notes/>
32. Elon Musk's xAI brings 1GW Colossus 2 AI training cluster online - Teslarati, accessed on March 12, 2026,  
<https://www.teslarati.com/elon-musk-xai-brings-1gw-colossus-2-ai-training-cluster-online/>
33. (PDF) EvoCUA: Evolving Computer Use Agents via Learning from Scalable Synthetic Experience - ResearchGate, accessed on March 12, 2026,  
[https://www.researchgate.net/publication/400003283\\_EvoCUA\\_Evolving\\_Computer\\_Use\\_Agents\\_via\\_Learning\\_from\\_Scalable\\_Synthetic\\_Experience](https://www.researchgate.net/publication/400003283_EvoCUA_Evolving_Computer_Use_Agents_via_Learning_from_Scalable_Synthetic_Experience)
34. (PDF) Step-GUI Technical Report - ResearchGate, accessed on March 12, 2026,  
[https://www.researchgate.net/publication/398806674\\_Step-GUI\\_Technical\\_Report](https://www.researchgate.net/publication/398806674_Step-GUI_Technical_Report)
35. NVIDIA and Nebius Partner to Scale Full-Stack AI Cloud, accessed on March 12, 2026,  
<https://nvidianews.nvidia.com/news/nvidia-and-nebius-partner-to-scale-full-stack-ai-cloud>
36. Navigating the global battle over AI independence - WP Intelligence, accessed on March 12, 2026,  
<https://wpintelligence.washingtonpost.com/topics/2026/03/09/navigating-global-battle-over-ai-independence/>