Envisioning the future of mobility powered by AI

Kearney Foresight
Advanced Mobility Institute

MIT Mobility Initiative

KEARNEY FORESIGHT
Advanced Mobility Institute

The Advanced Mobility Institute, part of Kearney's Foresight network, helps companies, governments, and nonprofits implement innovative mobility solutions. We work across the entire mobility ecosystem to modernize and improve transportation with a uniquely data-driven approach. Let us help you transform how people move. More accessible, more sustainable, more efficient.

MIT Mobility Initiative

The MIT Mobility Initiative (MMI) is a global platform to accelerate a safe, clean, and inclusive mobility system. MMI brings together researchers, industry leaders, policymakers, and entrepreneurs to advance the future of transportation through cutting-edge research, world-class education, technology commercialization, and high-impact engagement.

Executive summary

This report provides an overview of how artificial intelligence is shaping the mobility sector and how it could do so in the future. It also examines how executives and policymakers should think about these transformations.

The findings are based on a joint study conducted by the MIT Mobility Initiative (MMI) and Kearney Advanced Mobility Institute (KAMI). The research aims to equip public- and private-sector leaders with clear insights into where artificial intelligence is already delivering measurable value across mobility systems. It also outlines the key enablers that will be essential to accelerating adoption, integration, and large-scale impact in the years ahead.

The research team engaged with 55 organizations across the global mobility ecosystem, capturing a broad spectrum of perspectives from both advanced and emerging markets. This diversity ensured that insights reflected the full range of challenges and opportunities shaping the future of AI in mobility worldwide.

The study generated five key takeaways:

The AI landscape in mobility is fragmented but has vast potential. AI is already embedded across the mobility value chain, powering applications from network planning to new materials discovery to crowd monitoring. However, implementation patterns are scattered, with limited alignment between impact and feasibility. As a result, most value creation is confined to isolated pilots rather than scaled, integrated systems capable of transforming mobility at large.

The bigger the potential, the tougher the execution. AI's transformative potential grows as applications move from task-level tools to system-level integration, but so does complexity. Real transformation will depend on coordinated action across public and private stakeholders who are aligned on a common vision for safe, clean, and inclusive mobility.

The safety criticality of mobility requires careful design of the human-AI interface. AI excels at some tasks and lags in others, in ways that are not always intuitive to humans. In mobility, a safety-critical field, this dynamic "jagged frontier" heightens risk. Ensuring safety requires intentional hybrid systems that define how humans and AI interface, with some functions evolving toward AI-led or human-led models.

Ecosystem collaboration is the next primary driver of progress. Progress depends less on isolated innovation and more on coordinated execution. Unlocking AI's value requires system-wide alignment on strategic intent and joint action on four core enablers: data infrastructure, regulation, standards, and governance.

Multiple regional end states can emerge amid uncertainty over collaboration and trust. The future of AI remains uncertain as offering/industry maturity/collaboration and public trust stand at a crossroads. Uncertainty is amplified across regions, shaped by regulatory and institutional contexts. Multiple AI end states can emerge, requiring mobility enterprises to tailor their AI strategies.

The report is grounded in the premise that a fundamental test of any modern mobility system is how well it delivers on three societal priorities: safety, inclusivity, and sustainability. In all three, there are significant challenges still to be addressed. There is a growing sense that AI may enable transportation systems to overcome these challenges.

This research identified 39 AI application areas in the mobility sector. These were categorized into seven groups that span the entire mobility value chain, offering a useful structure for understanding how AI is transforming the sector. This report systematically evaluates these applications, assessing their potential impact for mobility systems.

The report concludes with an assessment of the primary enablers of successful AI adoption within organizations—enablers that some leading transit agencies and mobility enterprises are already deploying effectively.

Introduction: the priorities and pitfalls of modern mobility

Over the past three centuries, advanced technologies have repeatedly reshaped our understanding of mobility. Through the development of railroads, automobiles, and airplanes, we have experienced a series of technology-driven revolutions in transportation. This ongoing cycle of innovation continues with the arrival of broadly available artificial intelligence.

This report provides an overview of the most promising AI applications in the mobility industry, how mobility players can use AI to optimize their business, and the potential value creation of AI for the mobility industry and society.

Mobility systems provide access to the goods and services on which we depend—employment, healthcare, education, and more. It is not only a physical system of vehicles and infrastructure, but also an economic and social system that determines who has access to life's opportunities and who doesn't. A well-functioning mobility ecosystem connects people and markets efficiently, safely, and affordably; a failing one creates barriers to growth, inclusion, and well-being.

The fundamental test of any modern mobility system is how well it delivers on three societal priorities; safety, inclusivity, sustainability. Each represents an essential dimension of human welfare. In all three, there are significant gaps still to be closed—potentially with the assistance of AI and related technologies (see figure 1 on page 3).

Safety

First and foremost, safety refers to the capacity of a transportation system to prevent fatalities and serious injuries and to protect people even when errors occur. A safe system is designed to anticipate risk rather than to react to it, combining vehicle design, infrastructure, enforcement, and behavior.

One example where progress has been particularly slow is road safety, where the status quo is costly in every sense of the word. Each year, around 1.19 million people worldwide die in traffic accidents, and tens of millions more are injured or disabled (WHO, 2023). Roadway traffic injuries are the leading cause of death among children and young adults ages 5 to 29 years, while the economic cost of such incidents is approximately 3 percent of the GDP in most countries (WHO, 2023).

Although vehicle standards, road design, and emergency response have improved, the global target of halving deaths and injuries by 2030 appears out of reach; the continued persistence of accidents indicates that traditional engineering and enforcement measures have reached diminishing returns in the face of complex traffic dynamics and human error.

Figure 1
The mobility industry has three priorities, but there are significant gaps to close

Safety
Protecting people from harm and minimizing risks during travel

~1.19 million
road fatalities globally

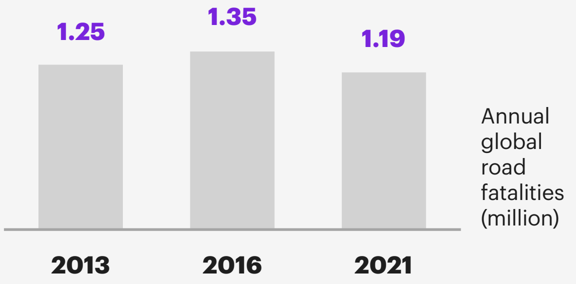

Inclusivity
Providing fair, affordable access and removing barriers for participation in mobility

>30%
of income spent on mobility by low-income and vulnerable households

Sustainability
Adopting technologies and practices that minimize environmental impact and conserve resources

~11%
share of urban transport in global energy-related CO_2

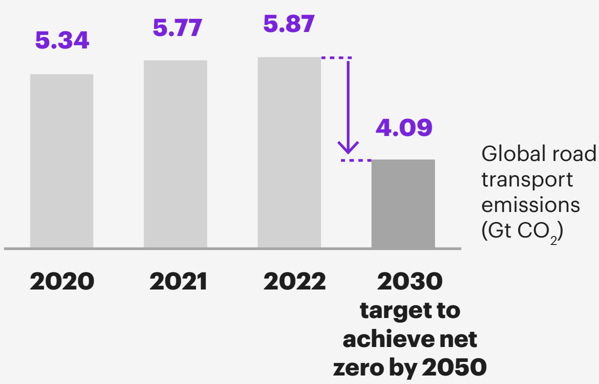

Sources: WHO, IEA,UITP; Kearney and MIT analysis

Inclusivity

This refers to whether mobility systems are practically affordable and equitably available to all segments of society. Access is measured not only in fares or service coverage, but also in how readily people from various economic and demographic categories may physically reach essential services and economic opportunities. Affordability is measured not only in terms of money (transit fares, vehicle prices), but also in terms of time (transit service delays, traffic congestion delays).

In many parts of the world, the burden of travel costs on low-income and vulnerable households is disproportionately high—often more than 30 percent of household budgets (UITP, 2025). This translates directly into lost access to jobs, education, and health care. In both developed and emerging cities, peripheral and low-income districts continue to receive fewer and less reliable services. The result is a systematic inequality of mobility that reinforces a broader social divide.

Then there are the physical barriers to accessible transport. Poorly designed infrastructure, limited connectivity, and unsafe facilities (such as bus stops and subway stations) too often deter participation by the elderly, women, and people with disabilities.

Sustainability

In this context, sustainability refers to the ability of a transportation system to meet mobility needs while keeping its environmental and resource impacts as minimal as possible, whether in terms of energy use, CO_2 emissions, air pollution, noise, or land consumption.

Urban transport remains one of the largest sources of greenhouse gas emissions, responsible for roughly one-tenth of all energy-related carbon dioxide (CO_2) output worldwide, with road traffic contributing the majority of emissions within the transportation category (IEA, 2023). Despite technological progress in areas such as powertrain electrification and fuel efficiency, global emissions from the sector continue to grow alongside rising travel demand and freight volumes (OECD, 2021).

Additionally, countries face institutional barriers to sustainability. Fragmented responsibilities among agencies and limited data-sharing reduce the efficiency of mobility operations, with negative effects on environmental resilience (World Bank, 2024).

To align the sector with the Net Zero Emissions by 2050 Scenario, road transport CO_2 emissions must decline by almost one-third by 2030, yet since 2015, they have only increased, rising by 3.5 percent, indicating that the sector is still not on track with global decarbonization goals (IEA, 2023).

To align with the Net Zero Emissions by 2050 Scenario, road transport CO_2 emissions must decline by almost one-third by 2030.

Mobility is an ecosystem

These gaps in safety, inclusivity, and sustainability are persistent, and they have deep structural origins. Mobility systems are highly complex systems, in which traditional interventions—such as building more roads, upgrading vehicle fleets, or adjusting traffic-safety laws—can deliver some incremental gains, but not the systemic transformation that would bring truly meaningful progress in any of these areas.

It is important to acknowledge that mobility is a multi-actor system, one that depends on the interaction of users, such as motorists, passengers, and delivery customers, and four broad categories of supply-side players that plan, fund, build, operate, and regulate assets and services (see figure 2). Amid all this variety, misaligned incentives, and potentially clashing interests, truly systematic transformation in the mobility sector is difficult to achieve.

AI as a systemic mobility enabler

Within the mobility sector, there is a growing sense that artificial intelligence may serve as a true systemic enabler, a capability layer that reconfigures how entire transportation systems perceive, learn, and act.

Private investment in AI has expanded enormously, and organizational use has moved into the mainstream; around three-quarters of companies now report using AI in at least one business function (QuantumBlack, 2024) (see figure 3 on page 6).

Figure 2
Urban mobility is an ecosystem shaped by demand and the interactions of five types of supply-side actors

Demand	**Mobility users**	— Private travelers: passengers and private-vehicle owners — Businesses relying on mobility (delivery, etc.)
Supply	**Policymakers, regulators, trade associations, and public transport authorities**	— City and national transport authorities and regulatory institutions — Industry associations and standard-setting bodies — Public agencies and international organizations, etc.
	Infrastructure developers and operators	— EPC contractors — Expressway, toll road, tunnel concessionaires — EV charging network operators, etc.
	Mobility operators, from transit agencies to rideshare firms	— Traditional public transit providers — Shared mobility (bikes, scooters, and cars) — Ride-hail and taxi fleets, etc.
	Entities that provide the equipment, services, and technology that the mobility sector depends upon, including manufacturers, suppliers, financial institutions, and tech companies	— Vehicle and rolling-stock manufacturers — Component suppliers — Financial institutions and investors — Technology firms, mobility start-ups, etc.

Notes: EPC is engineering, procurement, and construction; EV is electric vehicle.
Source: Kearney and MIT analysis

Figure 3

The rise of AI is creating new opportunities to address traditional mobility challenges

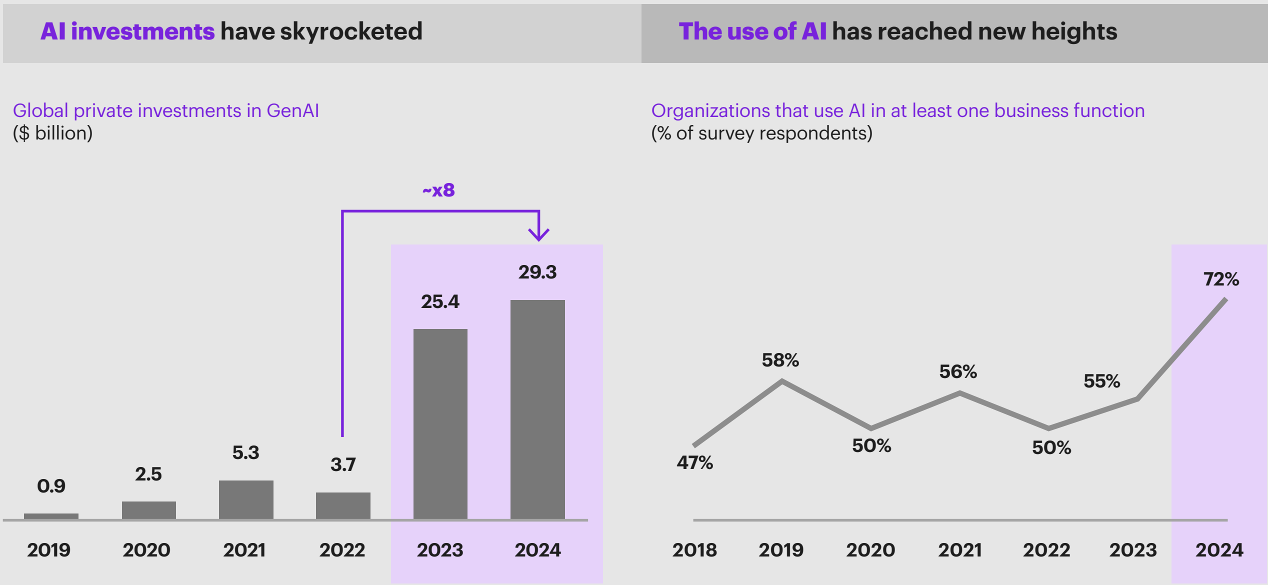

[1] Average causal impact a treatment has on an outcome, measured as the difference between the results of a treatment group and a control group

Source: Stanford HAI, QuantumBlack; Kearney and MIT analysis

AI has swiftly demonstrated its usefulness for mobility systems. Computer vision can process data from vehicles and networks far faster than any human operator and can detect patterns that would otherwise remain hidden. Predictive models can anticipate demand peaks, forecast incidents, or identify emerging risks before they escalate. Optimization algorithms can coordinate traffic signals and fleets in real time.

In short, AI provides a way to manage complexity at the speed and scale at which complexity actually occurs. And this has already generated meaningful gains across all three of the priority areas identified in the introduction.

On the **safety** front, AI can anticipate and prevent collisions and identify high-risk areas of the network to guide safer infrastructure design.

With regard to **inclusivity**, it can better measure and predict demand in underserved areas, support inclusive journey planning, provide multilingual assistance, and develop pricing models that make transport more affordable and accessible.

To enhance **sustainability**, AI can optimize traffic flow, plan low-emission routes, and better optimize the charging of electric vehicles (EVs) to align with cleaner energy use.

In a subsequent section of this report, we will explore the mobility applications of AI in greater detail. However, before moving further, we need to acknowledge the fact that considerable uncertainty persists regarding the technology's maturity, capability, and broader impacts (see figure 4).

Practitioners across the mobility sector have expressed various misgivings about the reliability of AI in various transportation contexts as well as the ultimate value of the immense investments that technology demands, the cost versus the benefit of AI, and much more.

AI systems excel across a wide range of tasks but continue to show structural limits. AI systems can 'hallucinate', providing plausible but incorrect information. Several leading systems show limits in abstract reasoning and remain vulnerable to security issues, such as prompt injection, in which a hacker introduces a prompt that induces an AI to generate a response different from the one intended by the user. As a result of these and other issues, fully scaled and integrated AI deployments in business are still not common, despite the widespread adoption of these technologies.

AI presents significant hazards just as surely as it offers exciting solutions. With regard to safety, it's not yet clear whether autonomous driving can be made acceptably safe at sufficient scale across varying road, weather, and traffic conditions. On inclusivity, the demonstrated biases of some AI models could, if left unchecked, lead to the creation of systems that disfavor individuals and communities from certain demographic or economic categories.

Regarding sustainability, the enormous energy and resource demands of AI data centers make it far from clear that the gains of this technology will outweigh the vast costs: the combined emissions of top AI systems are estimated to reach approximately 100 million tons CO_2 equivalent per year if they are left to grow at current projections (ITU, 2025).

At a technical level, AI models excel across a wide range of benchmarks, yet their complex reasoning functions—such as long-horizon planning and compositional logic—remain inadequate (Stanford HAI, 2025). Robustness is also a concern since AI systems can fail when conditions shift; for example, closing just 1 percent of the streets in an AI navigation test caused its accuracy levels to drop from nearly 100 percent to only around 67 percent (MIT, 2024).

Figure 4
Considerable uncertainty is associated with AI's maturity, power, and the broader impact

AI tools are abundant, but **scaled deployments** in business **can be rare**	AI systems excel across a wide range of tasks but continue to **show structural limits**	Rising **costs** and their **environmental footprint** are under scrutiny
Only 5% of custom GenAI tools that organizations investigate bring a return on investment. Pilots stall because of a lack of contextual learning and poor integration.	Frontier LLMs show limits in abstract reasoning and out-of-distribution generalization and are vulnerable to security issues (e.g., prompt injection).	The combined emissions (training, operations, etc.) of the top AI systems could reach 102.6 million tons CO_2 equivalent per year alongside rising investment demands.
Of all task-specific GenAI products offered by vendors, what % are companies exploring, piloting, or fully scaling?	Accuracy of AI tools in solving planning problems (% of instances correct, PlanBench benchmark)	Estimated carbon emissions from training select AI models, tons CO_2 equivalent

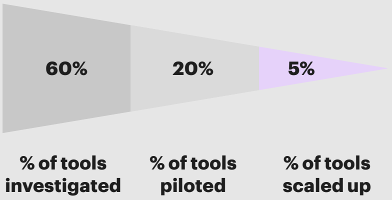

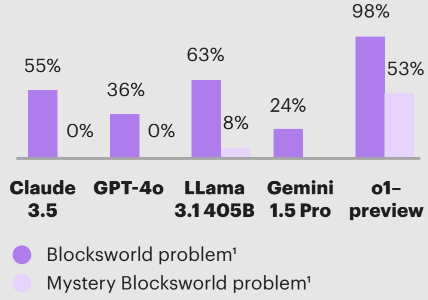

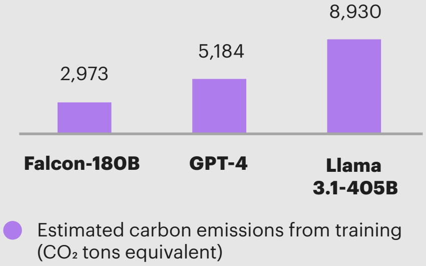

[1] AI planning problem where a robot arm must rearrange blocks to reach a goal configuration. Mystery Blocksworld is a harder version that uses the exact same logic but replaces all names with random words, testing if an AI understands the underlying planning problem or just memorizes patterns. Note: LLMs are large language models.

Sources: MIT Networked Agents and Decentralized AI (NANDA), HAI, ITU; Kearney and MIT analysis

Yet despite these and other uncertainties, momentum is building toward the heightened use of AI across the mobility sector. There are several reasons for this. Most fundamentally, the technological maturity of AI—particularly in areas such as computer vision, natural language processing, and reinforcement learning—has reached an inflection point at which industrial-scale deployment is increasingly feasible.

Applications such as AI-based navigation, real-time traffic modeling, autonomous driving, and predictive asset management are no longer confined to research labs or pilot projects; they are entering live commercial operations.

Another essential factor is the significant capital that investors and companies behind AI strategies are directing toward mobility-sector AI in recognition of the enormous potential for the development of innovative business models. The current surge in experimentation, while marked by uncertainty, also presents a unique opportunity for organizations to establish a first-mover advantage, define competitive standards, and secure a foothold in a potentially immense global market.

In addition, the societal imperative for improvement across the three priority areas of safety, inclusivity, and sustainability continues to build. Stricter climate targets, worsening urban congestion, more safety incidents, and persistent transport inequities are intensifying pressure on governments and companies alike to modernize the mobility system.

AI is not an automatic solution for any of these—and, as indicated earlier, it creates some additional problems of its own—but it is a potent new weapon in the policy and innovation arsenal, capable of optimizing fleets for carbon efficiency, making mobility accessible for underserved populations, and enhancing safety at scale.

On top of all of these factors, there is this simple fact: the AI policy landscape is in enormous flux, and this is therefore a defining moment for the future of the technology. Governments and private companies are drafting legislation, standards, and frameworks to regulate AI in transportation. This offers a rare chance for industry leaders to engage with regulators, shape balanced policies, and align innovation with societal values before more rigid systems are entrenched.

In light of all of these factors, leaders across the sector are realizing that now is the time to explore the potential of AI in mobility. Important questions about impact, costs, and social effects remain, and they are best answered through careful, evidence-based work, as we will demonstrate throughout the remainder of this report.

Purpose, scope, and methodology of the study

Our findings are the result of a combined effort of the MIT Mobility Initiative (MMI) and Kearney Advanced Mobility Institute (KAMI), which collaborated on a joint study to explore how AI is being used across the mobility value chain and where it can credibly generate meaningful advances.

This study is designed to support decision-makers in both the public and private sectors by identifying where value is already being created with AI and which enablers will bring AI to a new level of utility and effectiveness across various transportation modes.

From the outset, our research aimed to move beyond theory by grounding the findings in the practical realities of how AI is being developed, deployed, and governed across the sector. To that end, the team engaged with 55 institutions from all continents, ensuring diverse perspectives from both advanced and emerging mobility markets.

The participants represented four segments of the mobility value chain:

— Policymakers, regulators, and public transport authorities

— Infrastructure developers and operators

— Mobility operators

— Manufacturers, suppliers, financial institutions, and technology companies.

Within each segment, the study collaborated with leading organizations shaping the future of mobility, such as Google, Lyft, Uber Freight, Flixbus, Deutsche Bahn, NEOM, Dubai Taxi Company, and MIT itself, among others.

The participant base was highly international, reflecting the global nature of the challenges and opportunities addressed (see figure 5 on page 9). This geographic distribution ensured representation of distinct regulatory, economic, and technological contexts, providing a balanced and comparative view of how AI can transform mobility systems worldwide.

Figure 5

Our study engaged 55 leading mobility institutions via interviews, workshops, and surveys

Participants by region

55 **Organizations**

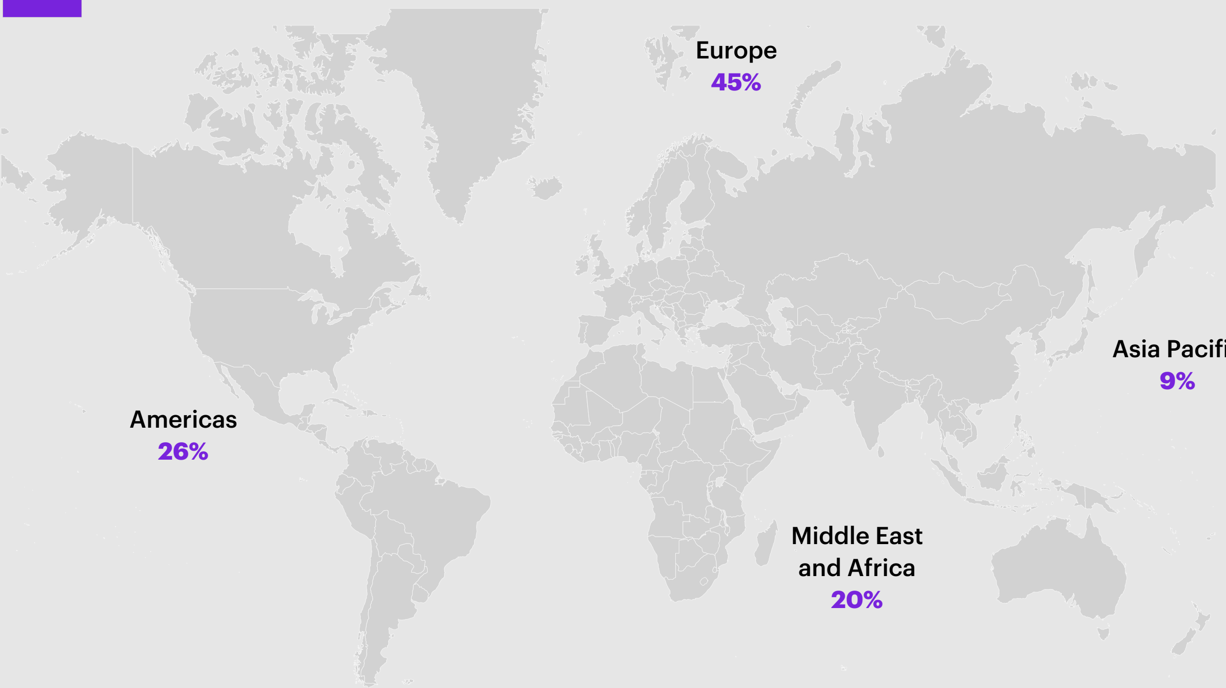

Policy makers, regulators, associations and public transport authorities

- Al Madinah Region Development Authority (MDA)
- EMT Madrid (Empresa Municipal de Transportes de Madrid)
- Hamburger Verkehrsver-bund (HVV)
- International Association of Public Transport (UITP)
- KSA Ministry of Transport and Logistic Services
- New York City Transit (MTA NYCT)
- San Francisco Municipal Transportation Agency (SFMTA)
- Toronto Transit Commission (TTC)
- traffiQ Frankfurt
- U.S. Department of Transportation (USDOT)
- Verkehrsverbund Rhein-Ruhr (VRR)
- Washington Metropolitan Area Transit Authority (WMATA)

Infrastructure developers and operators

- AlUla
- Cintra
- Edison
- EYSA
- Ferrovial
- Gensler
- Globalvia
- Iridium
- NEOM
- Parkin
- Remat Al-Riyadh
- WSP

Mobility operators

- Abertis
- Beacon Mobility
- Berliner Verkehrsbetriebe (BVG)
- ComfortDelGro
- Deutsche Bahn
- Dubai Taxi Company (DTC)
- Etihad Rail
- FlixBus
- Free To X
- Highland Electric
- Mobico
- Prasarana Malaysia Berhad
- Rail Europe
- RATP Dev
- Saudi Public Transport Company (SAPTCO)
- Seur
- SNCF
- Stagecoach
- Toll Group
- Tokyu Corporation
- Via

OEMs, suppliers, and financial and technology players

- Denso
- EV8
- Google
- LKQ
- Lyft
- Mutua Madrileña
- Torc Robotics
- Uber Freight

Sources: Kearney and MIT analysis

KAMI and MMI employed a parallel mixed-methods approach that integrated both qualitative and quantitative research designs (see figure 6).

The study proceeded in four phases:

Secondary research (November 2024–January 2025). This phase focused on building the theoretical foundation for the study. The team reviewed industry publications, academic research, and expert commentary and performed several workshops with MIT faculty and researchers to establish an analytical framework for AI in mobility and develop a long list of potential AI applications. Each application was explored across the significance of the underlying problem, the potential impact of AI in addressing it, the technical feasibility of implementation, and the degree of industry adoption; findings were benchmarked across modes, activities, and enabling capabilities. Based on this secondary research and subsequent team scoring and review, the study arrived at a prioritized list of 39 AI applications areas that shaped the next stages of the work.

Interviews and survey (February–July 2025). During this phase, the team engaged with the global mobility ecosystem to incorporate real-world perspectives and practical insights on AI applications into the study. The team conducted 55 semi-structured expert interviews using a predefined questionnaire to assess AI use and readiness across the mobility value chain. Each interview explored the respondent's role and organizational AI strategy, the positioning of AI from development to deployment, and the most relevant applications in practice. Discussions also examined key challenges and enablers for AI adoption, including regulation, technical maturity, financing, and talent. A short follow-up survey was then distributed to the same group of experts to obtain structured assessments of perceived impact and feasibility for selected applications. During this phase, the study team also conducted several working sessions and brainstorming workshops with MIT faculty and researchers to refine the theoretical framework and test its consistency against insights gathered from the interviews and survey results.

Figure 6
We drew on diverse sources and integrated qualitative and quantitative research designs

Study methodology

Secondary research

Develop an initial perspective on use cases, impact, feasibility, etc.:
- Review of industry publications, research documents, and expert input
- Benchmarking across the industry

Interviews and surveys

Assess use case relevance, impact, feasibility, and challenges, etc.:
- Extensive interviews and workshops with 55 organizations across the mobility value chain
- Survey of participants

Study team and cross-industry workshops

Generate key insights and scenarios:
- Regular expert and study team workshops
- Scenario planning workshops with selected participants from the mobility ecosystem

Source: Kearney and MIT analysis

Report development (August–October 2025). This phase consolidated the evidence base and applied the established analytical framework to contextualize the AI applications. Each AI application was assessed based on problem significance, expected impact, and technical feasibility, leading to convergent findings and cross-cutting themes. Participant feedback was also collected to test accuracy and relevance of the findings. In addition, a two-hour scenario-planning workshop was organized with 25 participants representing public transit agencies, mobility operators, technology firms, and academia. The session focused on identifying drivers of impact and adoption of AI in mobility over the next decade and outlining pathways for bringing AI in mobility to the next level.

Study scope. For the purposes of this study, we focus on surface-based urban passenger mobility: transport modes that operate predominantly within and around urban boundaries and serve the movement of people rather than goods.

Specifically, the scope includes three areas:

Road mobility. This includes micromobility (e-scooters and bikes), cars (private, shared, and automated), and buses. Trucks, while part of urban logistics, are not a primary focus due to the study's emphasis on passenger transport, though several of the study's findings also apply to trucks.

Rail mobility. This includes metro and other rail transit systems. Freight and high-speed rail, which operate primarily on intercity or long-distance routes, are not a primary focus, though many of the implications of the study are applicable to them as well.

Water mobility. This includes water taxis (intra-city) and water ferries (inter-city) where they form part of the urban passenger network.

Air mobility, such as traditional aircraft, electric vertical take-off and landing vehicles (eVTOLs), and parcel drones serving urban and peri-urban areas, are excluded from the scope of this study. Traditional freight and logistics are also excluded. Other modes, such as ropeways and aerial trams within cities, are acknowledged where they contribute to urban passenger transport networks.

Defining AI and its potential for modern mobility systems

At its core, "artificial intelligence" refers to the development of systems that can perform tasks typically requiring human intelligence, such as reasoning, learning, and problem-solving (MIT Media Lab, 2025). What sets modern AI apart from traditional software or related fields such as data science is its ability to learn and adapt with a certain degree of autonomy.

Thus, AI allows a machine or software to perceive its environment, learn or reason from that data, and then decide on actions to achieve certain goals. Crucially, AI systems can improve their performance over time as they gain more data and experience. They are constantly learning.

This recognition of AI's capacity for autonomous learning and self-improvement is crucial to our consideration of the technology's applications in the mobility sector, as is a clear taxonomy of AI's subvariants in the marketplace.

Artificial intelligence is an umbrella term, under which fall more specific fields, including machine learning (ML), deep learning (DL), and generative AI (GenAI). The easiest way to understand their relationship is as a hierarchy of subsets (IBM, 2025) (see figure 7):

Figure 7
AI refers to systems that can perform tasks typically requiring human intelligence

At its core, artificial intelligence refers to the development of systems that can perform tasks typically requiring human intelligence, such as reasoning, learning, and problem-solving.

AI definition framework (IBM)		Examples: General	Examples: Mobility
1950s	**Artificial intelligence** Human intelligence exhibited by machines	Chess and board-game programs	Early navigation algorithms
1980s	**Machine learning** AI systems that learn from historical data	Spam filters that learn from examples of spam vs. non-spam	Models forecasting metro ridership based on historical passenger data
2010s	**Deep learning** Machine-learning systems that mimic human brain function	Facial recognition systems in smartphones unlocking the device	Computer vision systems recognizing pedestrians and traffic signs
2020s	**Generative AI** Deep learning models (foundational models) that create original content	Image generation from text prompts	Systems imagining what a new vehicle design might look like

Agentic AI is an emerging concept (though its origins can be traced back to 1950s), now used to describe a new wave of systems that combine generative AI + autonomous reasoning, planning, and action at scale.

Sources: IBM; Kearney and MIT analysis

Artificial intelligence

Artificial intelligence (AI) is the broadest category, encompassing any technique that enables computers to mimic human intelligence. Early examples, predating the rise of machine learning, include rule-based systems such as chess and other board-game programs as well as early navigation and path-planning algorithms developed for robotics and logistics.

Machine learning

Machine learning (ML) is a subset of AI focused on algorithms that learns from data. Rather than being explicitly programmed for every scenario, ML systems improve through experience. A classic definition by ML pioneer Arthur Samuel describes it as "the field of study that gives computers the ability to learn without being explicitly programmed" (IBM, 2020). In practice, this means an ML model can be trained on historical data to statistically derive useful functions without a human engineer predefining all the rules. For example, ML can be applied to classify emails as spam or non-spam or—in mobility contexts—to analyze historical traffic patterns and vehicle-sensor readings to predict travel times or classify objects on the road.

AI is an umbrella term, under which fall a variety of more specific fields.

Deep learning

Deep learning (DL) is a subset of ML and refers to systems that use multilayered neural networks inspired by the human brain. These networks are composed of many interconnected "neurons" (mathematical functions) capable of learning hierarchical representations of data. The use of many layers of neurons distinguishes DL from other ML approaches (IBM, 2025).

The "depth" of these layers enables the system to identify very complex patterns. DL has been a primary driver of recent breakthroughs in AI—from image recognition to speech understanding—as it has evolved with larger datasets and advances in computing power.

Unlike earlier approaches that relied on manually designed features, deep neural networks learn directly from raw sensor data, achieving superior performance in most perception and prediction tasks. These networks have become the catalysts of everyday technologies such as facial recognition on smartphones, and they are increasingly central to mobility applications.

For example, they power the computer-vision systems that enable autonomous vehicles to detect pedestrians and traffic signs and interpret camera inputs to support routing and navigation decisions. DL neural networks also generate the prediction models that forecast conditions such as traffic congestion, public-transit demand, or network disruptions.

Generative AI

Generative AI (GenAI) is a considerable advance in AI technologies, one with considerable ramifications for the mobility sector. Formally, GenAI refers to "deep-learning models that can take raw data and learn to generate statistically probable outputs" (IBM, 2023).

Whereas most ML and DL applications in mobility have been analytical or predictive, GenAI (a subset of DL) focuses on creating new content or data. Rather than outputting a yes/no answer or a numerical prediction, generative models output new artifacts—a plausible sentence, a realistic picture, a snippet of audio, a design schematic, or other required item.

Among the most prominent subtypes of GenAI are large language models (LLMs) such as Open AI's ChatGPT or Anthropic's Claude. After training on vast amounts of text, an LLM can generate human-like textual responses, answer questions, or even draft reports (though not this one). Generative models are also increasingly capable of producing synthetic images or videos, an aptitude with obvious potential for such mobility applications as vehicle design.

Although GenAI tools have become ubiquitous and global private investment in GenAI has risen roughly eightfold since 2022 (Stanford HAI, 2025), large-scale business deployments are still rare: after reviewing hundreds of GenAI enterprise initiatives, MIT's Project Networked AI Agents in Decentralized Architecture (NANDA) found that only 5 percent reach production with measurable return on investment, with pilots often stalling due to a lack of contextual learning and poor integration (MIT NANDA, 2025).

Agentic AI is an emerging concept that combines GenAI with autonomous reasoning, planning, and action. These agentic models represent a new frontier, where AI not only generates content but also initiates, sequences, and executes tasks with limited human supervision—a concept we will explore in greater detail later in this report.

While the foundations of AI date back over 70 years, its capabilities have accelerated dramatically in the past decade, driven by advances in algorithms, computing power, and data scale.

Training computation capacity historically grew by about 1.5× per year, but since 2010, it has accelerated to roughly 4.2 to 4.5× annually (Epoch AI, 2025). This acceleration underpins the emergence of large-scale architectures such as ChatGPT-4, DeepSeek-V3, and DALL-E, each trained on unprecedented volumes of multimodal data, including text, images, and sensor records.

This growth in capacity has translated directly into measurable performance gains across demanding benchmarks. For example, on the Measuring Massive Multitask Language Understanding (MMLU), which assesses reasoning across diverse domains, accuracy rose from about 44 percent (ChatGPT-3, 2021) to approximately 91 percent (ChatGPT-5, 2025), exceeding human-expert levels in the tested disciplines (Stanford HAI, 2025).

Beyond benchmark performance gains, contemporary AI tools are unlocking entirely new capabilities, such as high-quality video generation and editing, motion synthesis, and multimodal understanding. AI today is entering a stage where its capabilities are approaching human-level performance across a growing number of complex tasks. This evolution lays the foundation for transformative applications in mobility.

In short, AI is no longer a distant promise. It is already reshaping mobility across the globe, perceiving, predicting, and acting across vehicles, infrastructure, and transportation networks, improving how people and goods move from place to place.

But the implementation of AI in mobility sector will initially be nuanced since the technology must navigate the sector's unique structural realities. Transportation is an inherently physical process, involving the movement of objects at speed through space. Therefore, every prediction or decision must at all times respect the immutable physical limits of infrastructure, vehicles, and the human body.

In addition, the management of transportation networks is distributed across governments, private companies, and individual users, with all actors interacting and creating perpetual ripple effects for one another. AI must therefore function across heterogeneous, interoperable systems rather than optimizing for single entities.

The “jagged frontier” and human–AI collaboration

The concept of AI’s “jagged frontier” captures one of the most interesting and important characteristics of current artificial intelligence systems, i.e. their dramatically uneven performance across tasks that appear to require similar levels of intelligence or difficulty to humans. That frontier is most readily apparent when AI systems demonstrate excellent performance on certain complex tasks while failing at seemingly simpler problems (see figure 8).

For example, large language models can write sophisticated poetry, solve complex mathematical proofs, and engage in nuanced philosophical discussions. Yet they also struggle (at least, as of this writing) with basic spatial reasoning, maintaining consistency across a conversation, or avoiding certain elementary logical errors that no human of standard intelligence would commit.

The jagged frontier seems to follow a certain pattern. AI systems perform well in fields where there are large amounts of high-quality training data and where success can be measured through pattern matching or statistical correlation.

Where they currently struggle is with tasks requiring common-sense reasoning, a clear understanding of cause and effect, or an ability to generalize from limited examples in ways that most humans find intuitive. Humans excel at tasks requiring common sense, ethical reasoning, creative problem-solving in novel domains, and understanding of context and nuance. We routinely draw upon these cognitive advantages in many of our daily mobility tasks, such as driving, riding a bicycle, or strolling down a sidewalk.

Figure 8
Overcoming the challenges and achieving a lasting impact will require system-level AI and optimal human–AI collaboration

The “jagged frontier” and human–AI interaction

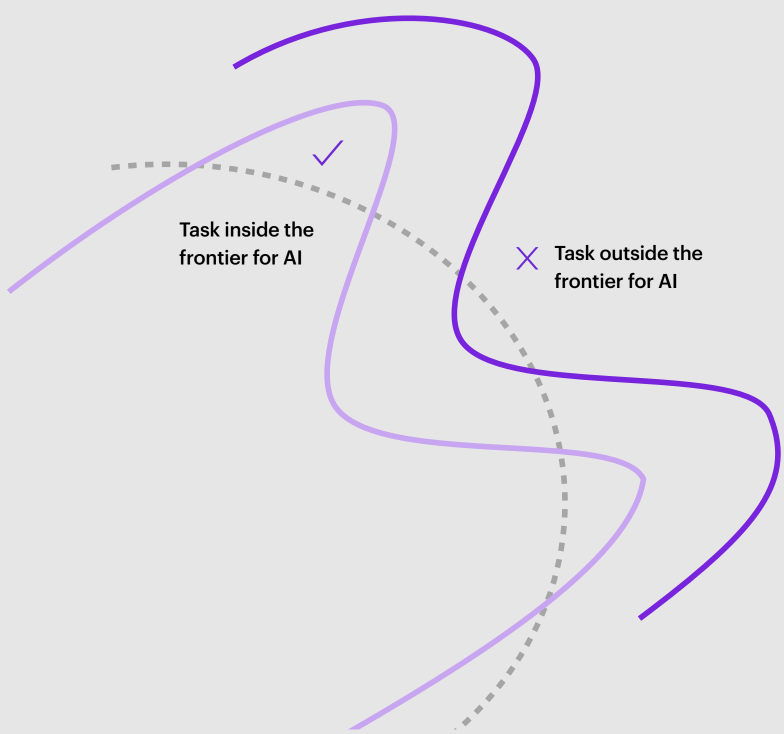

- - - Equal difficulty tasks for humans
— 2025 AI capabilities
— 2030 AI capabilities

Tasks that are of similar difficulty for humans can be of vastly different difficulties for AI. The frontier should guide human–AI interactions.

AI can struggle in such situations. An autonomous vehicle can easily "read" a billboard that contains a public-service ad calling for motorists to stop for school buses. However, if the ad were to include an image of a red octagonal stop sign, the system might misclassify it as a real road sign and stop unnecessarily in the middle of flowing traffic—a potentially catastrophic error that no lucid driver would make.

Functions that are inside are well within AI's demonstrated capabilities, the jagged frontier. One potential hazard for human teams is that they will be too slow to recognize AI's advantages within those functions, thereby putting themselves at a potential competitive disadvantage.

In one study, radiologists diagnosed X-rays either alone or supported by task-specific AI. AI alone was far more accurate than humans in this activity. Interestingly, adding AI support did not improve the accuracy of the joint human–AI teams since the radiologists tended to under-weigh AI's recommendations (National Bureau of Economic Research, 2024). Another study found that for tasks in which AI outperformed humans, adding a human to "team up" with the AI actually impaired performance.

By contrast, outside of the frontier, humans are generally better off without AI. Harvard Business School conducted a study evaluating the speed, quality, and accuracy of ideation and reasoning tasks, both inside and outside of the jagged frontier. For tasks outside the frontier, humans moved 18 to 30 percent faster with AI, but output correctness fell by 19 percent (see figure 9 on page 17).

AI shouldn't be seen as some kind of automatic linear upgrade to human decision-making and that a "human + AI" team won't automatically produce the best outcome in all circumstances.

Our respondents seemed to have some understanding of this already. One question asked in the Kearney–MIT study pertained to human–AI collaboration—specifically, which broad combination of human and AI control would function best within three years.

About two-thirds of respondents (65 percent) said the most optimal combination would be AI with some form of human oversight. Less than half that proportion (30 percent) said the best approach would be "human and AI working together." Only 4 percent said they thought the most effective model would be "AI alone," without any human guidance.

All of this has deep implications for mobility executives seeking to figure out how humans and AI systems can most effectively collaborate. Instead of viewing AI as a general-purpose replacement for human intelligence, the jagged frontier suggests that optimal performance emerges from a strategic division of labor that leverages each party's complementary strengths.

In other words, the most vital question is not simply whether to adopt AI, but how to strike the right balance in human–AI collaboration. Effective human–AI teams must identify the specific capabilities along the frontier where AI excels and design workflows that channel tasks appropriately.

We need to understand when to trust AI recommendations and when to rely on human capacities instead. Relatedly, we need to design AI with transparency and interpretability features that allow human partners to understand the basis for AI decisions and intervene when necessary.

We need to understand when to trust AI recommendations and when to rely on human capacities instead.

Figure 9

AI is not a linear upgrade to human decision-making, and a "human + AI" team won't automatically produce the best outcome

For the tasks inside the frontier, AI-aided humans can place too little weight on AI input

- In one study, radiologists diagnosed X-rays either alone or supported by task-specific AI. AI alone was far more accurate than humans; adding AI support did not improve accuracy of human + AI teams, as people underweighted AI's recommendations.
- In another study, when AI already outperformed humans, adding a human to "team up" reduced performance relative to AI alone.

AI was more accurate than **75%** of participants, but AI assistance did not improve the joint outcome.

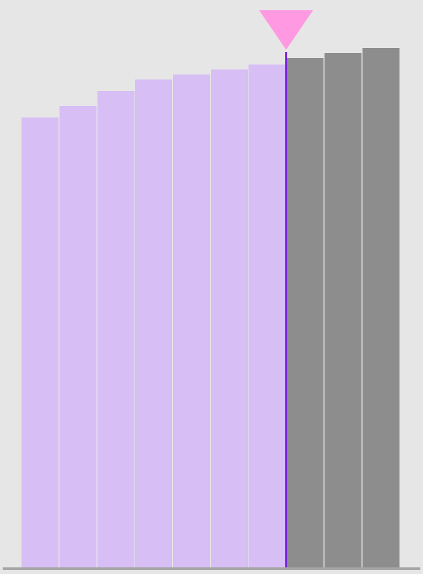

And for the tasks outside of the frontier, AI-aided humans are more likely to produce incorrect results.

- Harvard Business School performed a study evaluating speed, quality, and accuracy of the ideation and reasoning tasks inside and outside of the frontier.
- For the tasks outside the frontier, aiding human with AI reduced output correctness by ~19%, even though participants were 18 to 30% faster.

Outside of frontier output correctness

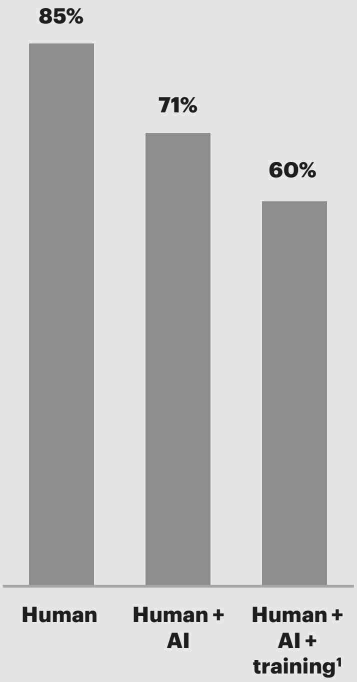

Yet, the industry remains reluctant to entrust decision-making to AI operating on its own.

- Kearney and MIT conducted a survey across the mobility ecosystem; one question was related to human–AI collaboration model.
- The majority sees human + AI setups as the most effective model for the future.

In 3y, applications of AI with better outcomes will be ...
(KAMI and MIT survey)

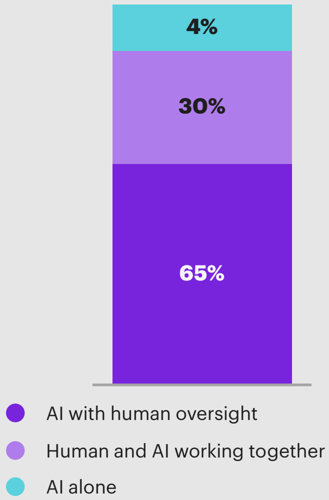

[1] Supplementary prompt engineering overview, which increased subject's familiarity with AI

Sources: "Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology," National Bureau of Economic Research; "When combinations of humans and AI are useful: A systematic review and meta-analysis," Nature Human Behaviour; Harvard Business School; Kearney and MIT analysis

Broadly speaking, there are three options for structuring the human–AI collaboration:

Human in the loop (HITL). The human has final responsibility for outcomes. AI performs tasks such as data processing, prediction, or recommendation, while the human interprets these findings and decides how to act on them. Mobility-sector examples of when this collaboration model might work best include AI maintenance diagnostics as well as vehicle design and prototyping.

Human on the loop (HOTL). In this collaborative model, AI executes within defined boundaries. Humans oversee outcomes and keep a lookout for exceptions and anomalies, overriding the system when necessary. Mobility-sector instances include adaptive traffic control and the supervision of autonomous fleets.

Human out of the loop (HOOTL). Here, AI systems execute decisions with little or no human oversight. In theory, AI's accuracy, safety, and ethics are ensured by design. Humans may stop the system or switch it to "safe mode" if anomalies are detected; in addition, systems undergo regular auditing. In mobility, this mode could apply to automated charging for electric vehicle fleets or train controls on closed-loop metro systems.

AI is constantly evolving, and the jagged frontier therefore continually shifts as AI systems improve and new training methodologies emerge. Capabilities that once represented peaks of AI performance become routine while new challenges arise.

The frontier's shape at any given moment reflects the current state of AI research priorities, available datasets, computational resources, and fundamental algorithmic approaches. As AI systems become more sophisticated, the remaining gaps may reveal uniquely human cognitive abilities that are difficult to replicate through extant AI technologies.

Understanding these limitations will help executives and policymakers develop a realistic expectation of AI's capabilities and ensure they design systems that enhance rather than replace human judgment in critical domains.

Unique challenges for AI in mobility

Despite the growing momentum and promise of AI in the mobility sector, its commercial deployment is challenged by a complex set of interrelated obstacles. These obstacles extend far beyond the frequently mentioned technical hurdles and reflect deeper structural, organizational, and institutional frictions. For example, safety is critical to mobility and any new tech will have to be safe in order to prevent cascading negative impacts.

Based on insights from our interviews for this study, five core challenge clusters emerge, each of which inhibits the full realization of AI's transformative potential in the mobility field.

Challenge 1: The data paradox of abundance and inaccessibility

AI thrives on data, but the mobility sector is frequently caught between abundance and inaccessibility. Although there is a plethora of data within each organization, at the systemic level, accessibility and quality remain crucial issues.

While some market participants, such as manufacturers, may guard vehicle or battery telemetry data closely, others face legal constraints in utilizing surveillance and transport-system data. Even where data exists, fragmentation across legacy IT systems, inconsistent formats, and slow digitalization efforts limit its usability. The result is a bottleneck that hinders AI model training, integration, and scaling.

Challenge 2: The commercialization gap arising from a lack of business models

While all interview respondents agreed that companies are increasingly investing in AI-enabled projects—particularly within their R&D and innovation efforts—the current share of AI-related projects still accounts for a low percentage of overall innovation and R&D portfolios across the mobility value chain. This relatively modest adoption rate is viewed as underwhelming, especially given the ongoing hype cycle surrounding AI.

Demonstrating a clear return on investment (ROI) is a major hurdle, complicated by lengthy business-to-government (B2G) sales cycles, slow user adoption due to trust concerns, limited scalability, and the difficulty of justifying massive long-term investments. For example, Argo AI collapsed in 2022, following over $3 Billion in investment from Ford & Volkswagen, because of the lack of a clear pathway toward profitable fully autonomous operations.

The intangible and difficult-to-quantify nature of AI-generated value often hinders broader implementation. Companies and agencies often struggle to integrate AI into viable business models, particularly when monetization requires changes across the entire enterprise or when success metrics remain undefined.

Without clearer paths to industrialization and commercialization, many promising use cases risk remaining merely academic or experimental. The buzz around AI risks creating disillusionment, highlighting a potential disconnect between strategic ambition and operational reality.

Challenge 3: Cultural and organizational resistance

Cultural and organizational lethargy is perhaps the most underestimated barrier to the effective utilization of AI. Successful AI deployment often requires rethinking workflows, retraining staff, and breaking down departmental silos. Yet across the mobility sector—among transit authorities, mobility service providers, equipment suppliers, and manufacturers alike—a reflexive resistance to change is all too common.

This resistance is driven by overgeneralized fears of automation, a lack of internal AI fluency, and the desire to protect legacy processes. Furthermore, in decentralized operations, aligning local teams around centralized AI strategies can be especially challenging.

Challenge 4: Governance, regulation, and ethical complexity

AI applications in the mobility sector face a fragmented and evolving regulatory landscape. Many interview respondents cited a lack of clear standards and policies for AI implementation. The deployment of AI models that rely on surveillance footage, real-time location data, and similar information is especially challenging due to European General Data Protection Regulation (GDPR) constraints, US federal data rules, and divergent international privacy regulations.

The legal ambiguity surrounding copyright, transparency, insurance, and accountability for autonomous systems exacerbates the problem. At the same time, ethical oversight mechanisms are often absent, leaving organizations to navigate questions of fairness, bias, and risk independently.

Challenge 5: Talent and infrastructure deficits

AI expertise remains scarce and expensive. Recruiting skilled talent, especially in machine-learning operations, computer vision, and edge computing is a formidable challenge since it involves competing with tech giants and startups alike. Public-sector entities and traditional mobility operators are unable to compete with tech firms on compensation, so they struggle to attract and retain mission-critical personnel. As a result of these talent deficits, many organizations lack the basic digital capabilities necessary to support modern AI systems at the requisite scale.

Having surveyed some primary features of the current AI landscape, we turn in the next section to a look at some of the foremost applications of AI in the mobility sector.

Applying AI to transform mobility

In the mobility sector, AI is being deployed for a growing range of functions. Broadly speaking, these applications so far have largely centered on perceiving data and predicting outcomes across vehicles, infrastructure, and transportation networks, with the goal of improving how people and goods move from place to place. But the ways organizations apply AI for these functions can vary depending on the operational scale and scope.

From task-specific tools to system-wide transformation

AI applications range from narrow task-specific implementations to broad systemic transformations, and a clear classification framework is essential as we map the landscape of AI capabilities and limitations in mobility.

By no means is the classification in this report all-encompassing or exclusive. There is some overlap between categories, but on a broad level, this classification works well for most transportation and mobility use cases of AI, and we shall use this framework for subsequent sections of this report.

Task-level AI applications

As mentioned, AI systems generally excel at discrete, well-defined tasks that can be clearly specified and measured. These applications typically involve pattern recognition, classification, prediction, or optimization within bounded domains.

The strength of task-level AI lies in its ability to achieve breakthrough performance within narrow domains through intensive training on large data sets. Once trained, they can ingest and parse large amounts of data of all types. These systems benefit from clear success metrics, abundant training data, and well-understood problem spaces where the rules and objectives remain relatively stable.

For mobility applications, examples include predictive maintenance tools that can accurately estimate remaining component lifespans, capacity planning algorithms that can predict usage levels on different transit routes, computer vision systems that can identify and classify objects on the road, and natural language processing models that can recognize sentiment and classify public transit riders' feedback.

Individual-level AI applications

Individual-level AI applications use systems that integrate multiple task-level capabilities to provide more comprehensive assistance, recommendations, or guidance for an individual human user. Route planners and virtual assistants are prominent examples of mobility use cases for this category.

Individual-level AI systems must navigate the complexity of human preferences, contexts, and behaviors—achieving hyper-personalization and tailored assistance for travelers while simultaneously maintaining privacy and abiding by user controls. The challenge at this level lies in understanding user intent, managing competing priorities, and adapting to changing preferences over time.

The effectiveness of these systems depends heavily on the quality of human–AI interaction design: is the combination more or less productive than human alone or AI alone? The system must be designed so that it balances automation with user agency and can accurately identify when it must let the human lead.

Enterprise-level AI applications

At the organizational level, AI systems can coordinate multiple processes, optimize resource allocation, and support complex decision-making across departments.

Here, AI can lend significant value through the integration of intelligence within mobility players—connecting network planning, dispatch, charging, and other basic functions. Enhanced learning leads to next-level network optimization, safety improvement, better utilization, and more efficient use of energy and other resources.

However, these systems face integration challenges, as they must interface with legacy infrastructure, fragmented data sets, and diverse stakeholder needs—all while maintaining operational continuity and introducing new capabilities. In fact, the complexities of organizational politics and regulatory compliance often do more to determine the success or failure of an AI application than the system's actual technical capabilities.

At this broader level of application, AI implementations require careful attention to governance, accountability, and transparency, including robust testing, monitoring, and fallback procedures.

System-level AI applications

At the highest, broadest level of scope, system-level AI applications could transform transportation infrastructure. Such systems promise to revamp metropolitan, regional or national mobility systems through improved coordination across modalities, enhanced capacity planning, and a more efficient use of power grids and other infrastructure.

Some leading examples of applications at this scale include autonomous integrated vehicle fleets and asset management systems, but far more ambitious projects are being planned for the years ahead.

System-level implementations require coordination across multiple organizations, jurisdictions, and regulatory frameworks. They often involve network effects, in which the value of the system increases substantially as adoption crosses a threshold but can be markedly less effective below that threshold. The transition to system-level AI deployment typically takes some time and requires sustained investment and policy development.

Successful implementations at this scale can generate enormous economic value and valuable insights, while failures can cascade across interconnected systems. This scale demands the most attention to safety, security, and alignment with human values.

This year, Dubai's Roads and Transport Authority (RTA) began work on a next-generation traffic-signal control system, incorporating AI, predictive analytics, and digital-twin technologies. The system will have the capability to adjust signal timings instantaneously in response to traffic conditions. It will employ a digital twin to simulate and test timing changes and enable priority control for public transport, emergency vehicles, and pedestrians. In its pilot program, it resulted in congestion reductions of between 16 and 37 percent.

System-level AI is the most complex to implement but has the potential to add the most value to society (see figure 10). At the other end, simpler task-level AI is easier to implement as it requires fewer integrations and will have fewer stakeholders involved but will be limited in the net value add to society. There will be minor variations within each category which we have represented via the cloud for each category.

A couple examples will illustrate how the value of AI in mobility evolves from the task level to the system level. In autonomous driving, AI can provide task-level assistance with computer-vision systems that can discern features such as pedestrians, roadway obstacles, and the lanes of a highway. But at the system level, AI can bring about entire autonomous integrated fleets, in which vehicles from different modalities are orchestrated within a single unified system.

Or consider AI's potential in capacity planning and scheduling. At the task level, algorithms can predict capacity utilization on specific routes. At the system level, AI can coordinate assets and test capabilities at municipal, regional, or national scale. The value gain here is as stark as the shift from more accurate scheduling on a single bus route to heightened resilience and efficiency across entire interconnected mobility networks.

We can illustrate this value evolution further through the example of the Washington Metropolitan Area Transit Authority (WMATA). The transit agency for the US national-capital region used MIT Transit Lab research to create a customer feedback and sentiment analysis tool to assess its Twitter and customer relationship management (CRM) system feedback in real time and classify its findings, along with generating corrective action to improve transit riders' experience (MIT, 2024).

At the task level, this work consisted largely of sentiment analysis: collecting passenger comments and complaints from social media feeds and internal CRM and classifying each post by such relevant categories as mode, route, and topic. At this level, the value created by the AI application resided in the heightened visibility and improved tracking of various customer-service issues.

Figure 10
Integrating AI tools into system-level applications impacts the whole mobility ecosystem

From task-specific tools to system-wide transformation

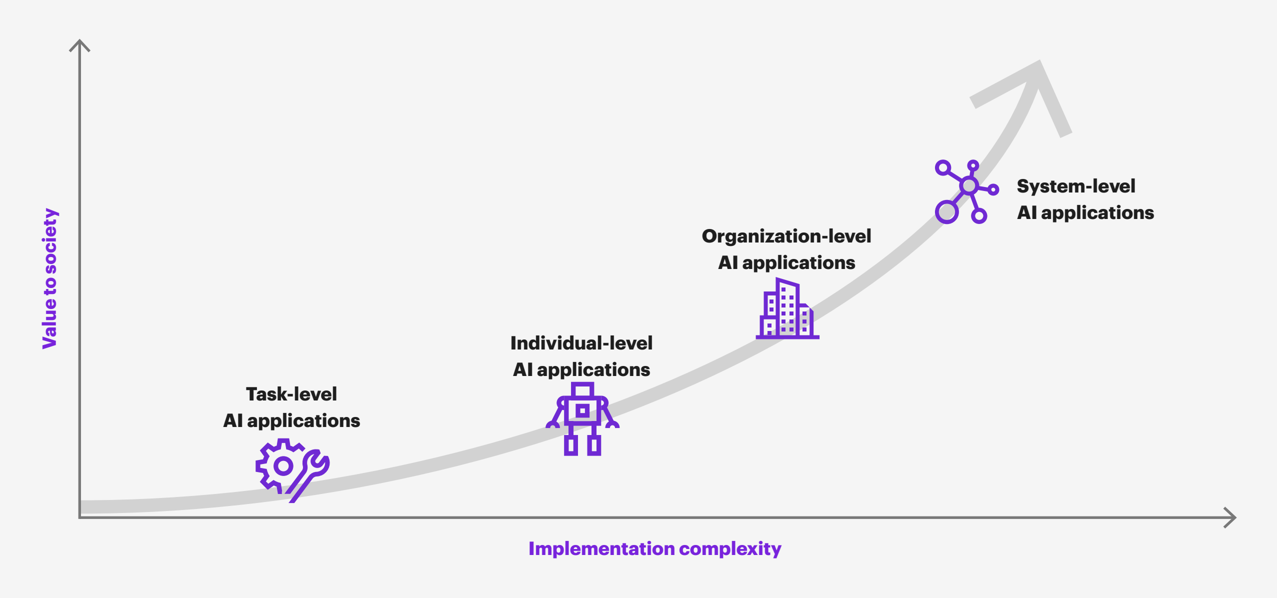

Sources: Harvard Business School; Kearney and MIT analysis

The next level up—the individual level—saw the use of AI for real-time mapping and tracking. Here, AI text analytics were combined with location data, creating live maps that could track the complaints' location and generate advanced performance analytics across the various lines and stations of WMATA's train and bus networks. The created value rises to actionable insights and decision support across the system.

From these levels, it is entirely possible to conceive of future states in which this work produces value at the organizational and even systemic level. At the organizational level, AI outputs could connect to WMATA's internal operations, such as those that govern maintenance and fleet planning, enabling automated responses such as the real-time dispatch of maintenance crews. The value gain here is clear, and potentially quite significant, as service reliability improves and operations become more efficient.

Finally, at the system level, the value creation could be greater still. One possibility would be for WMATA's AI network to be integrated with multiple metro-area mobility providers, such as Waymo, Uber, Capital Bikeshare, and Lyft. This interlinked AI could detect problems, demand surges, and coordinate actions across the entire system with potentially significant value gains in the form of system efficiency, customer experience, and accessibility.

AI applications across the three priority areas and beyond

Through extensive secondary research and over 50 expert interviews, we identified potential AI mobility applications across the three priority areas laid out in the introduction to this report: safety, inclusivity, and sustainability.

These domains represent the foundational challenges of the mobility sector that AI applications are designed to address. It's important to bear in mind, however, that the way AI addresses these domains may be a matter of indirect effect rather than primary intention. For example, RTA's AI-informed traffic system is primarily designed to reduce congestion; any resulting safety improvements are a consequence of that original intention rather than an explicit primary driver of them.

Next, we look at each of the three domains in turn.

Safety

This domain entails not only reducing traffic crashes, injuries, and fatalities, but also securing transport systems against operational failures and external threats. As urbanization intensifies and mobility systems become more interconnected, ensuring safety across different modes, infrastructure, and services becomes more complex.

Traditional safety mechanisms are often reactive, human-dependent, and insufficient for addressing the dynamic risks posed by high traffic volumes, mixed modes of transport, aging infrastructure, and unpredictable human behavior. AI can help by processing and absorbing vast amounts of real-time data and generating timely insights.

More specifically, AI can improve transportation safety by providing automated and autonomous vehicles with advanced perception and decision-making systems, offering predictive safety analytics for fleet operations, and enabling real-time hazard detection in infrastructure monitoring.

Several AI applications for the safety domain are already making a difference or will soon be on the way. Waymo operates robotaxi services in several US cities, managing routing, speed, and maneuvers in real time, using AI, lidar, radar, and cameras to perceive the environment and safely navigate without a human driver. The system continuously learns from billions of simulated and real-world miles, improving detection, prediction, and decision-making. So far, Waymo has logged 71 million rider-only miles without a human driver and has shown an 88 percent reduction in crashes compared with human drivers in similar conditions (Waymo, 2025).

Inclusivity

Applications in this area aim to ensure that technological advancements benefit all population segments, including individuals with limited mobility, underserved low-income communities, and regions with limited infrastructure.

AI-powered policy simulations, behavioral analysis tools, and journey planning systems help authorities and service providers create more inclusive mobility networks. Data-driven simulations, for instance, are being used to forecast the impacts of regulatory decisions and design policies that prioritize social equity alongside efficiency. Thanks to advances in AI capabilities, it is now much easier to design such interventions and gauge their impact.

That impact can be international in scale. AI tools have helped document 75 million miles of undocumented waterway across East Africa for a fraction of the time and cost of traditional methods (World Economic Forum, 2025). This helps provide essential information on ideal locations for bridges and other infrastructure in a traditionally underserved region of the world.

Sustainability

AI is being applied to support climate goals and reduce environmental impact. The mobility sector is under mounting pressure to reduce its carbon footprint, and AI can help by optimizing the placement of EV-charging infrastructure or managing traffic flow to reduce emissions.

For example, predictive algorithms that balance passenger supply and demand can reduce empty rides and optimize energy use during transport operations, contributing to cleaner mobility ecosystems. These applications support emissions reduction and facilitate the integration of renewable energy and electrified transport into broader infrastructure systems.

MIT researchers who simulated 6,011 signalized intersections across three major US metropolitan areas found that by transmitting intelligent speed commands to semiautonomous vehicles, AI can cut carbon emissions by between 11 and 22 percent (Jayawardana, Vindula, 2025). However, our respondents also stressed the need to balance the use of AI with the sustainability impacts of the technology itself—notably, its prodigious use of electricity, land, and water.

While these three priority areas were central to our study, we also deeply explored more business-centered applications of AI, inquiring into the cost-efficiency, productivity and profitability gains that AI makes possible across the mobility value chain.

Obviously, economic imperatives drive a large proportion of AI's mobility applications. Manufacturers and suppliers are leveraging the technology to accelerate vehicle R&D and optimize material use, while service operators are deploying AI for real-time fleet scheduling and predictive maintenance.

AI helps these solutions expand to meet evolving demand, enabling economies of scale and increasing revenue. These solutions reduce downtime, optimize resource allocation, and support smarter asset utilization, directly contributing to operational efficiency and bottom-line performance.

For instance, AI applications such as ride-sharing route optimization and intelligent demand forecasting exemplify how economic value can be driven by deep-learning algorithms. Uber Freight uses machine learning to create optimized truck routes, and the company has been able to reduce empty miles between 10 and 15 percent (Uber Freight, MIT CTL, 2025).

Deutsche Bahn is rolling out an AI-powered predictive-maintenance system for its rail infrastructure, using smart sensors and machine-learning models to monitor vibration, wear, and performance, predicting failures before they occur. This enables condition-based maintenance, reducing disruptions, lowering costs, and improving reliability.

Customer satisfaction emerges as a cross-cutting priority in both public and private mobility services. AI can enhance the passenger experience by enabling more personalized, responsive, and seamless services. AI-powered tools such as multimodal journey planners, dynamic pricing engines, customer-sentiment analysis, and personalized content generation tailor mobility offerings to individual needs and preferences across various modes of transportation and stakeholders in the value chain.

Assessing the impact, feasibility, and relevance of AI applications

Through our research, we identified 39 AI application areas in the mobility sector and categorized them into seven groups that span the entire value chain, offering a useful structure for understanding how AI is transforming the sector. Furthermore, the seven categories represent potentially divergent operational priorities that need to be aligned to unlock the full potential of AI and realize system-level gains (see figure 11 on page 26).

Vehicle design and manufacturing covers activities related to the development and production of vehicles and components. AI is applied in areas such as the discovery of new materials, virtual prototyping, energy management, advanced driver assistance systems, and autonomous driving to enhance safety, efficiency, and performance (see figure 12 on page 27).

Infrastructure design encompasses AI applications that support the planning, development, and resilience of mobility infrastructure. This includes location and network planning, demand simulation, and predictive modeling of disruptions using tools such as digital twins and GenAI (see figure 13 on page 27).

Seven categories across the mobility value chain reveal the operational priorities that can unlock the full potential of AI.

Traffic management and regulation includes AI applications in dynamic traffic control, parking optimization, tolling, and enforcement. These support safe, efficient, and equitable use of public road spaces, enhance rule compliance, and optimize urban mobility systems in real time (see figure 14 on page 28).

Operations refers to back-end mobility service processes such as fleet and route planning, capacity scheduling, energy optimization, personnel training, and service-reliability improvements such as bus-service deployment. AI contributes to these and other operational functions through capabilities such as predictive analytics, adaptive control, and simulation-based planning (see figure 15 on page 29).

Maintenance involves the upkeep of vehicles and infrastructure, with AI used for predictive diagnostics, inventory optimization, technician assistance, and anomaly detection through the Internet of Things (IoT) and computer-vision technologies (see figure 16 on page 30).

Safety and security focuses on applications that promote passenger and infrastructure safety, including crowd monitoring, fare-fraud prevention, and access control. AI enables real-time monitoring, anomaly detection, and rapid incident response (see figure 17 on page 30).

Customer engagement includes customer-facing functions such as multimodal journey planning, personalized offerings, dynamic pricing, and automated service responses. AI supports enhanced user experiences through tailored information delivery, accessibility solutions, and sentiment analysis (see figure 18 on page 31).

Figure 11

Based on discussions with actors across the mobility value chain, our study identified 39 AI application areas in seven categories

Application area overview

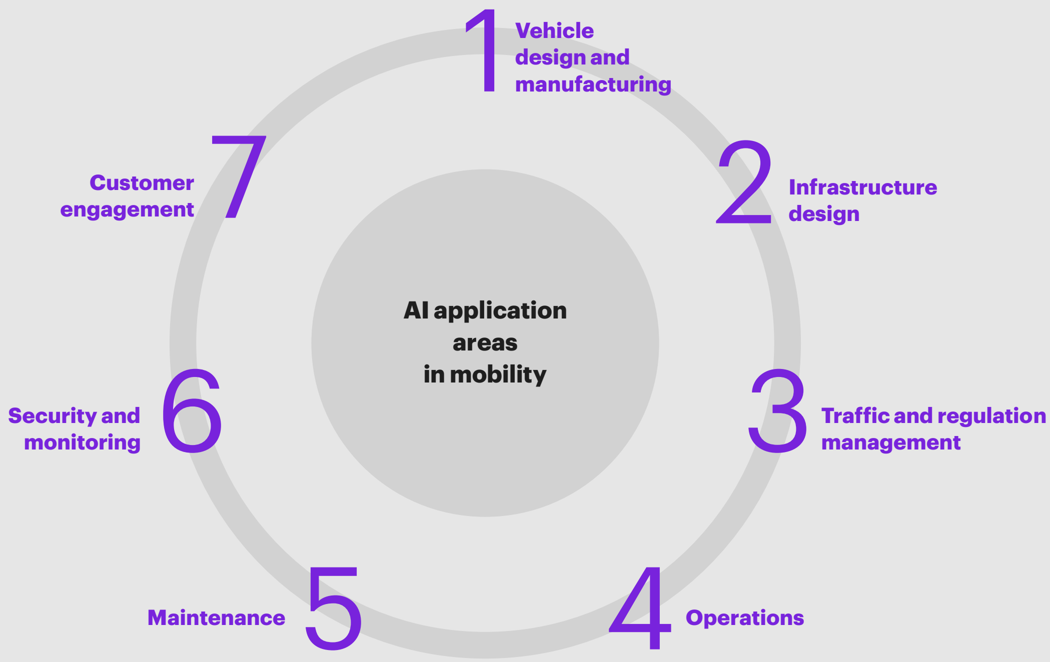

1 Vehicle design and manufacturing

1.1. Discovery of new materials
1.2. Design, testing, and prototyping
1.3. Energy management optimization
1.4. Advanced driver assistance systems
1.5. Autonomous driving

2 Infrastructure design

2.1. Location and network planning
2.2. Demand simulation
2.3. Disruption prediction

3 Traffic management and regulation

3.1. Adaptive traffic control
3.2. Real-time parking optimization
3.3. Advanced inner-city tolling
3.4. Toll enforcement
3.5. Parking enforcement
3.6. Automated traffic enforcement
3.7. Driver behavior monitoring
3.8. Smart policy design

4 Operations

4.1. Fleet and route planning
4.2. Capacity planning and scheduling
4.3. Route optimization
4.4. Daily fleet management
4.5. Personnel training and planning
4.6. Fuel/energy saving
4.7. Bus bunching reduction
4.8. Assistants for operations

5 Maintenance

5.1. Inventory and supply chain optimization
5.2. IoT-supported maintenance
5.3. Predictive maintenance
5.4. Maintenance training and assistants

6 Security and monitoring

6.1. Occupancy and crowd monitoring
6.2. Fare fraud prevention
6.3. Access control
6.4. Compliance and monitoring

7 Customer engagement

7.1. Seamless multi-modal integration
7.2. Enhanced travel information
7.3. Tailored customer offerings
7.4. Revenue management
7.5. Barrier-free and inclusive communication
7.6. Automated response management
7.7. Customer feedback analysis

Sources: Kearney and MIT analysis

Figure 12
Five AI application areas in vehicle design and manufacturing

Use case	Description	Use case examples
Discovery of new materials	Analyzing and discovering lighter, cheaper, and higher performance materials for batteries, chassis, vehicle components, etc.	**Toyota Research Institute** uses machine learning and automation for accelerated materials discovery for next-gen batteries.
Design, testing, and prototyping	Using AI-driven design and testing (e.g., simulation, generative design, virtual prototyping) to develop next-generation vehicles and parts	**Toyota Research Institute** pioneered an AI tool that converts text prompts into prototype sketches considering key engineering constraints.
Energy management optimization	Increasing battery life by charging cycle optimization, predicting battery health, optimizing charging routines, detecting anomalies, and identifying maintenance needs through data analysis	**Tesla** developed predictive battery management systems that adjust charging and discharging patterns based on vehicle usage to optimize battery performance.
Advanced driver assistance systems	Enhancing safety through accident prevention systems and advanced driver assistance, ADAS functionalities (Society of Automotive Engineers L1/L2)	**Singapore's Land Transport Authority** imposed AI-based driver assistance to buses including fatigue monitoring, blind-spot detection, and bus-lane enforcement.
Autonomous driving	Developing, testing, deploying autonomous vehicles and systems (cars, shuttles, buses, trains, trams, metro ...) (Society of Automotive Engineers[1] L3/L4/L5)	Robotaxi services without a safety driver on board: **Waymo** 2k vehicles, **Baidu/Apollo** 1.5k vehicles, and **Pony.ai** 1k vehicles

Sources: Materials Discovery, Toyota Research Institute; Toyota Research Institute Unveils New Generative AI Technique for Vehicle Design, Toyota USA Newsroom; Artificial Intelligence In Battery Health Monitoring, Advanced driver assistance system for buses, Land Transport Guru; "How robotaxis are trying to win passengers' trust," BBC, 18 November 2024; Kearney and MIT analysis

Figure 13
Three AI application areas in infrastructure design

Use case	Description	Use case examples
Location and network planning	Optimizing site selection and network layout for streets, stations, depots, EV charging hubs, and micro-mobility docking stations. Location choices might be optimized based on accessibility, forecasted demand, demographic, and traffic flow data. Can include retrospective analysis of the performance of the road network (speeds, journey times, routes, etc.).	**Sand Technologies' AI Network Planner** combines AI, machine learning, and digital twins to identify EV charging station locations that balance current and future demand, providing detailed insights into customer demographics and optimal network expansion configurations.
Demand simulation	Simulating passenger demand and movement patterns across transport modes and geographies. Includes use of anonymized mobile phone data, historical ridership, and land use data to simulate future demand in various infrastructure investment scenarios.	**Google Research** created **Mobility AI**, a program leveraging AI advancements in measurement, simulation, and optimization to provide transportation agencies with tools for data-driven policymaking, traffic management, and continuous monitoring of urban transportation systems.
Disruption prediction	Simulating disruptions due to extreme weather, accidents, or unusual events. Includes generative AI for simulating edge cases (e.g., obstacle detection, animal crossings) and testing infrastructure resilience. Supports events and emergency response planning and robust infrastructure design. Includes use of anonymized mobile phone data, historical ridership, and land use data to simulate future demand in various infrastructure investment scenarios.	**Swiss Federal Railways (SBB)** uses AI-driven data analytics tool to automatically identify train delays and their causes. This solution replaces manual disruption analysis, saving dispatchers up to 15% of their time and improving rail operations with data-driven insights from multilingual sources and big data.

Sources: "Powering the EV Future of Urban Mobility," Sand Technologies, September 2025; "Introducing Mobility AI: Advancing urban transportation," Google Research, 23 April 2025; "New data analytics software for SBB," Detecon; Kearney and MIT analysis

Figure 14
Eight AI application areas in traffic management and regulation

Use case	Description	Use case examples
Adaptive traffic control	Adjusting traffic signals dynamically for improved traffic flow, passenger and vehicle safety. and reduced congestion, including in response to incidents	**Transmove project in Hamburg** uses AI to forecast traffic for police traffic control and the public transport operations center.
Real-time parking optimization	Identifying free parking spaces, through sensor and camera data, and sharing information and navigation instructions with drivers	**Parkopedia's** machine learning predicts on-street and off-street parking availability; leveraging AI and big data.
Advanced inner-city tolling	Adjusting tolls and access based on real-time congestion, time of day, or vehicle-specific attributes, including enforcement of "mobility caps" (e.g., crossing limits per license plate)	**Singapore's** next-gen ERP 2 introduces satellite-based, gantry-less road pricing using GNSS.
Toll enforcement	Enforcing toll payments via real-time license plate recognition and cross-check with payment records. Ensures compliance with urban tolling, congestion zones, etc.	**Sighthound ALPR+** uses cameras and deep learning to recognize license plates, and real-time violation detection.
Parking enforcement	Monitoring and enforcing parking rules, including curb management, overstays, illegal parking, or zone-specific parking violations	**Medford, Massachusetts** uses AI-powered license plate recognition parking enforcement with ticket-by-mail, resulting in a 95% compliance rate.
Automated traffic enforcement	Detecting and penalizing traffic violations such as illegal bus lane use, red light running, and illegal turns using real-time cameras and AI vision systems	**MTA's ACE program** is a bus-mounted camera system issuing violations to vehicles occupying bus lanes and blocking bus stops.
Driver behavior monitoring	Monitoring driver behavior (e.g., harsh braking, speeding, fatigue) using telematics and sensor data; may include crash data assessment (weather, etc.)	**Alsa** buses integrate emergency braking, fatigue detection, cameras, speed monitor and driving, pedestrian detection systems.
Smart policy design	Developing data-driven mobility policies by simulating e.g., the socio-economic and environmental impact of measures to assess trade-offs and design evidence-based interventions	**Aimsun Live** digital twins suppor agencies in the UK in testing solutions and ranking response plans before deployment.

Sources: "Transition time for mobility in Hamburg," ITS International; "Parking Availability: How predictive and responsive data work together to provide drivers with the best possible parking experience," Parkopedia; Electronic Road Pricing, Ministry of Transport; Sighthound ALPR+, Advanced License Plate Recognition & Vehicle Analytics; The Canadian Parking Association: The AI Evolution in Parking: From Assistance to Autonomy; Medford's Citizen-Centered Parking Management Success; MTA Automated Camera Enforcement; Passenger transport safety: the vision that drives Alsa, IRU, World Road Transport Organisation; Aimsun & Yunex deliver digital twin for Tees Valley, ITS International; Kearney and MIT analysis

Figure 15
Eight AI application areas in operations

Use case	Description	Use case examples
Fleet and route planning	Forecasting long-term demand and scenario-based route planning to support strategic decisions on fleet size, type, and vehicle deployment	**Stagecoach** uses Optibus's AI platform to deliver smarter timetables and networks and keep up with travel demand.
Capacity planning and scheduling	Matching capacity (vehicles and staff) with expected demand by predicting peak, off-peak loads and aligning, e.g., shift planning, service frequency; includes events management	**West Coast Motors** launched AI scheduling based on driver preferences, relief vehicles, and break rules.
Route optimization	Computing and adjusting optimal pick-up, drop-off, and shared routes in real time for pooling and on-demand services, based on traffic, cancellations, and new requests	**UPS ORION** is an AI-powered route optimization software to increase efficiency, and reduce delay and fuel wastage.
Daily fleet management	Optimizing daily operations such as charging schedules, cleaning, inspections, and minor servicing, leveraging AI to minimize vehicle downtime and maximize service readiness	**Uber's** AI models minimize idle time and unproductive waits.
Personnel training and planning	Training crew and planning their shifts based on qualifications, performance, and demand forecasts. Includes adaptive learning paths and simulation-based training	**Alsa** introduced AI-derived behavior clusters; for targeted driver coaching and incentives, training sessions route/time-specific issues.
Fuel/energy saving	Using driver-assistance systems and route optimization to reduce energy and fuel consumption by improving driving patterns and vehicle dispatch	**Alsa** introduced a driver-efficiency program correlating telemetry with context to coach drivers and reduce fuel use (4–12% savings).
Bus bunching reduction	Dynamically recommending drivers and dispatchers to maintain headways and avoid vehicle clustering, including holding strategies and depot-based route interventions	**First Bus's** AI solution enables altering timetables and automatic adjustments by road congestion, resulting in 20%+ punctuality.
Assistants for operations	Provide seamless operation support, trouble-shooting, instructions, procedure walkthroughs, and other technical and logistical issues through "ops–copilots" and chatbots	**AC Transit** introduced conversational assistant to submit IT requests, get troubleshooting, automate password resets and basic HR.

Sources: "Stagecoach invests in new technology to help plan the bus networks of the future," Stagecoach Group News for Media and Press; "Using AI to create better bus driver schedules, faster," Optibus; "UPS & the Power of Route Optimization, Jones Elite Logistics; "Forecasting Models to Improve Driver Availability at Airports," Uber Blog; Alsa, "How AI is helping to prevent three buses turning up at once," BBC; "AC Transit Case Study: Transformative Impact of GenAI," Rezolve; Kearney and MIT analysis

Figure 16
Four AI application areas in maintenance

Use case	Description	Use case examples
Inventory and supply chain optimization	Forecasting spare part demand and optimizing inventory levels to ensure timely availability of critical components while avoiding overstock and shortages	**BMW** is using GenAI tools, including **Knowledge Navigator, Offer Analyst, and Tender Assistant,** to assist in the analysis and comparison of supplier offers as well as tender creation and management.
IoT-supported maintenance	Detecting issues such as graffiti, track obstacles, or infrastructure faults by using cameras and sensors to enable timely interventions and remote monitoring	**Metropolitan Transportation Authority** launched **TrackInspect,** using Google Pixel smartphones, retrofitted onto subway cars, to capture vibrations and sound patterns of tracks through built-in sensors and microphones. AI and machine learning analyze the data for potential track defects.
Predictive maintenance	Optimizing maintenance scheduling using, e.g., on-vehicle sensor data and digital twins	**Railigent X** brings together asset operations and condition data to detect faults, streamline maintenance, and support staff decision-making, to optimize maintenance and service reliability.
Maintenance training and assistants	Leveraging co-pilots to support technicians with step-by-step repair guidance, interactive troubleshooting, and access to manuals via voice, AR, or chatbot interfaces	**Penske** use AI-powered systems to guide fleet technicians with predictive analytics and fault diagnostics. AI provides step-by-step repair instructions and facilitates proactive, data-driven maintenance, boosting operational efficiency.

Sources: "BMW is using AIconic AI tools for its Purchasing and Supplier Network, Automotive Logistics; "AI is telling New York subway workers if that suspicious sound is a problem," StateScoop; "Smarter Maintenance: How Fleets Are Leveraging AI, TT; Railigent X, rail asset management, Siemens Mobility Global; Kearney and MIT analysis

Figure 17
Four AI application areas in security and monitoring

Use case	Description	Use case examples
Occupancy and crowd monitoring	Monitoring occupancy and passenger density in buses, trains, other vehicles, and stations or assets to ensure safety and improve operational decisions; includes real-time alerting to potential incidents	**Sofia Urban Mobility Center and Theoremus** developed onboard camera images classified into five occupancy levels in real time, replacing error-prone boarding/alighting counts.
Fare fraud prevention	Identifying fare evasion patterns using, e.g., biometric and behavioral profiling to prioritize passengers for inspection and ticket controls	**FGC Barcelona** partnered with **AWAAIT** to deploy an AI-powered fare evasion detection system called **Detector**, using real-time video analytics to identify fare evaders by analyzing camera footage around ticket barriers.
Access control	Preventing unauthorized persons from entering restricted areas (e.g., train tracks) via AI-based monitoring and alert systems	**Dubai Roads and Transport Authority** showcased smart gates allowing fare payment across metro/tram/bus/taxi via facial recognition.
Compliance and monitoring	Monitoring compliance with safety protocols (e.g., presence of fire extinguishers, use of protective equipment), detecting driver impairment (e.g., fatigue, DUI), identifying passenger misconduct (e.g., harassment, substance abuse, altercations), and recognizing lost or unattended items (e.g., luggage)	**Reply** developed **Driver State Monitoring** solutions, which can detect instances of impaired driving (drowsiness, distraction, driving under the influence, sudden illnesses, and aggressive driving).

Sources: "Customer technologies," Transport for NSW; Sofia Urban Mobility Center with Theoremus; "AWAAIT deploys its fare dodging detection system DETECTOR at Provença and Plaça Catalunya FGC's stations in Barcelona," Awaait Artificial Intelligence; "RTA showcases technology of paying public transport fares through facial recognition at GITEX Global," Government of Dubai; "AI-enhanced Driver State Monitoring goes beyond sensors," Reply; Kearney and MIT analysis

Figure 18
Seven AI application areas in customer engagement

Use case	Description	Use case examples
Seamless multimodal integration	Routing and planning journeys across multiple modes (bus, train, ride-hail, bike, walk) with real-time updates and seamless transitions as well as on the spot solutions when delays occur (re-routings, automatic adaptations, etc.)	**Transport for London's** TravelBot on Facebook Messenger provides journey planning and service updates via chat.
Enhanced travel information	Providing passengers with real-time information on location, crowding, delays, and next connections across modes	**PostBus** introduced an automated, multilingual acoustic passenger information (TTS) for planned / unplanned incidents.
Tailored customer offerings	Analyzing customer's mobility behavior to offer customized recommendations (e.g., ticket bundles) or identify changing patterns	**Citymapper's** app uses ML to personalize route options and alerts based on rider behavior and preferences.
Revenue management	Adjusting fares based on demand, time, or user behavior to increase revenue	**Uber's** marketplace uses ML-based dynamic pricing (surge) to balance rider demand and driver supply in real time.
Barrier-free and inclusive communication	Translating tools, audio assistance, and visual cues, e.g., for passengers with disabilities or language barriers	**SBS Transit and FingerDance** deploy SiLViA, a GenAI application that instantly translates spoken or written words into sign language.
Automated response management	Automating responses to service complaints, fare queries, and bookings via AI enabled voice and text interfaces	**Dubai Roads and Transport Authority's** multilingual rider chatbots respond to 20 to 30% of inquiries, automating common questions and complaint intake.
Customer feedback analysis	Analyzing customer reviews (sentiment, topic) and social media inputs to derive actionable service improvements, policy insights, adaptions of product offerings, etc.	**Transport for New South Wales** applies advanced analytics using Microsoft Azure to harness multimodal data streams to improve customer service.

Sources: "TfL launches new social media 'TravelBot,'" Transport for London; "Artificial Intelligence in Public Transport," UITP; "From Predictive to Generative: How Michelangelo Accelerates Uber's AI Journey," Uber Blog; "Annex A: SBS Transit Launches Mobility Innovation Centre for Smarter and Greener Journeys and Better Customer Experience," SBS Transit; Transport for NSW: MPR Program Customer Success Story," Agile Analytics; Kearney and MIT analysis

Figure 19 on page 33 shows where our respondents from across the mobility sector are currently focusing their AI efforts, based on our extensive interviews. As indicated, the level of interest across and within these seven categories is somewhat uneven, with some applications attracting more interest than others. Areas of particular emphasis include access control, fleet and route planning, predictive maintenance, and automated response management.

To drill down further on the potential importance of AI to the mobility sector and to more rigorously assess the strategic relevance of AI applications, we deployed a triangulated methodology, integrating insights from expert interviews, written survey responses, and secondary research. (See the introduction for more details on that methodology.)

This multisource approach enabled a robust evaluation of both feasibility and impact across the 39 identified AI application areas shown in figure 19 on page 33. In the figures that follow, each application is mapped along a two-dimensional framework, with the horizontal axis measuring feasibility and the vertical axis representing impact.

It's worth unpacking what we mean by these two dimensions because their interaction defines much of our analysis.

— **Feasibility** is how technically and organizationally implementable a use case is under current or near-term conditions. It reflects factors such as data availability, algorithmic maturity, integration complexity, and institutional readiness.

— **Impact** refers to the transformative potential of each AI application across dimensions such as operational efficiency, cost reduction, safety improvements, sustainability outcomes, and user-experience enhancement.

The relative positioning of a use case within this coordinate system offers an indicative snapshot of its strategic viability. The size of each dot in figure 20 on page 34 denotes the overall relevance of the use case in practice, measured by the number of times it was mentioned across all expert interviews.

These mentions serve as a proxy for industry attention and intent, signaling whether companies have already implemented, are planning to deploy, or are seriously exploring the given application. Therefore, dot size represents the approximate level of current prioritization among our expert interviewees.

This juxtaposition of dot location and dot size enables us to distinguish between the ubiquity with which an application is discussed, and its actual relevance in the mobility marketplace. While every effort was taken to ensure we interviewed a wide range and number of companies across the globe and across domains (to ensure all use cases have sufficient representation), it is possible that the mix of companies and number of interviews may lead to relevance being slightly skewed.

For example, a small dot that is positioned high on both the horizontal (feasibility) and vertical (impact) axes would represent an application that is high in potential value but may be flying under the radar. By contrast, a large dot that is positioned relatively low on either axis may indicate an application that may not be ready to deliver much actual value—at least, not yet.

As figure 20 indicates, the most commonly discussed mobility applications for AI are unevenly situated in terms of their overall impact, feasibility, and perceived relevance.

Figure 19

AI interest is uneven across categories with more focus on security and monitoring, O&M, and customer engagement

Relevance of use cases[1]

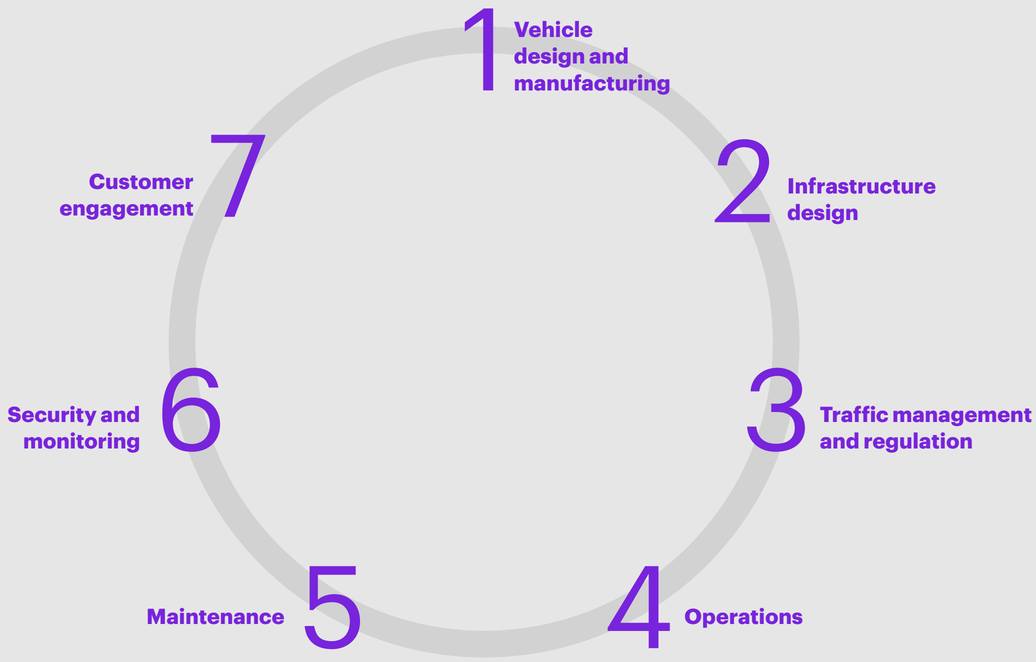

1. Vehicle design and manufacturing

1.1 Discovery of new materials	1
1.2 Design, testing and prototyping	4
1.3 Energy management optimization	2
1.4 Advanced driver assistance systems	4
1.5 Autonomous driving	13

2. Infrastructure design

2.1 Location and network planning	16
2.2 Demand simulation	5
2.3 Disruption prediction	10

3. Traffic management

3.1 Adaptive traffic control	18
3.2 Real-time parking optimization	2
3.3 Advanced inner-city tolling	1
3.4 Toll enforcement	4
3.5 Parking enforcement	2
3.6 Automated traffic enforcement	4
3.7 Driver behavior monitoring	9
3.8 Smart policy design	2

4. Operations

4.1 Fleet and route planning	17
4.2 Capacity planning and scheduling	13
4.3 Route optimization	8
4.4 Daily fleet management	7
4.5 Personnel training and planning	8
4.6 Fuel/energy saving	1
4.7 Bus bunching reduction	1
4.8 Assistants for operations	1

5. Maintenance

5.1 Inventory and supply chain optimization	1
5.2 IoT-supported maintenance	6
5.3 Predictive maintenance	20
5.4 Maintenance training and assistants	4

6. Security and monitoring

6.1 Occupancy and crowd monitoring	11
6.2 Fare fraud prevention	7
6.3 Access control	20
6.4 Compliance and monitoring	4

7. Customer engagement

7.1 Seamless multimodal integration	6
7.2 Enhanced travel information	5
7.3 Tailored customer offerings	5
7.4 Revenue management	12
7.5 Barrier-free and inclusive communication	3
7.6 Automated response management	20
7.7 Customer feedback analysis	5

[1] Using the number of interview mentions as a proxy

Sources: Kearney and MIT analysis

Figure 20

Use case prioritization across the mobility sector does not clearly follow either expected impact or feasibility

Impact, feasibility, and relevance assessment by application area

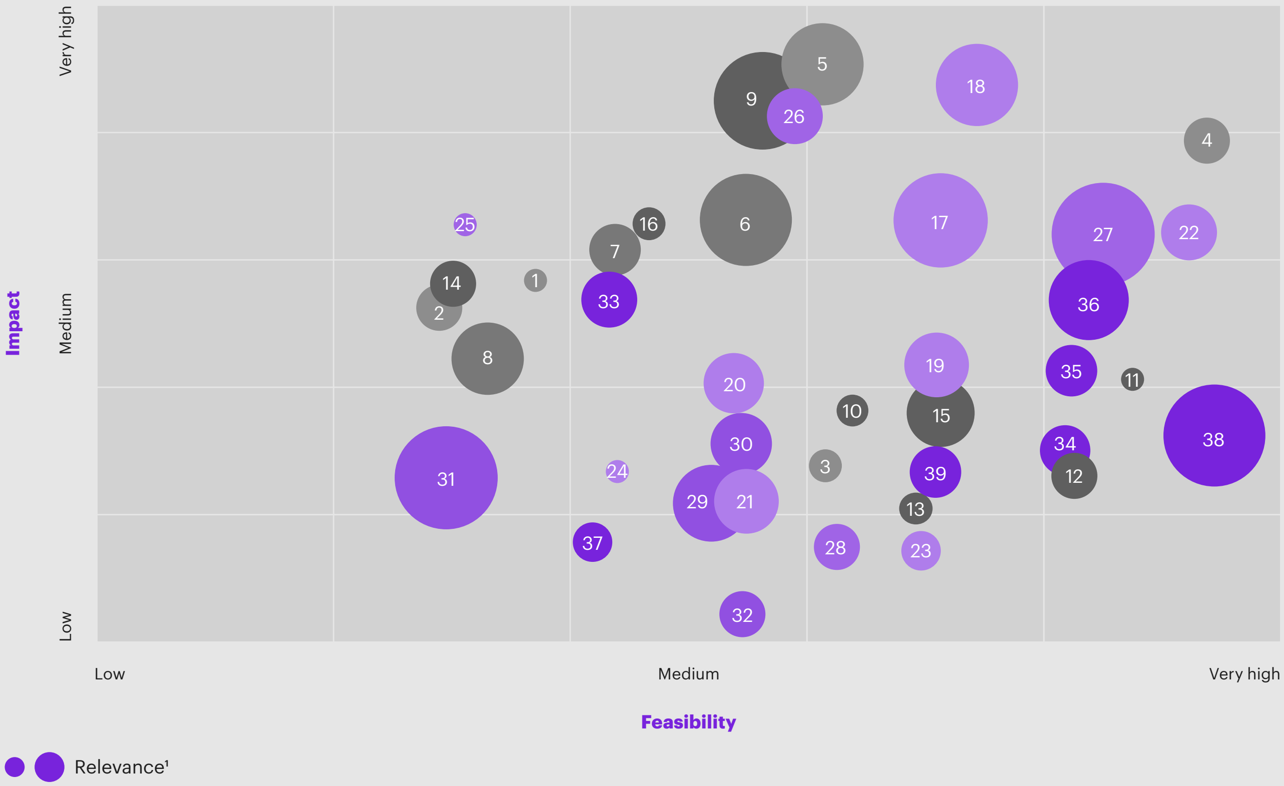

Vehicle design and manufacturing

1 Discovery of new materials
2 Design, testing, and prototyping
3 Energy management optimization
4 Advanced driver assistance systems
5 Autonomous driving

Infrastructure design

6 Location and network planning
7 Demand simulation
8 Disruption prediction

Traffic management and regulation

9 Adaptive traffic control
10 Real-time parking optimization
11 Advanced inner-city tolling
12 Toll enforcement
13 Parking enforcement
14 Automated traffic enforcement
15 Driver behavior monitoring
16 Smart policy design

Operations

17 Fleet and route planning
18 Capacity planning and scheduling
19 Route optimization
20 Daily fleet management
21 Personnel training and planning
22 Fuel/energy saving
23 Bus bunching reduction
24 Assistants for operations

Maintenance

25 Inventory and supply chain optimization
26 IoT-supported maintenance
27 Predictive maintenance
28 Maintenance training and assistants

Security and monitoring

29 Occupancy and crowd monitoring
30 Fare fraud prevention
31 Access control
32 Compliance and monitoring

Customer engagement

33 Seamless multimodal integration
34 Enhanced travel information
35 Tailored customer offerings
36 Revenue management
37 Barrier-free and inclusive communication
38 Automated response management
39 Customer feedback analysis

[1] Using the number of interview mentions as a proxy

Sources: Kearney and MIT analysis

This implies that the adoption of AI in mobility remains highly fragmented; impact and feasibility vary widely; and much of the value is still confined to isolated pilots rather than scaled, integrated systems—and the market is pursuing numerous task-specific AI applications without a focus on either value creation or feasibility.

High-impact, high-feasibility use cases, such as advanced driver assistance system, fuel and energy savings, or predictive maintenance, are ready for scaling up.

Others, including IoT-supported maintenance or adaptive traffic control, show real promise but require further investment, regulatory clarity, or organizational alignment. Some, particularly those focused on end-user engagement, remain experimental but could be highly beneficial in the long term if properly integrated into broader digital strategies.

Put simply, the sector is experimenting in many directions at once, often without a clear vision of how to progress toward system-level AI or which applications will ultimately achieve the desired returns on investment.

Another insight is that use-case prioritization across the mobility sector does not clearly follow either expected impact or feasibility. For example, access control—one of the largest dots on figure 20, befitting its status as one of the most mentioned applications in our study—scored notably low on both impact and feasibility.

Figure 21 below and figure 22 on page 36 offer insights across specific usage groups.

Vehicle design and manufacturing use cases show a high degree of variance in terms of both impact and feasibility. Autonomous driving and advanced driver assistance systems score high on impact and feasibility, while applications in design, testing, and prototyping score relatively low on both.

Infrastructure design is a somewhat different story. Here, use cases cluster along a diagonal line of middling feasibility and impact. Yet while none of these applications yet indicate transformative potential, their generally high level of current market relevance (as indicated by dot size) suggests at least the potential to win support across the mobility sector. So far, applications in location and network planning are particularly viable candidates within this category.

Figure 21
Autonomous driving and ADAS lead as high-impact, high-feasibility applications, while R&D-focused use cases lag

Vehicle design and manufacturing

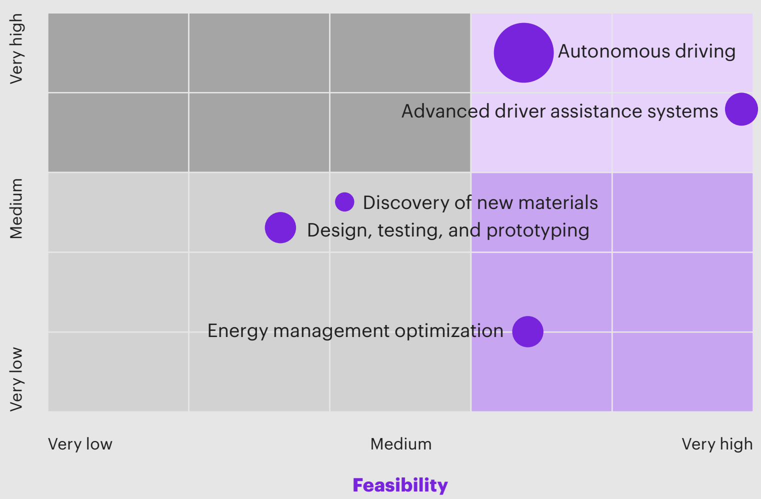

- **Autonomous driving** as a future game changer, with high impact and mid–high feasibility; true relevance likely understated given the respondent mix (with high share of operators/authorities/regulators vs. OEMs/suppliers)
- **Near-term focus is ADAS**, positioned in the high-impact/very-high- feasibility zone thanks to mature L1/L2 stacks and clear safety ROI
- **R&D-heavy items are longer horizon;** respondent mix may depress relevance

Peripheral ideas · Long shots · Quick wins · Game changers

Source: Kearney and MIT analysis

Figure 22

Location and network planning emerge as leading opportunities among moderately positioned infrastructure use cases

Infrastructure design

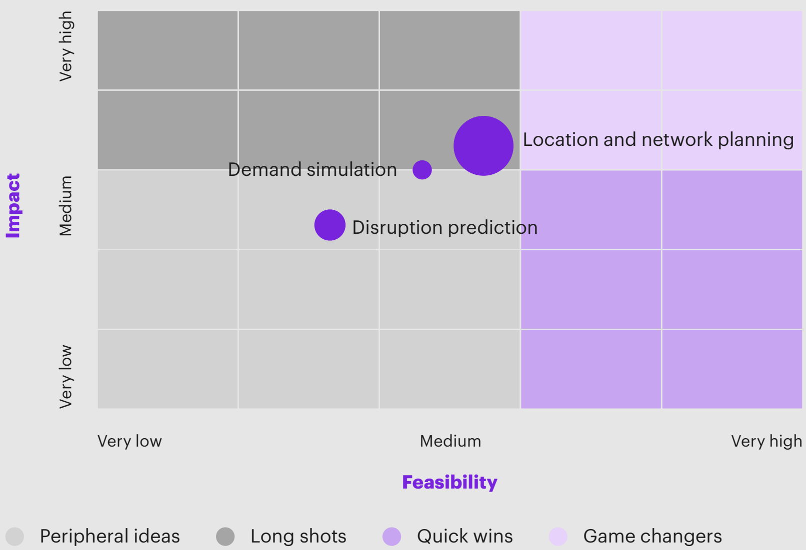

- **Use cases cluster along a diagonal** around mid-high feasibility and impact
- **No game changer cases;** however, generally high relevance indicates potential to win broad support across the mobility ecosystem
- Location and network planning as a **clear frontrunner**

Source: Kearney and MIT analysis

In the field of traffic management and regulation, average relevance per use case is surprisingly low, given that the survey included a significant number of regulators and operators that might be expected to value these functions more highly (see figure 23 on page 37). Applications tend to fall into two clusters: low-impact but readily attainable **quick wins** in enforcement areas such as tolling and parking and the somewhat more speculative—but potentially more impactful—**long shots** in adaptive traffic control and policy design.

Use cases in the operations category vary substantially in terms of impact but are generally seen as highly feasible (see figure 24 on page 37). Applications such as energy savings, capacity planning and scheduling, and fleet and route planning all scored high for both feasibility and impact.

Here, predictive maintenance looms large in relevance and feasibility, though its assessed impact is lower than expected; respondents may see limited value gains relative to current diagnostic systems. By contrast, IoT-supported maintenance has potentially high impact; its mid-tier scores for feasibility and relevance may reflect perceived difficulties with regard to extracting value from brownfield environments with legacy assets as well as the inherent challenge of dealing with mixed fleets, dispersed vendors, uneven connectivity, diffuse data sets, and other complexities (see figure 25 on page 38).

Security-related use cases are generally seen as having relatively low impact, mainly because they are not seen as serious drivers of either revenue or efficiency gains. Here, we get another look at the curious market position of access-control apps. While many of our respondents identified them as highly relevant to their operations, they actually score low on both impact and feasibility, possibly because of the high costs of hardware for security gates and other access-control infrastructure (see figure 26 on page 38).

Figure 23
Traffic management and regulation use cases cluster between attainable quick wins and speculative long shots

Traffic management and regulation

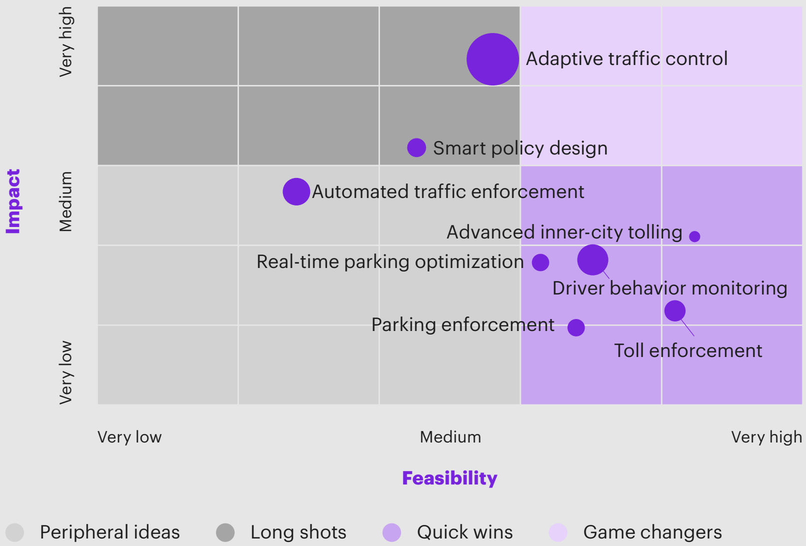

- **Two areas of use cases:** multiple low-impact quick wins in enforcement policy/behavior tools (tolling, parking, etc.) and long shots in adaptive traffic control and smart policy
- **Average relevance per use case** is rather low, even though the survey covered a significant number of regulators and operators. **Adaptive traffic control** stands out with very high impact and relevance for congestion, safety, and emissions.

Source: Kearney and MIT analysis

Figure 24
Highly feasible, high-impact operations use cases define near-term mobility performance gains

Operations

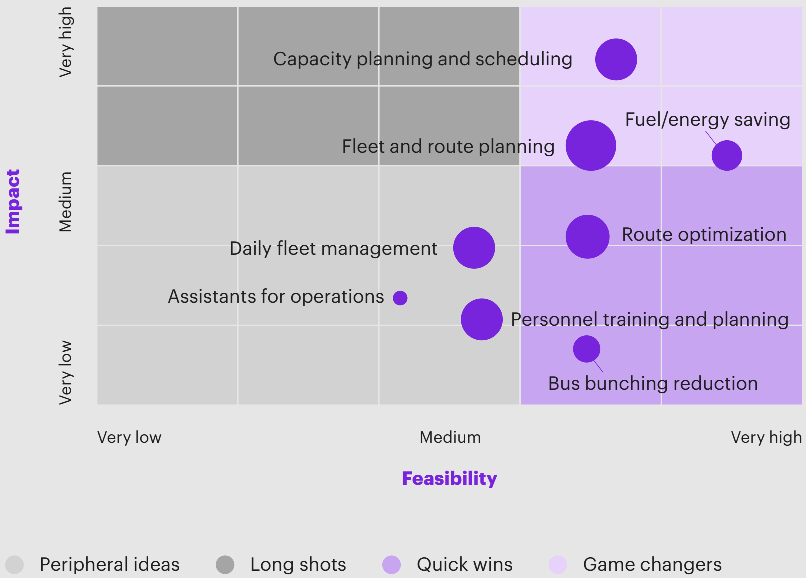

- Very high relevance and on average high feasibility for **operations use cases among the participants;** cluster is skewed to the right, so ecosystem will likely **prioritize by impact than feasibility**
- Impact **linked with control of core levers;** the top bubbles (capacity planning and scheduling, fleet/route planning, fuel/energy saving) reflect operator focus on KPIs they own daily

Source: Kearney and MIT analysis

Figure 25

Maintenance use cases are relevant and achievable but constrained by ecosystem complexity

Maintenance

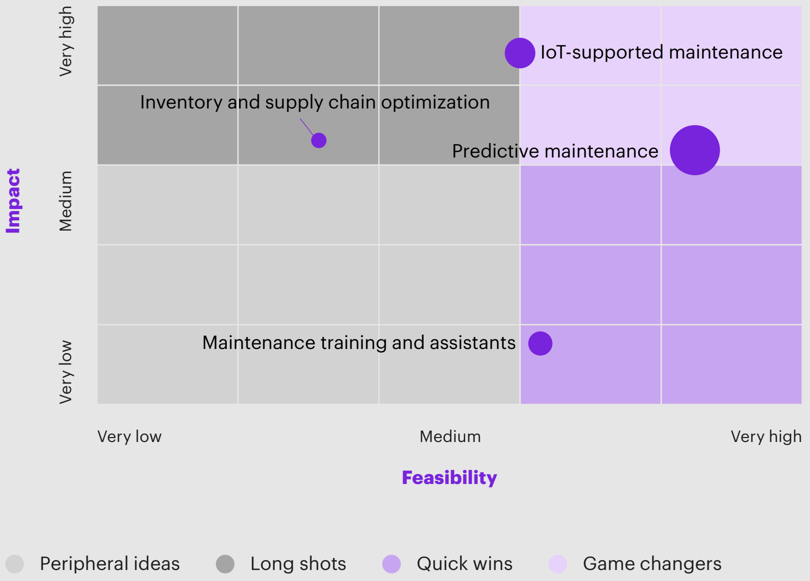

- **Predictive maintenance** carries the largest relevance and feasibility; the impact is lower than expected. Respondents may see limited value uplift over existing diagnostics/condition-based maintenance.
- **IoT-supported maintenance has potential for high impact; medium feasibility and relevance** may reflect the challenges extraction value in brownfield environments with legacy assets, mixed fleets/vendors, uneven connectivity/data harmonization, etc.

Source: Kearney and MIT analysis

Figure 26

Security-related AI applications show strong relevance but limited business impact due to hardware and cost barriers

Security and monitoring

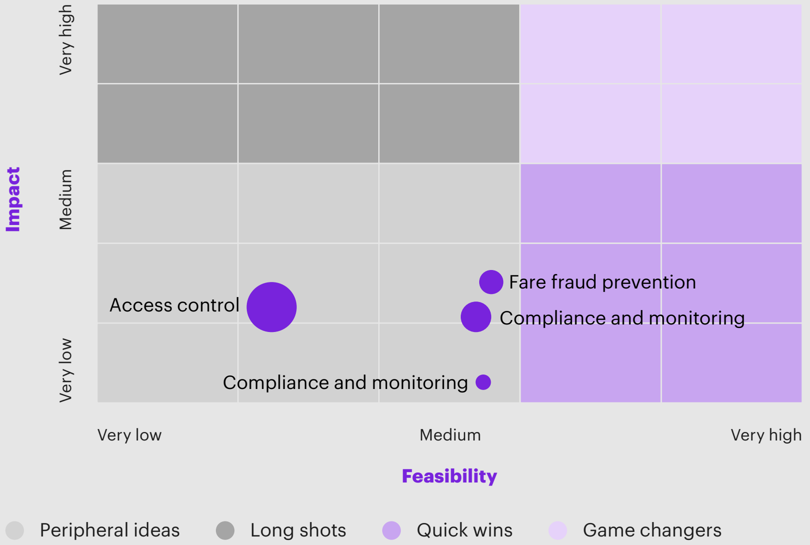

- **Security-related AI use cases are rated low on impact** despite medium feasibility; their adoption is driven by safety and compliance needs (which are critical for operators), but respondents likely don't view them as revenue/efficiency drivers
- **Access control** stands out with highest relevance yet scores low on both impact and feasibility; this perception could be driven by capex-heavy nature of gate hardware

Source: Kearney and MIT analysis

Customer engagement use cases are often highly feasible but of moderate impact. These applications play to LLMs' strengths and are relatively easy to deploy but may have only a limited effect on financial performance, possibly due to commoditization. Even a high-feasibility use case such as automated response management scored relatively low for impact because of concerns about the long-term effects of deploying such systems in the absence of adequate human customer-service interaction (see figure 27).

These charts reinforce an earlier observation: the current prioritization of AI use cases in mobility does not consistently align with their potential impact or feasibility. This misalignment highlights the lack of a coherent, system-wide framework to guide where—and how—AI can generate sustainable value across the mobility sector.

The fragmented AI landscape in mobility shows something deeply important: the industry is not yet approaching AI as an enabler of system-wide transformation. This is, to some extent, understandable: early adoption typically begins with task- and individual-level applications that address specific operational challenges. However, there is still a risk that AI initiatives will remain disconnected: developed as isolated task and individual-level tools rather than as components of an integrated, systems-level architecture.

In such a fragmented environment, the establishment of robust, thoughtful human–AI collaboration models becomes especially important. With multiple AI tools deployed across the sector, the value of each one depends on how effectively humans and AI systems interact. Without clear collaboration frameworks, even the most advanced AI tools risk underperforming since their impact ultimately depends on the quality of the human–AI interface.

Figure 27

Customer engagement innovations favor feasibility over impact, reflecting ease of deployment but modest financial returns

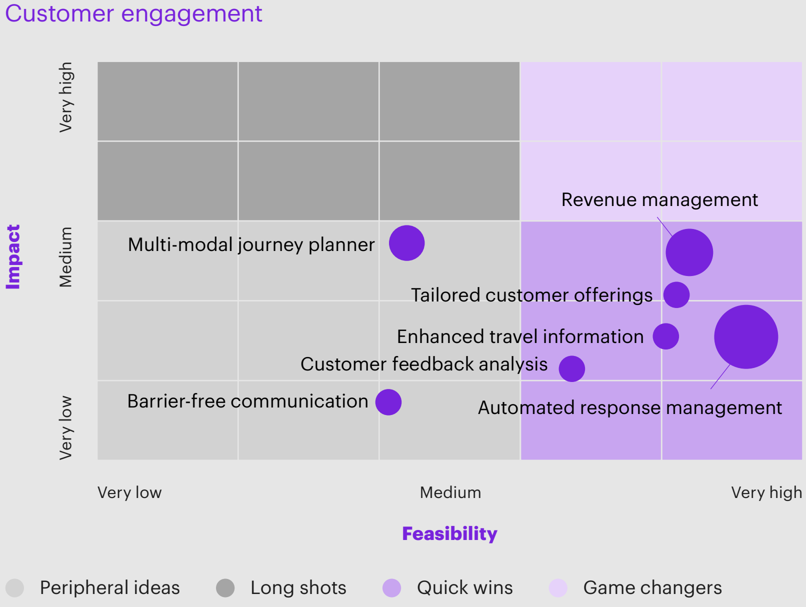

— **Customer engagement use** cases play to the strengths of large language models and are relatively easy to deploy; however, impact scores are generally mid/low. Respondents may view them as commoditized or see limited effect on financials.

— **Automated response management** scores highest relevance and feasibility but generally low impact; while AI can deflect queries, impact on long-term customer loyalty may be limited unless coupled with personalized engagement.

Source: Kearney and MIT analysis

How organizations should think about AI in mobility

In this concluding section of the report, we look at two predominant factors that mobility enterprises should weigh as they think about deploying AI, and we offer four scenarios for the future based on those factors. We then highlight the crucial enabling factors that will allow organizations to derive the greatest value from these transformative technologies.

It all comes down to maturity and trust

While countless trends will influence the course of AI in mobility, two dimensions will be decisive: the industry's maturity and collaboration around AI along with the level of public trust in the technology.

The maturity and collaboration dimension captures both the level of AI maturity the industry has reached and the extent of collaboration it can foster to connect individual tools into cohesive, reliable systems.

Certain questions rapidly present themselves. Are AI models ready to handle the complex and ever-changing demands of the mobility sector, including cyber threats and unexpected—but potentially disastrous—"edge cases"? Will AI models continue to improve in performance, or will they reach physical or economic limits, resulting in diminishing returns? Can the business and government sectors actually attain the necessary level of collaboration?

With regard to public trust, the picture is murky. The adoption of AI tools has skyrocketed since the introduction of ChatGPT in late 2022, with 100 million people logging on in its first two months. Yet surveys widely indicate a profound popular ambivalence toward AI, with roughly equivalent numbers of respondents embracing and rejecting the technology and some clear misgivings about certain mobility applications. For example, one survey by the American Automobile Association (AAA) found that only 13 percent of American drivers are willing to ride in an autonomous vehicle (AAA, Fear in Self-Driving Vehicles Persists, February 25, 2025).

Trust is a vital concept in mobility and transportation; users literally entrust their lives, and those of their loved ones, to the vehicles and infrastructure that enable them to move from point A to point B. The high velocities and usage densities of modern transportation networks mean that even a small error can easily result in death or serious injury.

For AI to attain a more prominent role in the operation of those vehicles and that infrastructure, it will need to earn higher levels of public trust, particularly on the vital question of whether effective AI-governance mechanisms are in place. Whether it will ever reach those levels is one of the central questions facing executives and policymakers in the mobility sector (see figure 28 on page 41 and figure 29 on page 42).

Figure 28

AI’s rapid evolution could last or lose steam before true maturity

AI technology is **advancing.** However, growth can collide with power, funding, or infrastructure **limits.**

AI technology is advancing steadily

— Systems performance vs. benchmarks is improving (e.g., scores on SWE-bench[1] rose by 67.3% in 2024).

— The inference cost is decreasing steadily (e.g., x280 decrease in 2022–2024 for GPT-3.5-like performance).

However, critical bottlenecks may lie ahead

— Data center energy consumption to double by 2030

— Annual AI spending projected to reach $1.3 trillion by 2029 (Can this pace be sustained?)

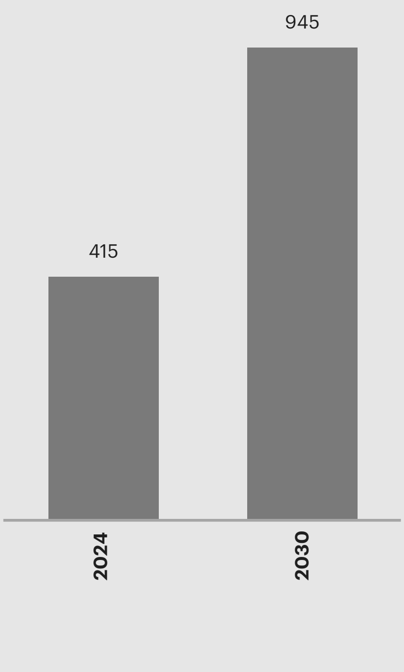

Data center energy consumption, TWh

Industry-specific AI systems increase **performance,** yet future **reliability** and **security** questions remain unresolved.

Industry-specific AI systems have surpassed humans based on technical performance

— Autonomous vehicles show statistically significant improvement on safety benchmarks vs. humans.

— Pilots of AI-optimized traffic signals report ~30% fewer red-light stops and ~10% lower emissions.

However, concerns remain over performance in rare events and cybersecurity risks

Some studies show that autonomous vehicles can be more likely to cause injuries in “edge case” crashes than regular cases

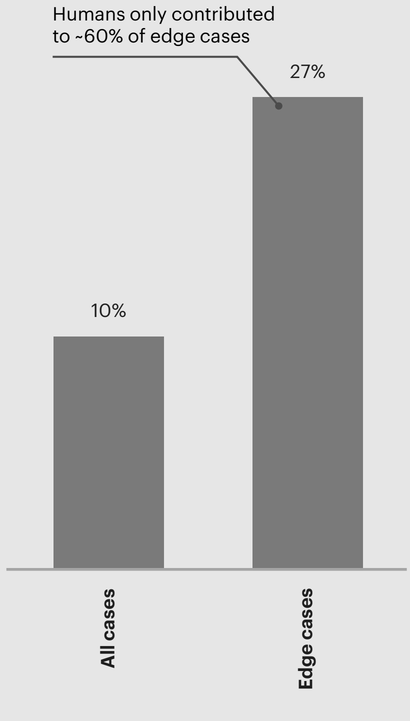

Share of AV crashes with injuries in all AV crashes, NHTSA (2021–2023)

Research on edge cases is scarce, more recent data not available

Growth momentum can be sustained with **system-wide governance and collaboration.** It has yet to take shape.

Governance activity is ramping up

— US policymakers broadly support AI regulations.

— Legislative mentions +21.3% in 2024 across 75 countries.

— The UN created a 39-member advisory body to address issues in the international governance of AI.

However, collaboration on AI within the ecosystem is far from optimal

OECD highlights policy, incentive, and interoperability gaps and calls for more data sharing.

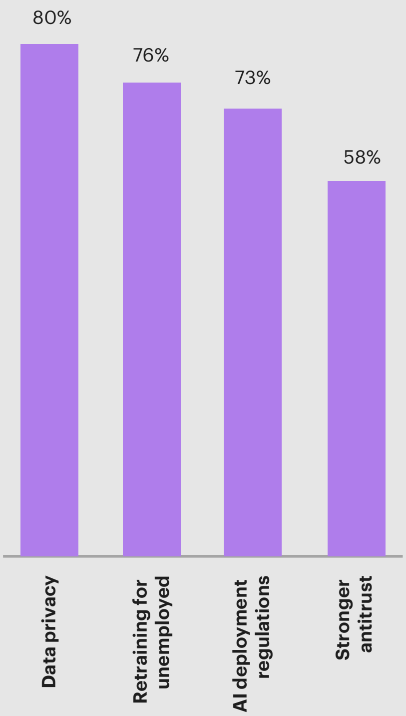

Favored AI policies by local US officials (% agree)

[1] Evaluation dataset used to measure how well LLMs can understand, debug, and fix real-world software code

Sources: HAI, IEA, IDC, OECD, IMF, “Safety in higher level automated vehicles: Investigating edge cases in crashes of vehicles equipped with automated driving systems,” Accident Analysis & Prevention, August 2024; Kearney and MIT analysis

Figure 29

Public trust in AI and its governance is divided; it could strengthenover time or gradually erode

Society is **divided on AI technology:** its adoption and use and whether the benefits outweigh the risks.

AI faces both enthusiasm and resistance from the public

— Surveys show sentiment is divided between embrace (30%) and rejection (35%).
— 42% say AI benefits outweigh risks; 32% say risks outweigh benefits.
— Some systems not yet trusted (e.g., only 13% of U.S. drivers would trust riding in an AV)

At the same time, the adoption of AI tools has skyrocketed since the introduction of ChatGPT in 2022:

— 100 million people logged onto Chat GPT in its first two months.
— 21% of respondents already use AI tools daily in their personal lives.

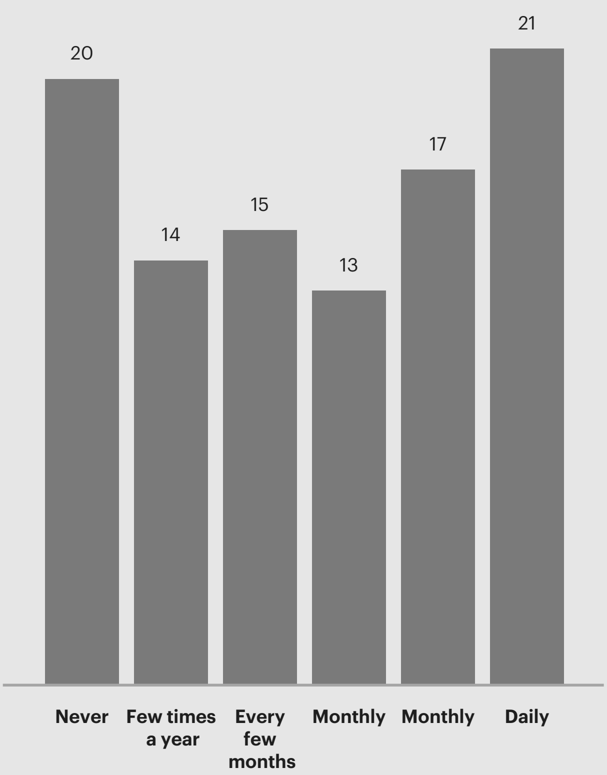

Frequency of AI use for personal, work, or study purposes (% respondents)

Public trust in **effective AI governance** will be key to stimulate adoption, yet that trust remains mixed today.

Business leads on public trust to introduce innovations

— 63% of respondents trust business, and 59% trust it to integrate innovation in society
— The public wants partnership: 60% say they'd trust business more on tech-led change if it partners with government—up 15% since 2015.

However, trust in public institutions on governing innovation is low

— 59% of respondents believe governments lack adequate understanding of emerging technologies to regulate them effectively.
    — An alarming sign as confidence in regulation correlates strongly with adoption
— Only 47% trust AI companies to protect their personal data.

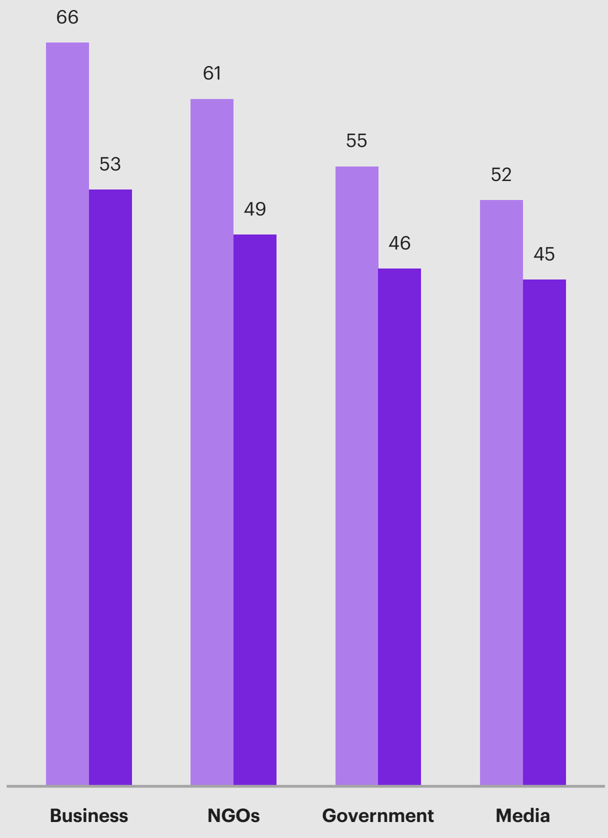

I trust each with the introduction of innovations (% respondents)

Low-income respondents   High-income respondents

Sources: Edelman, KPMG, AAA, HAI; Kearney and MIT analysis

Four scenarios of AI in mobility

Given the uncertainties across the two dimensions outlined above—and the fact that both dimensions are highly capable of shifting in either direction over the months and years to come—we see four potential future scenarios for AI in mobility over the next decade (see figure 30 on page 44). The qualities of offering maturity and public trust can be plotted on separate axes, to produce four scenarios, which we now examine in turn.

Scenario 1: Golden AI Age (high maturity, high trust)

Under this scenario of open, interconnected platforms, AI in mobility is no longer a novelty but a common feature of transportation infrastructure. High consumer trust and strong industry collaboration lead to some significant societal benefits (see figure 31 on page 45).

Shared standards and interoperable systems are established, and autonomous mobility systems become both more reliable and more common. System-level AI applications are successfully deployed, enabling the emergence of integrated multimodal public and private networks and real-time mobility optimization. Consumers and societies enjoy enhanced safety, improved accessibility, reduced congestion, and diminished carbon emissions across transportation networks.

Under this scenario, proper human–AI collaboration models and robust governance are established across system-wide AI applications, allowing organizations to effectively navigate the jagged frontier. “Human in the loop” applications such as digital-twin infrastructure planning emerge alongside “human out of the loop” apps such as passenger AI companions that can integrate journey planning, payments, accessibility, and real-time updates across modes and geographic locations.

Scenario 2: Fragmented Walled Gardens (low maturity, high trust)

This scenario features a high degree of public enthusiasm for AI but a stall-out in the development and adaptation of the technology itself. This leads to closed systems and limited interoperability as competing proprietary platforms create confusion, duplicative efforts, and uneven outcomes.

The progress that does occur tends to be clustered at the level of individual and organizational applications. While there might be some high-quality AI products and services throughout the sector, integration would be problematic as governance and institutional collaboration lag—a serious challenge for companies and transit agencies alike. Humans tend to over-rely on AI, undermining outcomes when tasks fall “out of frontier.”

Scenario 3: Minimal AI Impact (low maturity, low trust)

Poor industry coordination and consumer resistance result in minimal adoption and impact. Privacy concerns and lack of trust hinder the deployment of AI technologies, and what deployments do occur tend to be unreliable, opaque, and poorly integrated. Not surprisingly, consumers lose confidence and interest in additional applications of AI in mobility.

On the bright side, consumers and societies avoid the potential risks associated with the large-scale deployment of AI. But they also lose access to whatever potential benefits the technology could offer, beyond some functional applications at the task or individual level.

Figure 30

Given the uncertainties across the two dimensions, we see four potential scenarios for the next decade for AI in mobility

AI in mobility scenarios

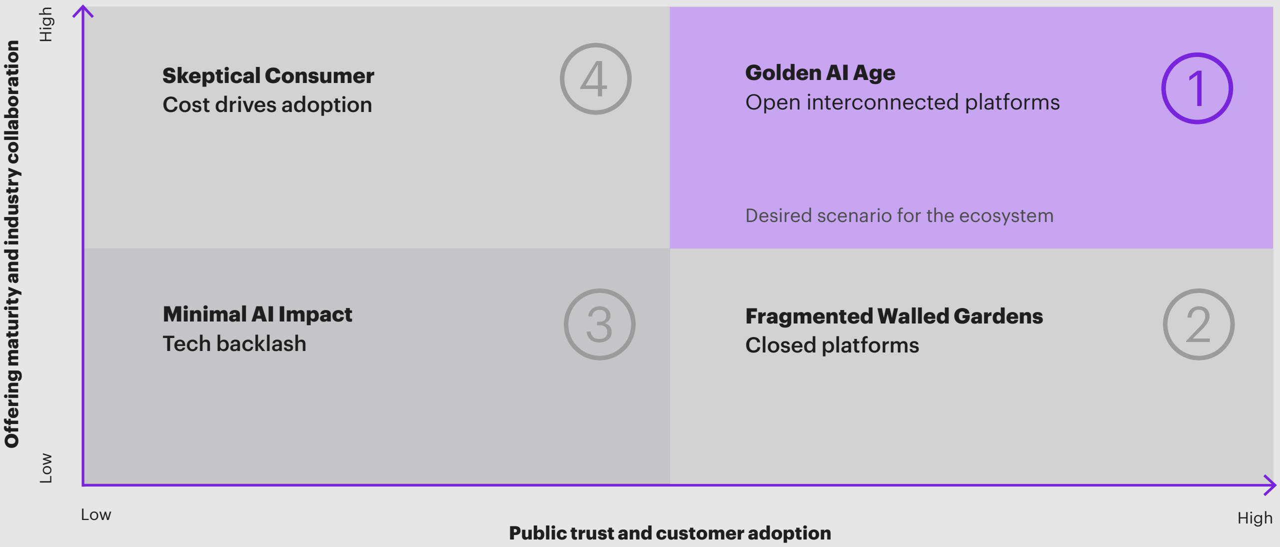

Scenario	Context	AI applications
1 Golden AI Age	— Consumers **trust AI** to manage their travel safely and transparently. — Ecosystem works with **shared** standards, interoperable **systems, and governance.** — AI has moved from novelty to **infrastructure.**	— **System-wide AI** applications are deployed and create new outcomes in mobility (zero fatalities, etc.). — Strong presence of fully autonomous mobility AI systems (HOOTL).
2 Fragmented Walled Gardens	— **Consumers adopt** AI-powered mobility, but governance and institutional alignment lag. — Players deploy **proprietary** ecosystems. — Competing platforms create duplication and uneven outcomes; some offering concerns are unresolved	— **Individual/organizational-level** AI dominates the landscape. — Humans often over-rely on AI, undermining mobility outcomes when tasks fall "out of the frontier."
3 Minimal AI Impact	— AI deployments are **unreliable, opaque, and poorly integrated.** — Consumers have **lost confidence and interest.** — Institutions lack the **competence** or resources to correct course.	— Limited AI implementation, mostly at the task or individual level. — Lack of trust and inadequate governance keep humans in **a strong supervisory role** over AI.
4 Skeptical Consumer	— AI tech is sophisticated and reliable, but adoption is supply-driven. — Ecosystem participants deploy mature systems, e.g., **for cost, safety, and sustainability** gains. — The public remains uneasy or disengaged.	— **Organization-level AI** becomes the standard across mobility. — Optimal human + AI collaboration models are deployed in organizations.

Sources: Kearney and MIT analysis

Figure 31

The Golden AI Age scenario is the optimal outcome for the ecosystem, driven by system-wide AI and effective human–AI collaboration

Golden AI Age

System-level AI applications are successfully deployed, enabling superior mobility outcomes:
— Integrated multimodal public and private networks
— Real-time mobility optimization
— High safety, minimized journey time, low emissions, equitable access, etc.

Proper human–AI collaboration models and robust governance are established across system-wide AI applications **to navigate the "jagged frontier"**

Possible system-level AI applications and collaboration models

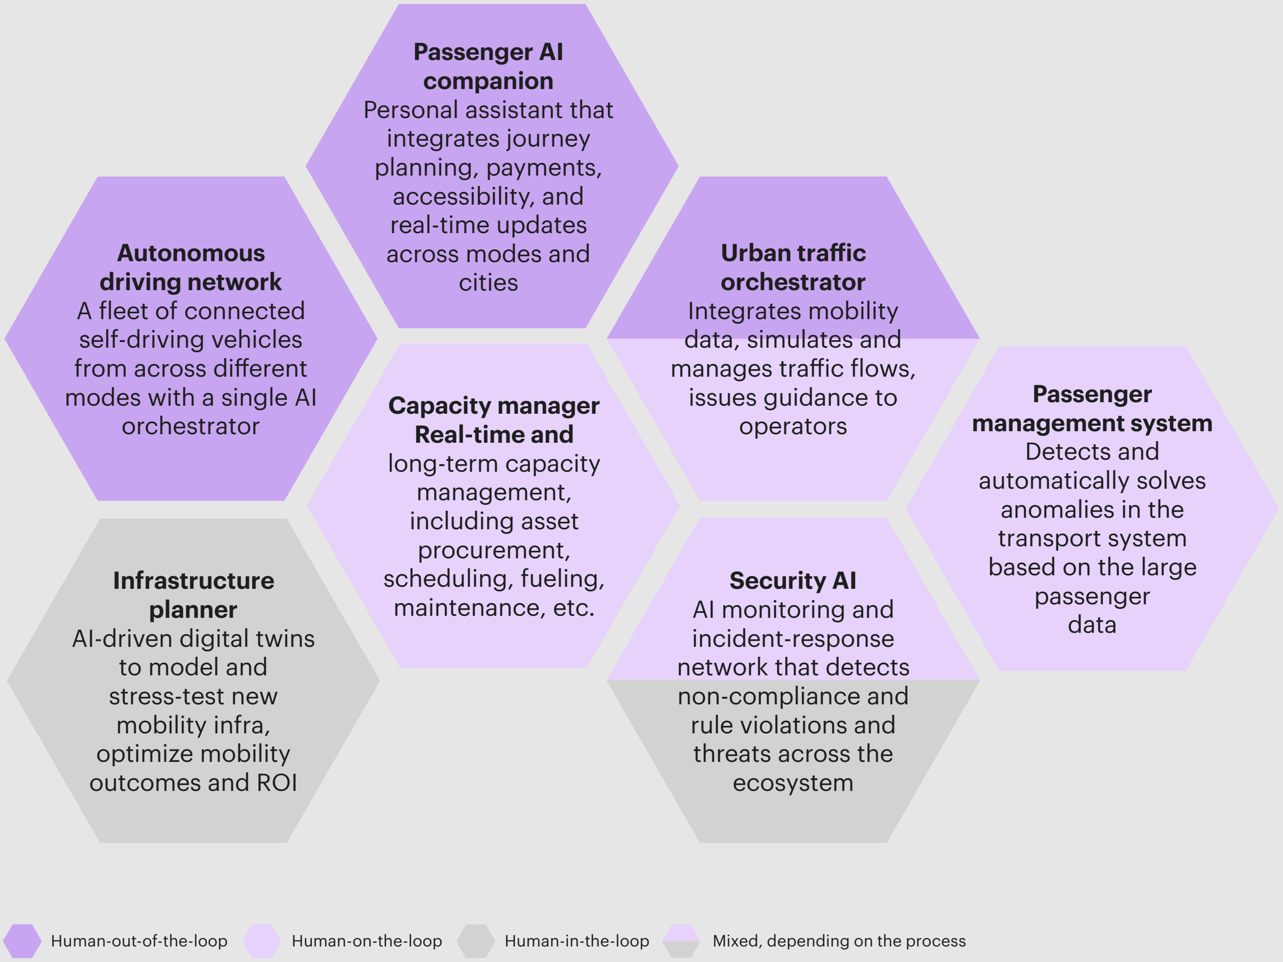

Sources: Kearney and MIT analysis

Scenario 4: Skeptical Consumer (high maturity, low trust)

What if AI applications attain a high level of maturity, but the public still doesn't trust it? That's the premise behind this scenario, in which rollouts of AI mobility apps are supply-driven and somewhat piecemeal, driven mainly by the industry's need to shave costs or to pluck some low-hanging safety or sustainability gains.

Applications at the organizational level, deploying ever more sophisticated human–AI collaboration models, set the standard across the mobility sector, with public trust still too low for broader systemic breakthroughs.

Interestingly, most of our participants indicated that this scenario might be the likeliest of the four over the next decade—and that, given enough time, this state of play would naturally evolve into a high maturity/high trust condition as public attitudes toward AI shift.

It is important to recognize that these four scenarios may coexist: one may prevail in one part of the world, while another takes hold elsewhere. Varied regulatory and institutional contexts are pushing regions onto different development trajectories, and this could lead to a future where multiple AI scenarios play out simultaneously.

For example, the United States has largely embraced a market-driven approach, which plays to the country's strengths as a traditional magnet of top-tier technical talent and leader in fundamental research. China, by contrast, has undertaken a state-driven model: a "whole of nation" strategy featuring a growing focus on internal talent development and data training. The European Union has sought to differentiate itself in the market through the establishment of a comprehensive regulatory framework as well as through the quality and scale of its research and industrial sectors.

Regardless of the locale in which it operates, AI can build toward system-wide transformation by building on existing applications. Fragmented task-level and individual-level AI tools can become more integrated through shared data rails, interoperable platforms, orchestration layers, and joint governance frameworks.

As noted earlier, the success of this integration will ultimately depend on the establishment of proper human–AI collaborations across existing AI applications, with a keen awareness of where those applications sit at any given moment along the jagged frontier. The ecosystem also needs to establish proper human–AI collaboration models across existing AI applications (see figure 32 on page 47).

Starting AI enablement from within organizations

Now that technical capability is no longer the main obstacle in most cases of commercial-scale deployment of AI applications, scaling AI in the mobility sector requires a set of organizational, infrastructural, and strategic enablers. Several leading enterprises, ranging from transit authorities to private operators, are beginning to address these enablers and implement AI solutions successfully.

Below, we share the five primary enablers of successful AI adoption within transit organizations.

Enabler 1: Strategic vision and executive leadership

AI integration is most successful when it is supported by clear strategic intent and vision. Organizations with dedicated digital transformation teams, long-term AI road maps, and C-level top-down leadership and support are better positioned to align cross-functional resources and prioritize investments. In the Middle East, for instance, national AI agendas and smart mobility mandates (such as Saudi Vision 2030) have prompted alignment between the public and private sectors around AI objectives.

What's needed is a clear vision of what AI can solve for. Reducing traffic fatalities from 1.19 million per year to zero with the use of AI, for instance, could be a societal and overarching vision for all nations and organizations (WHO, 2023). Such an ambition may sound far-fetched, but it has already spawned a major mobility company, Waymo, which grew directly from an Alphabet "moonshot" project to solve the puzzle of safe, broadly accessible autonomous driving. As seen in that example, leadership commitment can foster extraordinary internal momentum, with potentially market-shifting results.

Figure 32

The ecosystem also needs to establish proper human–AI collaboration models across existing AI applications

Illustrative

	"Inside" the frontier		"Outside" the frontier
Vehicle design and manufacturing	**1.4.** Advanced driver assistance systems **1.3.** Energy management optimization	**1.2.** Design, testing, and prototyping **1.5.** Autonomous driving (Note: L4)	**1.1.** Discovery of new materials **1.5.** Autonomous driving (Note: L5)
Infrastructure design		**2.1.** Location and network planning	**2.2.** Demand simulation **2.3.** Disruption prediction
Traffic management and regulation	**3.6.** Automated traffic enforcement **3.4.** Toll enforcement **3.5.** Parking enforcement **3.2.** Real-time parking optimization	**3.7.** Driver behavior monitoring **3.1.** Adaptive traffic control	**3.8.** Smart policy design
Operations	**4.3.** Route optimization **4.1.** Fleet and route planning	**4.6.** Fuel/energy saving **4.2.** Capacity planning and scheduling **4.4.** Daily fleet management **4.7.** Bus bunching reduction	**4.8.** Assistants for operations **4.5.** Personnel training and planning
Maintenance	**5.2.** IoT-supported maintenance	**5.3.** Predictive maintenance **5.1.** Inventory and supply chain optimization **5.4.** Maintenance training and assistants	
Security and monitoring		**6.4.** Compliance and monitoring **6.2.** Fare fraud prevention **6.1.** Occupancy and crowd monitoring	**6.3.** Access control
Customer engagement	**7.2.** Enhanced travel information **7.6.** Automated response management **7.7.** Customer feedback analysis	**7.4.** Revenue management **7.1.** Seamless multimodal integration	**7.5.** Barrier-free and inclusive communication **7.3.** Tailored customer offerings

Source: Kearney and MIT analysis

Enabler 2: Internal champions and use-case focus

The success of AI in mobility depends on the targeted, high-impact use of AI applications backed by internal stakeholders who understand domain workflows and the capabilities of these rapidly evolving technologies.

We are already seeing examples of this in the field. Transit agencies that deploy task-specific tools, such as feedback and complaint parsing using LLMs or predictive maintenance using DL, have demonstrated clear performance benefits. The autonomous-delivery company Gatik focused narrowly on middle-mile driverless trucking for fixed business-to-business routes, enabling operations with Walmart and other customers.

Enabler 3: Optimal human–AI collaboration model

Industry pioneers are increasingly seeking assistance in applying modular, standardized, human-centered AI designs and narrowing the scope of each application to facilitate validation, oversight, and iterative improvement. To bring this about, they are working their way toward ever more sophisticated balancing between the roles of humans and AI.

Such balancing can take a variety of forms, as noted earlier. Take one example from actual business practice: AC Transit uses computer vision to detect and report parking violations; while AI collects evidence, staff members ultimately decide whether to issue a citation to the vehicle owner.

Enabler 4: Data infrastructure and integration

Access to diverse, high-quality, and timely data is essential to the development of effective AI mobility solutions. Leading organizations are investing in robust data pipelines, proprietary data sources, and real-time telemetry systems that integrate infrastructure and vehicle-based data.

Legacy infrastructure is being upgraded with sensors, cloud platforms, and open application programming interfaces (APIs), enabling seamless data flow across planning, design, and operations. The Swedish transport company Einride developed the Saga platform, providing a data layer for the teleoperation of driverless vehicles, reskilling drivers into "pod operators" who manage fleets remotely.

Enabler 5: Ecosystems, collaborative innovation, and governance

AI maturity and deployment often depend on systemwide participation by numerous actors, both public and private. Many of our respondents stated that their own AI solutions are sourced externally. Strategic partnerships—such as those we've already seen between Waymo and Uber or between Mercedes Benz and Microsoft—enable speed, innovation, and domain-specific AI models.

Collaboration could also take the form of linkups between private companies and nonprofits, universities, or public agencies. For example, the European Union's funding instruments and pilot grants offer valuable support for early-stage testing and scaling.

Some organizations are even reshaping their procurement models altogether, bundling tenders for platforms, consulting, and data services to drive integrated AI adoption. One notable instance of imaginative collaboration is a consortium of Aurora, FedEx, PACCAR/Volvo, and governmental regulators, to align on critical questions of safety validation and reporting.

Ultimately, stakeholders across the mobility ecosystem must act together to elevate AI to the next level (see figure 33 on page 49).

These and other developments, including many that we have identified throughout this report, generate grounds for optimism that the mobility sector may be on the verge of finding the right balance between caution and innovation in introducing the broader use of AI applications.

Figure 33
Elevating AI in the mobility ecosystem hinges on collaboration

Ecosystem actions to enable the Golden AI Age scenario

1	Create strong alignment on goals of AI adoption and system trade-offs (economics, equity, etc.).
2	Establish and commit to collective governance mechanisms (AI board, operations council, etc.).
3	Co-create a robust mobility AI regulatory framework.
4	Establish shared data rails, including data standards, interoperability, privacy protection, data sharing, etc.
5	Fund and execute necessary smart infrastructure projects (e.g., IT infra, sensors, connectivity, etc.).
6	Win public trust with communication, citizen participation, transparency.

Focus for the actors along the value chain

Policymakers, regulators, authorities	— Lead vision and regulation, coordinate standards, enable testing. — Enable safe and secure mobility data sharing protocols. — Incentivize or lead smart infrastructure projects. — Lead public trust campaigns.
Infrastructure developers	— Enable sharing of standardized infrastructure data. — Instrument mobility assets for data collection and AI use. — Lead smart infrastructure development.
Mobility operators	— Set up and engage in pilots and sandbox environments. — Enable operational data sharing. — Support and integrate in smart infrastructure. — Lead public trust campaigns.
OEMs, suppliers, financial and technology players	— Lead standards definition, and develop ecosystem platforms. — Facilitate cross-ecosystem data sharing. — Support/fund smart infrastructure. — Lead joint "innovation labs" for AI projects across the ecosystem.

Source: Kearney and MIT analysis

Authors

Christian Gasparic
Partner, Dubai
Kearney
christian.gasparic@kearney.com

John Moavenzadeh
Executive Director
MIT Mobility Initiative
moavenza@mit.edu

Andres Mendoza Pena
Partner, Chicago
Kearney
andres.mendozapena@kearney.com

Bhuvan Atluri
Associate Director of Research
MIT Mobility Initiative
bpatluri@mit.edu

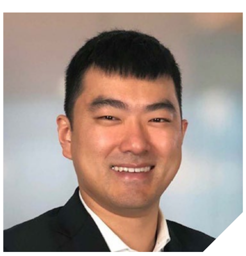

Benny Zheng
Principal, Chicago
Kearney
benny.zheng@kearney.com

Marvin Greifenstein
Researcher
MIT Mobility Initiative
marvingr@mit.edu

Matthias Feldotto
Principal, Düsseldorf
Kearney
matthias.feldotto@kearney.com

Valerii Kurepanov
Consultant, Dubai
Kearney
valerii.kurepanov@kearney.com

The authors wish to thank Ammar Agha, Jason Irias, Jeff Kolker, and Maximilian Marienhagen for their valuable contributions to this report.

MIT Mobility Initiative (MMI)

The MIT Mobility Initiative (MMI) is a global platform to accelerate a safe, clean, and inclusive mobility system. MMI brings together researchers, industry leaders, policymakers, and entrepreneurs to advance the future of transportation through cutting-edge research, world-class education, technology commercialization, and high-impact engagement.

mmi.mit.edu

About Kearney

Since 1926, Kearney has been a leading management consulting firm and trusted partner to three-quarters of the Fortune Global 500 and governments around the world. With a presence across more than 40 countries, our people make us who we are. We work impact first, tackling your toughest challenges with original thinking and a commitment to making change happen together. By your side, we deliver—value, results, impact.

kearney.com

MIT Mobility Initiative

KEARNEY