AI and AI Governance in Automotive Dealerships

A STAR Whitepaper for Franchised Automobile Dealers

Executive Summary

Artificial Intelligence (AI) is transforming the automotive retail landscape, offering franchised dealers new ways to streamline operations, personalize customer experiences, and drive profitability. Auto dealerships are leveraging AI for everything from intelligent marketing and customer relationship management to predictive maintenance and inventory optimization. In fact, recent surveys show an overwhelming majority of dealerships plan to increase AI investments, as many have already seen significant boosts in ROI from early AI initiatives [1] [2]. AI-powered tools—such as chatbots for customer service, machine learning models for demand forecasting, and computer vision for vehicle inspections—are quickly becoming part of the dealer toolkit.

However, **with these opportunities come significant risks and responsibilities**. AI systems can introduce **new challenges around data privacy, security, transparency, and ethics**. Dealerships handle sensitive customer information and must ensure AI solutions do not misuse data or inadvertently violate regulations. The **risks of ungoverned AI** include potential privacy breaches, biased decision-making, misinformation (AI "hallucinations"), and loss of customer trust [3]. Compounding this is a fast-evolving regulatory environment in the United States and abroad, which is increasingly focusing on AI oversight and consumer protection. Dealers must navigate state and federal rules (from privacy laws to upcoming AI-specific regulations) and be prepared for compliance requirements that demand accountability in AI usage.

AI governance is therefore essential. Just as a dealership wouldn't deploy new machinery without safety protocols, AI tools require **clear policies, employee training, and oversight**

---

[1] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

[2] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

[3] Five Ways AI Will Disrupt Car Dealerships in 2025

mechanisms to ensure they are used ethically and effectively [4]. This paper, published by STAR (Standards for Technology in Automotive Retail), provides an educational overview of AI concepts and practical guidance on governance tailored to auto dealers. It covers foundational AI definitions (including a new section on *Agentic AI*—AI systems with greater autonomy), outlines dealership-specific use cases and benefits, identifies key risks, and offers **technical and non-technical recommendations** to mitigate those risks. It also summarizes relevant regulatory guidance so dealers can stay ahead of compliance obligations. Throughout, we include checklists and best practices to help dealership leaders ask the right questions and implement governance frameworks that balance innovation with caution.

In short, **the goal of this white paper is to equip franchised automobile dealers with the knowledge to harness AI's potential responsibly**. Dealers who proactively adopt AI with strong governance will likely **outpace competitors in efficiency, customer satisfaction, and innovation**, while those who do nothing risk falling behind in an AI-driven future [5]. By understanding AI and implementing the right controls, dealerships can embrace new technologies with confidence—improving their operations and customer experiences without inviting undue risk.

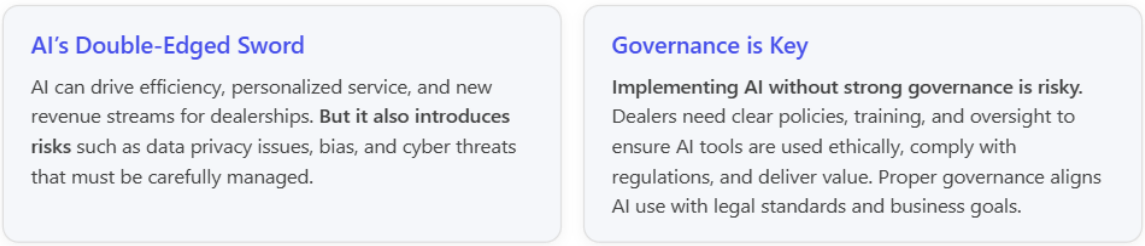

Introduction: AI in Automotive Retail

AI — the simulation of human intelligence in machines — is already at work in automotive retail, though sometimes behind the scenes. Whenever a dealership website's chatbot answers a shopper's question, or a system predicts when a car will need service, **AI is the engine making it possible**. In simple terms, AI enables computers to **learn from data, recognize patterns, and make decisions** with minimal human intervention. Key technologies under the AI umbrella include **machine learning, natural language processing, computer vision,** and more. These technologies allow AI to perform tasks like understanding spoken or written language, analyzing images, and optimizing complex decisions.

---

[4] From Ideation To Autonomy: A Strategic Framework For Scaling ... - Forbes

[5] Five Ways AI Will Disrupt Car Dealerships in 2025

Why it matters now: Customers today expect quick, personalized service. They are already accustomed to AI-driven recommendations and virtual assistants in other shopping experiences [6]. Competitors (including large dealer groups and online auto retailers) are investing in AI to gain efficiencies and insights. In this context, **AI is becoming a competitive necessity for dealerships**. It can automate mundane tasks (freeing up staff for higher-value work), uncover sales opportunities, and improve accuracy in everything from lead follow-up to stock management. A few high-impact examples:

- **Automating routine tasks:** AI can handle repetitive administrative work (appointment scheduling, data entry, follow-up emails), reducing errors and saving staff time.
- **Enhancing customer engagement:** AI-driven chatbots and virtual assistants provide instant responses to customers 24/7, improving lead capture and customer satisfaction by being available when staff are not.
- **Data-driven decision making:** Machine learning models can analyze past sales, local market trends, and customer preferences to help forecast demand and optimize inventory and pricing strategies.
- **Service optimization:** AI can predict maintenance needs by analyzing vehicle data, allowing dealers to schedule service proactively and keep customers happy by preventing breakdowns [7] [8].

Despite these benefits, many dealers are still figuring out how to start with AI. It's important to remember that **AI is not a single off-the-shelf solution but a spectrum of technologies and applications**. Some AI tools are relatively simple (e.g., a basic chatbot script), while others are far more complex (e.g., an autonomous agent that manages a whole workflow). Understanding the different types of AI and their capabilities is a foundation for making smart adoption decisions. The next section demystifies key AI concepts and terminology relevant to dealerships, including the emerging concept of *Agentic AI*, which represents the next frontier of autonomous, goal-driven systems.

Key AI Concepts and Terminology

To engage in meaningful AI discussions and make informed decisions, dealership stakeholders should understand a few **core AI terms**. The following table summarizes some key concepts and how they relate to automotive retail:

---

[6] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

[7] Five Ways AI Will Disrupt Car Dealerships in 2025

[8] Five Ways AI Will Disrupt Car Dealerships in 2025

Term	Definition	Relevance to Auto Dealers
Artificial Intelligence (AI)	The broad field of creating machines or software that simulate human intelligence–learning, reasoning, and problem-solving. Includes subfields like ML, NLP, computer vision, etc.	AI is the engine behind modern dealership tech, powering features from smart CRMs to vehicle recognition systems. It enables automation and smarter decision-making in sales, marketing, and service.
Machine Learning (ML)	A subset of AI where algorithms improve through experience. ML models learn from training data rather than being explicitly programmed with fixed rules. Common types: Supervised (learning from labeled examples), Unsupervised (finding patterns in unlabeled data), and Reinforcement Learning (learning via feedback/rewards).	ML underpins many dealer-focused AI tools. For example, ML models can learn to predict vehicle demand or identify customer purchase likelihood based on historical data. Dealers don't need to tweak code; the model adapts as it learns, though knowing the source and quality of training data is critical.
Generative AI	AI that creates new content (text, images, videos, code, etc.) in response to prompts. These models (like GPT-4, DALL-E, etc.) learn patterns from vast datasets and generate original outputs.	Generative AI can automate content creation for dealers: writing vehicle descriptions, crafting marketing copy, creating chat responses or even generating promotional imagery. It offers creativity and efficiency, but may produce errors or "hallucinations", so outputs need review.
Conversational AI	AI systems designed to engage in human-like dialogue, understanding natural language and responding appropriately. Examples include chatbots and voice assistants.	Dealers use conversational AI for customer service chatbots on their websites and virtual assistants (text or voice) that guide customers through inquiries or even in-car systems. It provides scalable customer engagement, but dealers should disclose when AI is interacting and ensure conversations are recorded for quality control.
Agentic AI	Highly autonomous AI "agents" that can make decisions and take actions toward achieving goals independently of constant human guidance. These systems have the ability to adapt, invoke other tools	Emerging in 2025, agentic AI is deemed the "next big trend". In auto retail, an agentic AI might handle a multi-step process end-to-end (e.g., monitoring inventory and automatically reordering popular cars or

	or services, and persist towards objectives over time.	managing an entire customer inquiry through to resolution across systems). While promising for efficiency, it carries new risks since the AI has more freedom. Strong guardrails and human oversight are needed before deploying autonomous agents in dealership operations.
Artificial General Intelligence (AGI)	A hypothetical future AI that possesses human-like broad intelligence, capable of understanding or learning any task a human can. Not yet achieved; current AI is "narrow" (specialized).	AGI and the even more theoretical Artificial Superintelligence (ASI) often appear in futurist discussions. For practical purposes, dealers need not worry about AGI today—it remains speculative. The focus now is on narrow AI applications that address specific dealership needs.

Table: Key AI terms and their relevance. Dealers should focus on narrow AI applications while keeping an eye on emerging agentic AI capabilities.

Agentic AI – A New Class of Autonomous Systems

One concept warranting special attention is **Agentic AI**, because of its growing prominence and its potential impact on how workflow tasks might be automated in the near future. **Agentic AI** refers to AI systems endowed with a degree of *agency* – meaning they can perceive their environment, make decisions, and act upon those decisions over extended sequences, all with minimal human intervention. In essence, these are **autonomous AI agents** that can string together multiple actions to achieve a goal.

Agentic AI goes beyond single-task chatbots or prediction models. For example, imagine an AI agent that monitors your service department scheduling system: it detects a slow day next week, identifies customers due for maintenance, automatically drafts personalized outreach messages, and schedules appointments – all on its own. Such an agent would need to integrate language understanding (to communicate), decision logic (to select which customers to contact and when), and perhaps even access to dealership systems (scheduling and CRM).

This level of autonomy is **cutting-edge**. Gartner has named agentic AI as the #1 strategic technology trend for 2025 [9], highlighting how it enables operations to run with less direct human micromanagement. Early examples in other industries include AI agents in cybersecurity that hunt threats and respond in real-time, or in software development where

[9] From Ideation To Autonomy: A Strategic Framework For Scaling ... - Forbes

agents assist by writing and testing code. In automotive retail, agentic AI is just starting to appear in pilot programs. Some forward-looking dealerships are experimenting with autonomous sales assistants that can cross-reference inventory with customer web activity and initiate contacts, or service bots that coordinate parts ordering when inventory runs low.

Benefits: If successful, agentic AI could dramatically improve efficiency by handling complex, multi-step tasks. It's like having a proactive virtual employee. These agents can work 24/7, react faster than humans to certain triggers, and potentially discover novel solutions to problems by exploring data links that a human might miss.

Risks: With greater autonomy comes greater risk. By making decisions without immediate human approval, an AI agent might commit errors or actions the dealership wouldn't intend. For instance, an overeager sales agent might push too hard and spam customers, or an autonomous system might make inventory purchases outside desired budget limits. There's also the expanded **attack surface** concern: an agentic AI connected to various systems could be exploited if not properly secured (imagine a malicious actor manipulating an agent's API access). The **lack of human oversight in the loop** means dealers must set strict boundaries and failsafes. Agentic AI systems should defer to humans when there's uncertainty [10] – for example, requiring human confirmation before finalizing a large order initiated by an AI. Additionally, detailed logging and "reasoning transparency" become crucial so that if an agent does make a mistake, the dealership can audit what happened [11].

Practical takeaways: For most dealerships, truly agentic AI is not a today issue, but it is on the horizon. It's wise to be aware of the concept because software vendors may soon market agentic features in their products. Dealers should approach these with **healthy skepticism and governance**. Insist on the ability to configure limits (what the agent can and cannot do) and start with pilot programs that keep a human in the loop. Over time, as trust builds and proper safeguards are proven, agentic AI might take on more routine decision-making. This phased approach ("crawl, walk, run" in adopting autonomous agents) is recommended to manage risk while exploring the potential [12] [13].

---

[10] From Ideation To Autonomy: A Strategic Framework For Scaling ... - Forbes

[11] From Ideation To Autonomy: A Strategic Framework For Scaling ... - Forbes

[12] From Ideation To Autonomy: A Strategic Framework For Scaling ... - Forbes

[13] From Ideation To Autonomy: A Strategic Framework For Scaling ... - Forbes

AI Use Cases and Opportunities for Dealerships

AI's value becomes more tangible when looking at **specific use cases** in a dealership setting. Below are several areas where AI is already making an impact or holds strong potential. Each use case is accompanied by examples of how it works and the benefits it can bring:

- **Sales and Marketing Optimization:** AI is a powerful tool for capturing and nurturing leads. For example, **AI-driven chatbots** on dealership websites can engage visitors in real time, answer common questions, and schedule test drives or service appointments. This ensures no lead goes unattended, even after hours, and provides instant responses that today's customers expect. AI can also analyze CRM data to score leads, helping salespeople prioritize whom to call first. On the marketing side, AI algorithms can personalize email or ad campaigns based on customer behavior patterns. Dealers using AI for marketing and advertising have been able to significantly boost engagement – as of late 2024, about 34% of dealerships were utilizing AI in marketing, a number expected to grow rapidly [14] [15]. The result is a more efficient marketing spend and higher conversion rates, as messages are tailored and sent at the right time to the right prospects.
- **Customer Service and Experience:** In the showroom of the future, a customer might be greeted by an AI system that already knows their preferences. **Virtual assistants and voice interfaces** can be deployed via kiosks or mobile apps to answer customer queries ("Which SUVs have 3rd-row seating?") or guide them to the appropriate department. AI can also power self-service portals where customers schedule services or check vehicle order status through natural language queries. These tools enhance the experience by providing immediate, interactive service without always needing staff intervention. Crucially, they free up human employees to handle more complex inquiries or offer a personal touch where it matters. Most dealers view AI as an enhancement, not a replacement, to their staff's role [16] [17] – for instance, by handling initial FAQs, AI allows your service advisors and salespeople to focus on in-depth consultations.
- **Inventory Management and Pricing:** Inventory is the lifeblood of sales, and balancing it is a constant challenge. **Machine learning models** can analyze sales history, local market data, seasonal trends, and even social media sentiment to

[14] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

[15] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

[16] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

[17] Five Ways AI Will Disrupt Car Dealerships in 2025

forecast which vehicles will be in demand [18]. This helps in ordering the right mix of models and trims. ML can also assist in dynamic pricing – adjusting vehicle prices or incentives based on factors like time on lot, market supply, and customer interest levels. Some advanced dealer platforms offer AI that suggests when to transfer cars between dealerships or when to run a special promotion to move aging stock. Optimizing inventory with AI leads to fewer overstocked vehicles (reducing floorplan costs) and less chance of missing a sale due to a stockout. In essence, AI helps put the right car at the right price in front of the right customer at the right time.

- **Service and Maintenance:** Dealership service departments are leveraging AI for **predictive maintenance** and improved diagnostics. Modern vehicles generate heaps of data (through OBD-II ports, telematics, and connected car apps). AI systems can parse this data to predict component failures or maintenance needs *before* the customer experiences a problem [19]. For example, an AI might learn that a certain model of car tends to need a battery replacement around a certain mileage combined with certain usage patterns; the system can flag upcoming appointments for those cars to suggest a pre-emptive battery check or replacement. This proactive approach minimizes breakdowns and enhances customer trust ("the dealership helps me avoid problems"). Additionally, computer vision AI can assist technicians by analyzing vehicle photos or camera feeds for damage or wear (imagine an automated tire tread depth estimator via image). Such tools speed up inspections and upsell opportunities (like recommending tire changes at the optimal time). Overall, AI in service can lead to higher throughput (by streamlining diagnostics), increased parts and service revenue (through predictive insights), and improved customer satisfaction (through preventative care).
- **Fraud Detection and Finance:** In the finance and insurance (F\&I) office, AI can be used to detect fraudulent documents or identify high-risk loan applications by spotting anomalies that a human might miss. Dealerships handling numerous credit applications could employ machine learning models that flag potential identity theft or loan stacking issues using pattern recognition. Similarly, for online transactions or remote sales, AI vision tools can verify IDs against selfies to reduce fraud in digital retailing. While this is a more specialized use case, it's increasingly relevant as more of the car purchase process moves online. Another finance-related AI application: optimizing loan or lease offerings. AI can match customer profiles with the financing products they're most likely to be approved for (or interested in), improving the F\&I conversion rates.
- **Operational Efficiency and Other Areas:** AI can contribute to nearly every department. In parts departments, AI forecasting can ensure optimal stock levels of

[18] Five Ways AI Will Disrupt Car Dealerships in 2025

[19] Five Ways AI Will Disrupt Car Dealerships in 2025

high turnover parts. In HR, AI tools can help schedule shifts or even aid in hiring-screening resumes (though caution is needed to avoid bias in AI hiring tools). In accounting, AI can automate invoice processing or detect irregularities in expenses. Even compliance tasks, like monitoring camera feeds for safety (did the employee in the shop put the car on the lift correctly?) can leverage AI. The overarching theme is **automation of routine tasks** and surfacing of actionable insights from data, which in turn frees managers and employees to focus on strategic and interpersonal aspects of the business.

Benefit Highlights: Across these use cases, dealers stand to gain multiple benefits:

- *Efficiency and Cost Savings:* Automation reduces manual workload and errors. Tasks get done faster and often with fewer people-hours, translating to cost savings or capacity to handle more business.
- *Improved Decision-Making:* AI can digest far more data than a person and find patterns invisible to humans, leading to smarter decisions (what inventory to stock, which leads to pursue, how to price vehicles, etc.).
- *Personalization:* AI enables a more tailored experience for customers. Communications and recommendations can be customized, which tends to increase customer engagement and satisfaction.
- *24/7 Operation:* AI doesn't sleep – services like chatbots or monitoring systems work around the clock, meaning your dealership is "active" even when your employees have gone home.
- *Innovation and Competitiveness:* Early adopters of effective AI differentiate themselves. They can offer new services (like an AI-powered virtual showroom or voice-activated info kiosks) that set them apart from competitors. In a world where, as one industry observer put it, **"embrace AI or risk being left behind"**, the innovative dealership can capture tech-savvy customers and establish a forward-thinking brand image [20].

Of course, every technology implementation should be driven by a clear business need or problem to solve. As a dealer, before jumping on the AI bandwagon, ask: *What issue am I trying to address or which process do I want to improve?* Whether it's reducing customer wait time or improving used car acquisitions, keep the goal in focus. AI is a means to an end, not an end in itself. In the next sections, we'll explore the **risks involved** in using AI and how to govern these technologies so that they truly deliver on their promise without unintended consequences.

---

[20] Five Ways AI Will Disrupt Car Dealerships in 2025

Risks and Challenges of AI Adoption

While AI offers enticing benefits, it's crucial for dealerships to approach it with eyes wide open to the risks and challenges. Ignoring these could turn an AI initiative from a boon into a potential liability or PR disaster. Here we outline key risk areas and concerns, many of which underscore why strong AI governance (addressed in the next section) is so important:

- **Data Privacy and Security:** Dealerships deal with sensitive personal information (customer contact info, financial details for credit checks, driver's license scans for test drives, etc.). If an AI system mishandles this data, it can lead to privacy breaches. For instance, sending customer data to a third-party AI service (like a public cloud AI or an external vendor tool) *could violate privacy laws* if not done carefully. There's also the risk of **data leaks** – an AI could inadvertently expose private information (imagine a chatbot that, due to a bug, reveals one customer's data to another). Cybersecurity is another facet: AI tools could be new targets for hackers. A vulnerability in an AI service might open a backdoor into your network. According to industry analysis, the more an AI system knows about customers, the more vigilant a dealer must be in safeguarding that information [21]. In short, AI or not, the dealer remains responsible for protecting customer data. Failing to do so can result in regulatory penalties (fines under laws like California's Consumer Privacy Act or others) and loss of customer trust if a breach becomes public.
- **Regulatory and Compliance Risks:** The legal landscape around AI is evolving quickly. In the US, there isn't a single "AI law," but regulators are applying existing laws (consumer protection, anti-discrimination, data protection) to AI contexts. The Federal Trade Commission (FTC), for example, has warned companies about misleading or unfair use of AI – if an AI makes claims or decisions that can't be explained or that cause consumer harm, the FTC might consider it a deceptive practice. Specific to auto retail, the FTC's updated Safeguards Rule (under Gramm-Leach-Bliley Act) requires dealers to assess and oversee how customer data is protected by their service providers [22], which would include AI vendors. Some states are moving ahead with AI-targeted regulations; Colorado passed a law (effective 2026) requiring developers of high-risk AI to mitigate algorithmic discrimination [23]. Internationally, Europe's **EU AI Act** is introducing stringent requirements on AI systems, especially those in "high risk" domains (which could include creditworthiness evaluations or AI driving features). A dealer operating only in the U.S. might not feel Europe's rules directly, but if you use software from global

[21] Five Ways AI Will Disrupt Car Dealerships in 2025

[22] Risk Assessment Questionnaire – STAR

[23] Federal AI Mandates and Corporate Compliance: What's ... - Cogent

companies, those companies may change their products due to the EU Act, or if you deal with international customers, you could indirectly be affected. **Franchise considerations** are also relevant: automakers may issue their own AI policies for franchisees to follow (to ensure consistent brand experience and compliance across the network). Non-compliance with any applicable rule—be it a data privacy law, an AI transparency requirement, or OEM policy—poses legal and financial risks. Keeping track of and adapting to these regulations is a challenge that must be actively managed.

- **Ethical and Bias Issues:** AI systems learn from data, and data can reflect biases (or lack context). An AI might inadvertently exhibit discrimination – e.g., a loan recommendation algorithm might disfavor certain demographics if the training data had biases. Or a sales lead scoring AI might undervalue customers from a certain ZIP code due to historical sales data, not actual individual interest. These outcomes can be unethical and also illegal (violating equal opportunity laws). Another ethical issue is **transparency**: AI can be a "black box," making decisions that even developers find hard to explain. If a customer asks, "Why did your AI deny me this service offer?" The dealer should have some accountability or answers. Without transparency, it's hard to justify decisions or debug problems. There's also **misinformation risk**: generative AI might fabricate answers (known as *hallucinations* or confabulations). If, say, a chatbot confidently gives a customer incorrect information about a financing deal, that's a problem. Dealers must be mindful that AI doesn't have human judgment; it will do what it's told (or what it learned to do) even if that produces nonsensical or harmful results.
- **Lack of Control / Operational Risks:** Handing tasks to AI means you relinquish some direct control. If not properly configured, an AI might deviate from your intended business rules. For example, an automated pricing AI might start pricing too aggressively low, hurting margins, because it single-mindedly optimizes volume. Or an AI tool could go down or malfunction, disrupting an important process (imagine your appointment scheduling AI fails, and customers can't book service slots online, and you aren't aware until much later). Reliance on AI also raises **business continuity questions**: if your staff have become dependent on an AI and it fails, do you have a manual fallback plan? The flip side of "cost of doing nothing" is also **the cost of doing something poorly** – a misbehaving AI can create new headaches. Additionally, if a dealership implements an AI without proper vetting, it might inadvertently introduce **Shadow IT/AI** risk: employees might use unauthorized AI tools (perhaps to make their work easier) which haven't been approved by IT or legal. This can lead to data spilling into unmanaged platforms.
- **Integration and Skills Challenge:** Many AI solutions won't work in isolation; they need to interface with your other systems (DMS, CRM, inventory management, etc.). Integration can be complex and, if done incorrectly, can create security gaps or data errors. There's also the challenge of having staff with the right skills to manage and

interpret AI. If nobody on your team understands how the AI makes decisions, you could be flying blind or become overly dependent on vendor support. Training people to work alongside AI (for example, teaching sales staff how to interpret lead scores from an AI tool) is an often-underestimated challenge. And culturally, there can be resistance: some employees might mistrust or fear the AI (worrying it might replace them or doubting its suggestions). Change management is therefore part of the risk—ensuring your team buys into using AI as a helpful assistant, not an enemy.

- **Cost and ROI Uncertainty:** Implementing AI systems can be expensive, not just in licensing software, but in data preparation, training, and process changes. There's a risk of investing in an AI solution that doesn't pay off because it was overhyped or not the right fit. Without careful planning, a dealership could spend significant money on AI and get little value, perhaps because the data wasn't sufficient quality to train the model well, or staff didn't adopt the new system. Always consider conducting a pilot and measuring results before scaling AI deployment. And be wary of one-size-fits-all claims; the best AI solution for one dealership may not yield the same elsewhere if processes differ.

Consequences of these risks can range from minor inconveniences to major crises:

- Data/privacy mishaps can lead to lawsuits or fines as well as reputational damage (customers hearing that their dealer had a data breach).
- Regulatory violations could result in penalties or forced cessation of a practice (e.g., being ordered to stop using a non-compliant AI tool).
- Ethical lapses or AI errors could alienate customers or lead to bad publicity (imagine a news story: "Dealership's AI loan assistant disqualified minority buyers unfairly" – a nightmare scenario).
- Operational failures impact revenue and trust – if the AI that prices your inventory goes haywire for a week, that could mean lost profit or chaos in sales.

It might sound daunting, but these risks **can be managed**. The next section of this paper focuses on **governance and risk mitigation strategies** – essentially, how to harness AI's benefits while keeping these potential pitfalls in check. A well-devised AI governance plan addresses the above concerns proactively. It is also worth noting the *risk of inaction*: doing nothing about AI has its own cost, as competitors and customers move forward. In many ways, the biggest risk is being unprepared – either unprepared for adopting AI (and doing it badly) or unprepared for a world where others have adopted it (and you're left behind). As one industry analysis put it, *the cost of doing nothing in the face of AI advancement is itself a strategic risk* [24]. With that in mind, let's turn to how dealers can develop the proper guardrails and practices – in other words, AI governance.

[24] Five Ways AI Will Disrupt Car Dealerships in 2025

AI Governance and Risk Mitigation Strategies

AI governance refers to the **framework of policies, practices, and oversight mechanisms** that ensure AI is used responsibly, ethically, and in alignment with an organization's goals and obligations. For dealerships, implementing AI governance is about putting the "rules of the road" in place for your use of AI – much like you have standard operating procedures for sales or service processes. Governance doesn't have to be onerous; it's there to protect your business and customers while enabling innovation. Below, we outline key components of AI governance for dealerships, including both **technical controls and organizational (non-technical) measures**. Think of this as a checklist or playbook to strengthen AI usage in your dealership.

1. Strategy and Policy Framework

Begin with a clear strategy: **why and where is the dealership using AI?** Identify the problems or opportunities AI will address (e.g., "improve lead conversion by responding faster via chatbot" or "reduce service no-shows through predictive reminders"). This helps keep AI deployments goal oriented.

Next, develop an **AI Policy** or a set of guiding principles. This policy should define what AI means for your organization and outline acceptable uses. For instance, it might state that *"AI will be used to augment employee capabilities, not replace human decision-making in final outcomes without review"*. It should also cover data usage boundaries (e.g., *"Customer personal data will not be shared with AI providers without ensuring compliance with privacy requirements"*). Essentially, set the *guardrails*: what the AI can do, and what is off-limits. STAR recommends publishing an **"acceptable use matrix"** for AI at the dealership [25], clarifying which departments can use which types of AI tools and for what purposes, and where human sign-off is required.

In your policy, address issues like **data privacy, accuracy, transparency, and accountability**. For example, decide whether employees are allowed to use public generative AI tools (like ChatGPT) for work. Some organizations ban input of any sensitive data into such tools; others allow use with caution. If allowed, provide training so staff know how to use it safely (more on training later).

It's also wise to require a **validation step for AI outputs** in critical processes. For instance, if an AI writes responses to customer emails, you might require a human to quickly review them before sending until the AI has proven highly reliable. Make it company culture that AI suggestions are helpful but *not infallible*. Having a written policy and getting leadership buy-

[25] AIGovWG_Running_Documentation

in (and ideally, communicating a summary to all employees) sets a tone from the top that AI is to be used thoughtfully.

Finally, include in your strategy a consideration of the **business case** for each AI application. What ROI or benefit do you expect? How will you measure success? This ties into governance because it prevents chasing AI hype without clear purpose. Every AI project should have success metrics (e.g., reduce response time by X%, increase upsells by Y%). Periodically review if those are being met. If not, reassess or pivot; governance includes the decision to turn off an AI tool that isn't delivering or is causing problems.

2. Data Management and Privacy

Since data is the fuel for AI, govern your data diligently. Good **data governance** practices should parallel AI governance. Dealerships should know *what data is being used by the AI, where that data comes from, and where it flows*. Key steps include:

- **Inventory and classify data:** Identify which data sources your AI systems draw from (CRM databases, DMS records, customer web forms, etc.). Classify data by sensitivity. For example, distinguish public data (vehicle specs), internal data (sales figures), and sensitive personal data (customer contact, financial info). AI models trained on or processing personal data should be flagged for special oversight, given privacy implications.
- **Permissions and consent:** Ensure you have the right to use the data for AI analysis. If you plan to use customer data in an AI tool, check privacy policies and any consent language. Customers might need to be informed if, say, their chat with you is analyzed by an AI. If unsure, consult legal counsel. Also, if you are combining your first-party data with third-party data (e.g., buying some marketing lists or using industry datasets for training an AI model), be clear on the usage rights of that third-party data.
- **Data hygiene and quality:** The old saying "garbage in, garbage out" applies strongly to AI. Clean, accurate data leads to better AI outcomes. Design processes to regularly clean up data (deduplicate customer records, correct errors, update outdated info). The dealership should **"understand data hygiene"** as a discipline: who owns each data set, how often it's updated, and how errors are rectified [26]. For AI specifically, if you're training any models (perhaps a custom model using your own dealership data), spend adequate effort on preparing the training data and removing biased or irrelevant records.
- **Limit data sharing:** A big risk is sending confidential dealership or customer data to an external AI service without protection. If you use cloud AI providers, prefer those that allow **data isolation** (so your data isn't used to train others' models) and have

[26] AIGovWG_Running_Documentation

robust security certifications. Whenever possible, **avoid using production private data in public AI tools**. If you want to experiment with a public generative AI (like asking it to rephrase an email), at least anonymize or alter any personal details you include. Even better, explore options to use AI services that can run on your own systems or are provided by vendors who contractually assure data privacy. Establish controls: for instance, block employees from pasting sensitive info into external AI websites or set up approved internal AI tools as alternatives.

- **Retention and ownership:** Clarify who owns the outputs of AI (usually the dealership but check vendor terms). Also, if an AI process generates new data (like recordings of chatbot conversations, or analysis files), manage those outputs securely and have retention policies. Don't keep data longer than needed, especially personal data, unless value justifies it – this reduces exposure.
- **Monitor data leakage:** As part of security (next section), monitor network logs for any unusual data exfiltration that could be AI tools sending data where they shouldn't. Some companies implement DLP (Data Loss Prevention) software to catch sensitive info leaving the network via unapproved channels.

In summary, treat data as an asset with rules. By **documenting data flows** and decisions (e.g., "we will use customer email addresses in an AI-powered email campaign, but we will not send any credit score info to the AI service"), you create a reference that both technical teams and management can understand. Given that many AI risks tie back to data misuse, having strong data governance is a foundational mitigation.

3. Risk Assessment and Security Controls

Incorporate AI into your existing risk management and IT security processes. For each AI application, ask **"What's the worst that could happen if this AI misbehaves or is compromised?"** Then plan controls accordingly. Some key considerations:

- **AI Risk Assessment:** Perform a targeted risk assessment for AI systems, much like you would for a new software system. Identify potential failure modes (e.g., erroneous output, downtime, unauthorized access, biased decisions) and who/what could be harmed by them (customers, the business, compliance...). The STAR AI Working Group suggests building and maintaining **risk management protocols related to AI use** [27]. This could mean adapting a standard IT risk assessment template to include AI-specific questions (like: Is there a risk of model bias? Could inputs be manipulated (prompt injection attacks) to yield bad behavior? Are outputs reviewed for accuracy before action?
- **Testing and Validation:** Before fully deploying an AI solution, test it. If it's a vendor product, ask about their testing. You might run the AI in parallel with human

[27] AIGovWG_Running_Documentation

performers to compare results for a period of time. Also, consider "red-teaming" the AI – intentionally input tricky or bad data to see how it responds. For example, what does your sales chatbot do if a user starts asking inappropriate questions or provides clearly false information? Does it handle it gracefully? Testing helps catch issues early. Also plan how you will continuously test or audit the AI over time. AI models can "drift" as data changes, so periodic re-validation is wise.

- **Security of AI Systems:** If an AI application has access to your systems, treat it like any other privileged system. Apply **least privilege** principle – give the AI or its service account access only to the data and functions it absolutely needs. For instance, if an AI scheduling assistant needs to read your calendar to book appointments, it should not also have full CRM database access unless necessary. Use strong authentication and API keys for integrations. Additionally, protect any AI model or API endpoints you expose (some AI might run on premises; ensure it's behind proper authentication). Regularly patch AI software and libraries; many AI frameworks are open-source and can have vulnerabilities like any software.
- **Monitoring and Incident Response:** Put monitoring in place to track the AI's actions and performance. If you have an AI agent making changes (like adjusting prices or sending communications), enable logging such that every action it takes is recorded (what it did, based on what input, and when). This will be invaluable if you need to investigate a problem. Define clear steps for an **incident response** involving AI. For example, if the AI outputs a glaringly wrong recommendation that reaches a customer, how should staff respond? Or if the AI system is breached by a cyberattack, how do you contain it? Make sure your general incident response plan covers AI scenarios. Also, consider setting up alerts for unusual AI behavior – e.g., if the volume of messages the AI sends in an hour exceeds a threshold, perhaps trigger an alert to IT.
- **Vendor and Third-Party Risk:** Many dealerships will acquire AI capabilities via vendors (software providers, DMS add-ons, etc.) rather than building everything in-house. It's essential to vet these vendors for security and compliance. Request and review their **security certifications or assessments** (SOC 2 reports, ISO 27001, etc.). The STAR organization provides a **Risk Assessment Questionnaire** exactly for evaluating service providers' data security and privacy controls [28] [29]. Dealers can use such standardized questionnaires to ask AI vendors about how they protect data, whether they use sub-processors, how they handle breaches, etc. Also check the vendor's terms of service: ensure they aren't claiming broad rights to use your data beyond providing the service. Ideally, include contractual clauses about data use, privacy, and AI ethics (for instance, you might require that the vendor's AI model is

[28] Risk Assessment Questionnaire – STAR

[29] Risk Assessment Questionnaire – STAR

not trained on any data that would be unlawful or that they will assist you in responding to any regulatory inquiries about the AI). In short, **choose trusted partners** and don't hesitate to ask tough questions. If a vendor can't explain how their AI works at a high level or refuses to share anything about their model's training or accuracy, that's a red flag.

- **Shadow AI Prevention:** As mentioned under risks, employees might experiment with AI tools outside official channels (especially tech-savvy staff trying to be helpful). To mitigate this, provide approved (and preferably easy-to-use) AI tools so that employees aren't tempted to use unsanctioned ones. Educate them about the dangers of unvetted apps. IT departments can update policies to block known risky AI web apps or browser plugins. It's about balancing **encouraging innovation** (letting power users prototype useful AI solutions) with oversight (ensuring IT knows and can support those solutions if they prove helpful). A governance program might set up an intake process: if an employee finds a cool AI tool, they can submit it for review rather than using it quietly under the radar.

In essence, treat AI systems with the same seriousness as you treat critical IT systems – because they will be making or influencing decisions. By systematically assessing and controlling risks, you can prevent most foreseeable issues or at least catch them quickly before they escalate.

4. Legal and Regulatory Compliance

Given the dynamic regulatory environment around AI, it's important for dealers to stay informed and compliant. Some best practices:

- **Stay Updated on Laws:** Assign someone (or a team) the responsibility of monitoring AI-related legal developments. This could be part of the role of a compliance officer or the IT director or even an external legal advisor. Keep an eye on federal agency guidance (FTC, NHTSA, CFPB if AI is used in financing, etc.), state laws, and international rules if applicable. The **state of AI regulation is in flux** [30], so today's best practice might need updating next year. A practical approach is to periodically (say, twice a year) review any new laws or guidance. Resources like NADA's regulatory alerts, industry webinars, or partnering with counsel who follows tech law can be useful.
- **Consult Legal Early:** When implementing a new AI tool, run the idea by legal counsel, especially if it touches customer data or could impact customers' decisions. For example, using an AI to make lending decisions may invoke fair lending laws – definitely a case for legal input. Or using biometric AI (like facial recognition for customer check-in) could trigger state biometric privacy laws (like Illinois' BIPA).

[30] AIGovWG_Running_Documentation

Legal review isn't to create roadblocks but to ensure you deploy AI in a compliant way (or add the right disclosures). It's better to bake compliance in from the start than to retrofit it after a regulator raises a concern.

- **Terms of Service and Jurisdiction:** Pay attention to the country of origin for AI software. Using tools from outside the U.S. might subject data to foreign jurisdictions. For instance, if you use an AI service based in another country, could the data be accessed by that foreign government under their laws? One recommendation is to **scrutinize terms of service/use** for any AI product [31]. They may contain clauses about data usage (some AI providers say they can use your inputs to improve their models – which might be unacceptable if inputs include customer data). Also, be mindful of where data is stored (seek vendors who offer data residency in U.S. or in countries with strong data protection). The notion of "consequences of using technologies that originate in other countries" [32] basically reminds dealers that if something goes wrong with that tool, *you* as the dealer still hold the bag for any fallout. So extra caution is merited with foreign or untested vendors.
- **Frameworks and Standards:** Aligning with industry-standard frameworks can help demonstrate compliance due diligence. For example, the new **ISO/IEC 42001:2023** is an international standard for AI management systems, akin to an ISO quality standard but specifically for AI governance. While a dealer might not get certified in ISO 42001 (that's more likely for large organizations or AI developers), knowing its existence can guide best practices. More accessible is the **NIST AI Risk Management Framework (RMF)** [33] which provides a voluntary set of guidelines to ensure trustworthy AI (covering principles like transparency, fairness, accountability). Dealerships could adopt relevant parts of NIST's framework as a checklist (e.g., do we have a process to map out AI risks? Are we measuring and monitoring our AI's performance? etc.). Doing so not only improves your governance but could be a positive signal to regulators or business partners that you take responsible AI seriously.
- **Privacy and Consent Notices:** Ensure your privacy policy (the one you present to customers, often required by law if you collect personal info) is updated to mention AI usage if applicable. For example, if customer chats may be recorded and analyzed by AI, say so. If you use AI to make decisions (like credit offers or deal recommendations), consider a brief notice like "automated decision-making technology may be used." Under some laws (like EU's GDPR or possibly upcoming

[31] AIGovWG_Running_Documentation

[32] AIGovWG_Running_Documentation

[33] Federal AI Mandates and Corporate Compliance: What's ... - Cogent

U.S. laws), individuals have rights when significant decisions are automated – including being informed and sometimes opting out or requesting human review. While U.S. law is not there yet broadly, California's privacy laws, for instance, have begun including the concept of automated decision transparency. It's good practice (and simply honest) to let customers know how their data is being used, especially if AI is involved in a way they might not expect.

- **Liability and Insurance:** Check with your insurance provider whether use of AI has any impact on your policies or if there are specific coverages you should have. For example, if you deploy a relatively autonomous AI and it causes damage (financial or even physical, in some hypothetical case of an AI instructing something incorrectly), does your general liability or cyber insurance cover that? Some insurers are starting to offer specific riders or policies for AI-related risks. Also, if your AI usage is significant, having legal counsel and/or an **AI governance committee** can help show you took reasonable precautions, which could mitigate liability in the event of a lawsuit.
- **Recordkeeping:** Maintain documentation of your AI governance efforts – policies you've set, training you've done, vendor due diligence records, etc. If a regulator ever questions your AI use, having that documentation can demonstrate that you are managing AI responsibly (which might reduce penalties or even prevent enforcement). For example, if an issue arises from an AI, you can show, "We followed these steps, we identified this risk, here's how we tried to prevent it, and here's how we responded once we found the issue."

In short, treat compliance as a continuous process in the AI journey. Your aim is not only to follow the letter of existing law but to uphold the spirit of using AI in a fair, transparent way that respects customers.

One proactive step is to have **legal or audit expertise on-call** – possibly via a retainer – specifically for privacy and AI matters [34]. Small dealerships might lean on industry associations for this; larger ones might have in-house counsel get up to speed on AI. The investment in guidance is likely far cheaper than fighting a legal battle or regulatory fine later.

5. Training and Awareness

People are at the center of any successful technology adoption, and AI is no different. Proper training and cultivating an informed culture are critical governance measures:

- **Employee Training on AI:** Ensure that all staff who use AI tools (or whose work will be affected by them) receive training. This training should include basic concepts of how the AI works (they don't need to know the algorithms, but should know, for

[34] AIGovWG_Running_Documentation

example, “this lead scoring tool analyzes past sales data and assigns a score of 0-100; higher means more likely to buy”). Emphasize the known limitations of the AI. If, for instance, the AI sometimes confabulates (makes up answers), show examples so employees recognize it. Teach them the company’s policies on AI use (from step 1 above). For example, if your policy says, “do not paste customer social security numbers into any AI tool,” make sure that’s clear in training with rationale. Training can be in-person workshops or online modules, and should be updated as tools evolve or new ones come in. **Training is not one-time**; consider refresher sessions, especially if a notable incident happened elsewhere that you can learn from.

- **Specialized Skill Development:** For roles heavily involved with AI (like a marketing manager using an AI analytics dashboard, or a service manager using an AI diagnostic tool), provide deeper training. This might involve the vendor providing a webinar or hands-on demo for your team. The goal is that users become not just competent but confident in using AI and in interpreting its outputs. They should know how to override or correct the AI when needed, and how to feed improvements (e.g., flagging when the AI gave a bad suggestion, if the system allows feedback). You might even identify internal “AI champions”- individuals with an interest in the tech who can help others and liaise with leadership on AI performance and issues.
- **Ethics and Scenario Planning:** Incorporate ethical reasoning into training. Pose scenarios to staff: “If our sales chatbot is asked a tricky question it can’t handle, what should you do?” or “If the AI recommends denying a loan to a customer that you feel is qualified, how do you double-check its reasoning?” Encourage a mindset of **human-AI collaboration**, where employees know they are the ultimate decision-makers. The AI is a tool, not the boss. Some companies adopt a slogan like “AI + Human > Human or AI alone” to stress teamwork between employees and AI.
- **Customer-Facing Transparency:** Train customer-facing staff (salespeople, BDC, service advisors) on how to communicate about AI to customers. For example, if a customer asks “How was this trade-in offer determined?”, a salesperson should be able to explain in plain language if an AI appraisal tool was used (“We use a system that compares thousands of data points – like recent auction prices for similar models and your car’s condition from photos – to come up with a fair offer, and then we review it to make sure it makes sense with the local market.”). Being transparent and confident in explaining AI’s role can build trust. Conversely, if employees are blindsided by questions about your AI usage, it could undermine credibility.
- **Leadership and Culture:** Train your management and leadership as well. They need to understand the strategic importance and limitations of AI to support it appropriately. Leaders set the tone: if they openly endorse the responsible use of AI and allocate time and resources for training and oversight, employees will take it seriously. Consider establishing an internal **AI governance committee or an “AI**

Lead" role [35]. This could be someone who coordinates AI efforts, keeps track of all AI projects, and ensures cross-department learnings. It doesn't have to be a new hire; it could be a tech-savvy manager who is given this responsibility as part of their role.

- **Knowledge Sharing:** Encourage staff to share experiences with AI tools – what works, what issues they've faced. Perhaps periodic meetings or an internal chat channel on AI usage where tips can be exchanged (and the governance team can chime in reminders if needed). As AI in dealerships is relatively new, fostering an environment of continuous learning will help you adapt your practices effectively. One department might discover a risk or a trick that others can benefit from if communicated.

Ultimately, **well-trained employees act as human safeguards and enhancement for AI**. They can catch AI's mistakes, provide the empathy and judgment AI lacks, and help tune the system to be more effective. Training also helps reduce fear – when people understand the tool, they're less likely to distrust or misuse it. A dealership that is educated about AI at all levels will handle the technology more deftly, reaping benefits while avoiding pitfalls.

6. Ongoing Monitoring and Continuous Improvement

Governance is not a "set and forget" exercise. As you deploy AI, continuously monitor how it's performing and be ready to improve processes:

- **Performance Metrics:** Define KPIs for your AI initiatives (as mentioned in the strategy). Monitor these regularly. If a chatbot was supposed to increase lead capture by 20%, is it doing so? If not, why? Perhaps customers are asking questions it can't answer – that insight might inform you to expand its knowledge base or route certain queries to humans faster. Use dashboards or reports from your AI tools to keep an eye on usage trends, success rates, error rates, etc. This is akin to how you'd monitor any business process; the presence of AI doesn't change the need for good old performance management.
- **Feedback Loops:** Establish a way for employees and customers to give feedback on AI-driven interactions. For example, after a chatbot interaction, you might ask the user "Was this helpful? \[Yes/No]". Internally, if salespeople notice the lead scoring AI is consistently missing obvious hot leads or mis-prioritizing, they should have a channel to report that. Use this feedback to refine the system or engage with the vendor for improvements. AI models might need retraining or adjusting parameters, which your vendor can often do if you provide them with concrete examples of where the AI went wrong.

[35] AIGovWG_Running_Documentation

- **Governance Reviews:** Periodically (perhaps quarterly or semi-annually), the cross-functional team or AI lead should review the overall AI governance status. This might include:
    - Ensuring all new AI projects went through the approval checklist.
    - Reviewing any incidents or near-misses (e.g., "We found that the service bot was double-booking appointments. We fixed that bug – how do we prevent similar issues?").
    - Updating risk assessments with any new information.
    - Check compliance against any newly effective laws or standards.
    - Discuss upcoming AI opportunities or changes in vendor landscape.

These governance reviews can be relatively short meetings if things are smooth, but they ensure continuous attention.

- **Adaptation:** Be prepared to adapt your governance documents (policies, procedures) as you learn. Perhaps you discover a need for a new guideline ("We should explicitly forbid using AI to communicate warranty terms without legal review"). Add that to the policy. Or maybe employees found the acceptable use matrix confusing – clarify it. Governance should evolve with both external changes and internal lessons learned.
- **Auditability:** If possible, keep records of decisions made by AI (this was mentioned under logging). This can aid in internal audits. For example, a dealer might audit a sample of AI-made pricing decisions against a human manager's decisions to see if they align with strategy. If you find anomalies, address them. Also, if regulators or OEMs ask questions, you have an audit trail to show what your AI did and that you were monitoring it.
- **Leverage External Resources:** Many industry groups and alliances are creating AI governance resources (for instance, NADA or Alliance for Automotive Innovation may share best practices, or general resources like Partnership on AI toolkit). Stay connected with industry forums to learn what peers are doing. Often, risks or solutions one dealership encounters will be similar for another.
- **Sunset and Refresh:** Occasionally, decide if an AI tool should be phased out or replaced. Maybe a new version or competitor offers better accuracy or features that align with updated compliance needs. Don't let inertia keep outdated AI in place. Conversely, if an AI tool has proven extremely useful and trusted, you could consider expanding its scope in a controlled way (with new risk assessment, of course).

A central theme here is **iterative improvement**. Start small, govern tightly, learn and expand. This approach ensures that the AI journey is smooth and yields positive results. Dealers who follow a disciplined governance approach will find they can innovate with AI faster in the long run—because each success builds confidence and a track record, making it

easier to justify the next AI project, and each governance practice becomes a standard part of doing business.

Navigating the Regulatory Landscape

(This section provides an overview of relevant regulations and standards affecting AI use, as a reference for dealers.)

The regulatory environment for AI is rapidly evolving, and while auto dealers may not be the primary target of many AI laws, they are certainly impacted by broader data and consumer protection regulations. Here's a snapshot of what to be aware of:

- **Federal Oversight in the U.S.:** Currently, there is no single federal law in the U.S. that exclusively governs AI. Instead, existing laws and agencies fill the gap. The **Federal Trade Commission (FTC)** plays a key role by enforcing consumer protection and fair competition laws on AI implementations. The FTC has made it clear that it expects businesses to use AI truthfully and fairly – for instance, not to make bogus claims about AI capabilities or to allow biased or discriminatory outcomes. The FTC's concept of "truth, fairness, and equity" in algorithms means dealers should avoid AI that could deceive consumers or treat them unfairly (like differential pricing that might be seen as discriminatory). The **Consumer Financial Protection Bureau (CFPB)** also watches AI in lending decisions (if you use any AI in credit or F\&I products, ensure it complies with fair lending rules). Meanwhile, the **National Highway Traffic Safety Administration (NHTSA)** is concerned with AI in vehicles (e.g., autonomous driving systems), which is more on the automaker side, but if dealers deal in vehicles with AI features, staying informed on those guidelines is wise (for customer safety education and liability reasons).
- **State Laws and Privacy Regulations:** Many states have data privacy laws (California's CCPA/CPRA, Virginia's CDPA, Colorado's Privacy Act, etc.) that, while not AI-specific, regulate personal data use. Using AI on customer data means you must adhere to these laws' requirements like providing disclosure of data usage, honoring opt-outs (e.g., if a customer opts out of data selling/sharing, you shouldn't be feeding their data into an external AI model that constitutes "sharing"). Notably, **California CPRA (2023)** introduced a concept of Automated Decision making and gives consumers rights to access and opt-out of certain automated decision processes – regulations on this are still being fleshed out. Also, as cited earlier, **Colorado passed an AI-specific law** (for AI developers) mandating care to avoid discriminatory AI outcomes [36]. While that targets creators of AI more than users, dealers should track such developments because they signal a direction of regulation (and if you build or

[36] Federal AI Mandates and Corporate Compliance: What's ... - Cogent

heavily customize your AI, you might fall under such requirements). Additionally, state laws like Illinois' **BIPA** (biometric privacy) could become relevant if using AI for facial recognition or fingerprint ID in any dealership process (e.g., some service kiosks use driver's license scanning). The safest approach is to treat any biometric or highly personal data with explicit consent and robust protection or avoid using such features unless truly adding value.

- **International Influence – The EU AI Act:** The European Union's **AI Act** (adopted in 2024, with phased implementation starting 2025-2026) is the world's first comprehensive AI law. It categorizes AI systems by risk level and imposes strict requirements on high-risk AI (which includes systems in areas like credit scoring, employment decisions, and safety components of products). If a dealer software vendor operates globally, they might adjust their product to comply with the EU AI Act by building in more transparency or documentation, which could benefit you too. If you have any operations or customers in Europe (unlikely for most U.S. franchise dealers, but possible for high-end collectible sales or online sales overseas), you'd need to consider compliance. The EU AI Act also bans certain AI practices outright (like social scoring of individuals by governments, or AI that manipulates human behavior in harmful ways). Philosophically, it's influencing global norms – emphasizing the need for risk assessments, logging, human oversight for high-risk uses, etc. U.S. businesses are watching it closely. It's a reminder that robust **documentation and transparency** in AI (knowing what data went in, how the model was designed, what it's intended to do) is a good practice that may soon be expected everywhere.
- **Industry Standards and Guidelines:** Outside of laws, there are industry efforts to self-regulate. We mentioned **NIST AI Risk Management Framework** – not a law, but a respected guideline [37]. Adhering to it can prepare you for future compliance and just generally make your AI initiatives safer. It encourages practices like mapping AI's purpose and risk, measuring and monitoring its performance, and having governance structures ("Govern" function in the framework) – many of which we've discussed in this paper. Automotive industry groups might also produce AI guidelines (for example, the Alliance for Automotive Innovation might release dealer-focused recommendations). **STAR's own standards** (through this working group and related efforts) aim to help dealers and vendors align on best practices – for instance, STAR has standardized risk assessment questionnaires (as referenced earlier) and is exploring guidelines specific to AI in retail automotive.
- **Franchise Agreements and OEM Policies:** Don't forget to check if your OEM (the vehicle manufacturer) has any policies on dealership use of certain technologies. As OEMs roll out connected vehicle services and digital retailing platforms, they may

[37] Federal AI Mandates and Corporate Compliance: What's ... - Cogent

have rules about how customer data is used or what add-on software (including AI tools) dealers can integrate with those systems. For example, if an OEM provides an official virtual assistant in the car, they might discourage dealers from implementing a competing voice assistant that could confuse customers. Also, if the OEM's systems (like a lead management system) incorporate AI, they might pass down compliance requirements to dealers or require certain disclaimers be given to customers. It's part of governance to ensure any local AI initiative doesn't inadvertently conflict with your franchise obligations.

- **Future Outlook:** U.S. lawmakers are actively discussing AI regulation, so in the next few years we may see more concrete rules. Possible areas are requirements for transparency (like labeling when content is AI-generated), biases audits for AI used in important decisions, and safety certifications for AI in products (cars increasingly, for autonomous tech). Keeping a proactive stance—documenting your AI processes, insisting on ethical AI from vendors, and being transparent with consumers—will put you in a good position to comply with any new mandates. It's much harder to bolt on governance later under regulatory pressure than to have it from the start.

In summary, dealers should **approach compliance as both a necessity and an evolving target**. By combining common-sense ethical practices with awareness of legal trends, you can avoid trouble and even use compliance as a selling point ("Our dealership respects your data – here's how we use AI responsibly..." could be a trust message in marketing someday).

For further reading, several resources offer detailed guidance on AI governance and compliance: IBM's "What is AI Governance" FAQ, the Partnership on AI's toolkit for navigating AI norms, and NADA's legal advisory memos are great places to start for those wanting to dive deeper.

Conclusion: Embrace AI with Responsible Innovation

AI technology presents a transformative opportunity for franchised automobile dealers. From automating routine tasks and uncovering business insights to elevating customer experience, **AI can be a gamechanger in automotive retail**. Dealerships that thoughtfully integrate AI will likely find efficiencies and competitive advantages that translate into real ROI – as evidenced by industry surveys reporting revenue lifts of 20% or more from AI initiatives [38]. Moreover, adopting AI can future-proof dealerships against the growing digital expectations of consumers who increasingly demand the kind of personalized, seamless interactions that AI can help deliver.

---

[38] Survey Reveals Expected Surge in AI Budget Allocations Among Car ...

However, **successful AI adoption is not just about technology – it's about trust, oversight, and strategy**. This paper has highlighted that while the upside of AI is significant, the risks are very real if left unmanaged. Issues of data privacy, fairness, transparency, and security cannot be an afterthought. They must be addressed up front through robust governance. Dealers should take comfort that they don't need to be AI experts to manage these risks; by applying the governance principles outlined – clear policies, data controls, risk assessments, compliance checks, and staff training – even smaller dealerships can innovate safely. Resources from industry bodies like STAR, NADA, and others can provide templates and support so you're not starting from scratch.

One recurring theme is the **"human in the loop"** – AI works best with humans guiding it and reviewing it. Your dealership's expertise and customer understanding, combined with AI's speed and analytical power, is a winning combination. Employees should be empowered by AI, not replaced, and in turn they must be vigilant stewards of the AI tools they use.

As you plan your AI roadmap, start with pilot projects that are high value but manageable in scope. Ensure you have measurement and governance around those. Learn from them, then scale up. Celebrate quick wins (like the chatbot that boosted appointment bookings, or the inventory algorithm that reduced lot age by 10%) but also scrutinize any missteps as learning opportunities. Make AI governance part of your dealership's continuous improvement culture.

Finally, consider designating an **AI champion or governance lead** internally to keep momentum. This person (or committee) can coordinate efforts, keeping everyone aligned and informed. They can also serve as the point of contact with external entities – whether coordinating with an OEM's digital team, answering customer questions about your AI usage, or participating in industry forums to stay ahead of trends.

In conclusion, **the cost of doing nothing with AI is growing**. Dealerships that remain idle risk inefficiencies, higher costs, and dwindling relevance as the industry evolves. Meanwhile, those that embrace AI thoughtfully can achieve more efficient operations, better customer satisfaction, and new revenue opportunities. As one industry magazine aptly noted, *"by 2025, the most successful automotive retailers will be the ones that balance AI-powered efficiency with old-fashioned human know-how"* [39]. It's not about choosing man or machine but leveraging the best of both.

STAR, as the publisher of this guide, encourages all dealers to view AI not with fear, but with **pragmatic optimism**. With proper governance, AI can be a trusted partner in your business, not a wild card. By laying the groundwork now—educating yourself and your staff, putting

[39] Five Ways AI Will Disrupt Car Dealerships in 2025

policies in place, and starting small—you set the stage for sustainable growth and innovation.

In the race towards an AI-enhanced future, dealers who combine innovation with responsibility will earn the trust of their customers and the respect of their peers. This balanced approach will ensure that AI becomes a pillar of success in your dealership's strategy for years to come. Embrace the journey, stay informed, and let governance be your guide rail as you accelerate into the future of automotive retail.