

# AI-Driven Infotainment: Advancing Contextual and Personalized Automotive Systems

Ronak Indrasinh Kosamia

Atlanta, GA  
rkosamia0676@ucumberlands.edu  
0009-0004-4997-4225

## Abstract:

In-vehicle infotainment has traditionally delivered static user interfaces and limited adaptability, despite the automotive industry's broader push toward connected and semi-autonomous systems. Recent advances in artificial intelligence (AI) and machine learning (ML) suggest that infotainment can evolve into a contextually aware ecosystem—adjusting display layouts, anticipating user needs, and coordinating multi-regional or multi-modal features. This paper outlines a framework for AI-driven infotainment that unifies occupant classification, environmental triggers, and cloud-based analytics to provide real-time personalization. We focus on occupant-centric gating of features to reduce driver distraction, integrate predictive maintenance alerts at opportune moments, and exploit partial offline caching for robust operation in connectivity-limited regions. Preliminary evidence, including pilot user tests, indicates that occupant-based UI adaptation can bolster user acceptance while safeguarding against information overload. The approach also highlights potential synergy with e-commerce microservices, advanced route planning, and region-specific customizations. By bridging occupant recognition, environment variables, and learning-based modules, the proposed system underscores how future automotive infotainment can deliver higher levels of convenience, safety, and global scalability.

**Keywords:** AI-driven infotainment, automotive systems, occupant classification, predictive maintenance, driver distraction, connected vehicles, microservices, offline caching, machine learning, multi-regional deployment.

## I. INTRODUCTION

### A. *Motivation and Scope*

Recent years have seen automakers integrate data-driven modules in behind-the-scenes telematics or over-the-air (OTA) update mechanisms, but occupant-facing features remain comparatively static [4]. Meanwhile, user expectations for personalized digital experiences—shaped by smartphones and streaming services—push automotive solutions toward real-time adaptation. Consider a scenario in which the driver seat occupant is recognized with high confidence: the infotainment system can automatically minimize extraneous media suggestions to reduce driver distraction, highlight route guidance, or present fuel or battery status if driving an electric vehicle (EV). By contrast, a verified passenger might see a broader array of streaming content or e-commerce prompts, effectively gating features that are unsuited for active drivers [2] [5].

From a global standpoint, major OEMs face additional constraints, such as multi-lingual content, varied regulatory limits, and inconsistent connectivity across regions [6]. Infotainment systems that unify occupant-based gating with environment triggers can adapt to local commerce partners, route data, or partial offline usage. For instance, if occupant classification deduces a “touring passenger” profile in a connectivity-limited region, the system might rely on cached maps or local promotions, deferring real-time synergy until a stable network reappears [7]. Bridging occupant logic, environment data, and microservices for e-commerce or

predictive maintenance demands a well-structured aggregator pipeline and advanced concurrency control. By shaping infotainment around occupant needs, the system can optimize when and how to present new features without overwhelming drivers or forcing extensive user setup.

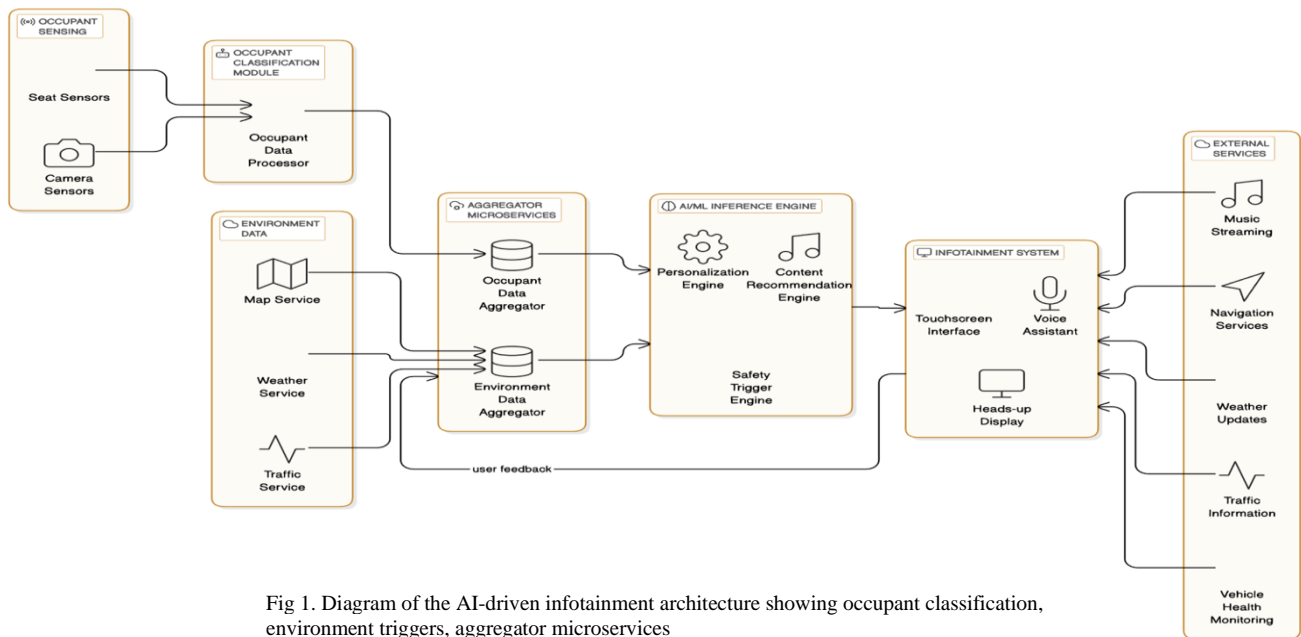


Fig 1. Diagram of the AI-driven infotainment architecture showing occupant classification, environment triggers, aggregator microservices

### B. Evolving Role of Infotainment in Connected Vehicles

While conventional infotainment units revolve around static home screens for navigation, phone, or media, next-generation systems aim for continuous context adaptation [3][8]. On the hardware side, head units commonly rely on embedded SoCs with limited GPU capabilities, historically insufficient for large-scale neural inference. However, partial or compressed ML models can now run occupant classification or local speech recognition with modest overhead, especially if model re-training occurs in the cloud [1]. This approach opens the door to occupant-based personalization—like seat or climate presets, user preference continuity, or occupant-based commerce tie-ins—updated nightly or periodically.

Simultaneously, telematics and backend aggregator services have progressed to handle predictive maintenance or real-time route adjustments [5]. By merging occupant detection with these aggregator-based modules, the infotainment system can time alerts or promotional pop-ups to coincide with occupant states that minimize distraction, such as displaying battery health updates only when the occupant is parked or recognized as a passenger. This occupant-based gating resonates with recent guidelines on driver distraction, in which dynamic elements must not intrude if the occupant is actively driving at moderate/high speeds [9,10]. The synergy also fosters environment-aware triggers; for instance, if occupant classification detects a new occupant or a shift from a solitary driver to multiple passengers, the system might reconfigure the interface to highlight group media or route suggestions that incorporate scenic stops.

### C. Linking Occupant Context with Predictive Maintenance

Predictive maintenance in consumer vehicles historically relied on simplistic triggers (e.g., mileage-based service intervals, static engine lights). Advanced ML-based approaches can track engine performance, battery usage, and driver habits to forecast likely component failures [11]. If occupant classification identifies a stable driver occupant during low-speed or parked scenarios, the system can discreetly present maintenance forecasts, booking suggestions, or upcoming part replacements. This measured approach addresses user annoyance or safety concerns. Indeed, occupant-based strategies can enable short maintenance status pop-ups only if occupant remains idle for ~30 seconds—reducing the risk of forcing a driver to read complex alerts mid-traffic [4][10]. Over time, aggregator data merges occupant usage logs with aggregated maintenance results, refining the predictive models. Such occupant synergy was rarely integrated in older generation infotainment, indicating a fresh frontier for user acceptance and brand differentiation.

#### D. *The Global Lens: Regional Localization and Offline Challenges*

The push to unify codebases for multiple regional markets confronts challenges around different language packs, local commerce networks, and regulatory constraints [6]. AI-driven occupant classification might further complicate matters if local laws vary on camera-based occupant detection or user data storage. By focusing occupant recognition on ephemeral seat sensor data or hashed camera embeddings, the proposed system can remain compliance-friendly while still personalizing features [12]. For example, occupant detection can direct local-lingual content if occupant regularly sets that preference, or show region-specific commerce if occupant is recognized as a passenger with historically positive interactions. Meanwhile, partial offline caching ensures occupant-based triggers continue functioning, even if a rural region's connectivity falters [1]. The aggregator microservices might queue new occupant model updates or region expansions, applying them once reconnected.

Table I. Comparing occupant-based gating scenarios across different occupant states, speeds, environment triggers, and recommended UI actions.

Scenario	Occupant State	Vehicle Speed	Env Trigger	Recommended UI Action
In-Car Commerce at Highway Speeds	Driver Confirmed	65 mph	None/Normal Weather	Lock commerce store. Show disclaimer: "Feature disabled above 0 mph."
Attempting Navigation Reroute in Heavy Traffic	Driver Uncertain	30 mph	Traffic Jam	Show partial nav instructions; require occupant re-confirmation (voice or seat sensor) if driver is unrecognized.
Child Detected with High Cabin Temperature	Child Seat Occupant	0 mph (parked)	Internal Temp > 90°F	Force climate adjustment for safety. Prominent alert: "High cabin temperature detected. Cooling system engaged."
Offline Mode with Unknown Occupant	Unknown	Varies	No Internet	Display minimal UI. Offer disclaimers: "Offline Mode: occupant identity limited, certain features unavailable."
Payment for Parking Garage Entrance	Driver Identified	5 mph	Parking Facility	Autoprompt payment if occupant is recognized and account on file; otherwise, block or require manual input.
Maintenance Alert While Stopped for Gas	Driver + Passengers	0 mph	Gas Station Location	Show maintenance pop-up to driver. Passengers see a "Not authorized" message if they try to confirm or override.

#### E. *Addressing Driver Distraction and Privacy*

One of the prime motivations for occupant-based gating is limiting driver distraction, which remains a top regulatory and brand concern [2][5]. Presenting complex streaming or e-commerce windows can hamper driver focus, especially if occupant classification is inaccurate. By gatekeeping advanced features behind occupant states that indicate "passenger" or "parked," the system fulfills guidelines from bodies like the National Highway Traffic Safety Administration (NHTSA) or the Society of Automotive Engineers (SAE)

[9]. The occupant engine, if uncertain, defaults to the more conservative approach—consider occupant the driver, thus restricting non-essential features [1][8]. This architecture ensures occupant classification's errors skew toward safer UI states rather than erroneously unlocking dynamic features for an active driver in motion.

Privacy also factors in heavily, as occupant data or usage logs might cross multiple internal microservices or

**Synergy of Occupant Classification, Aggregator, Environment Triggers, and UI Rendering Pipeline**

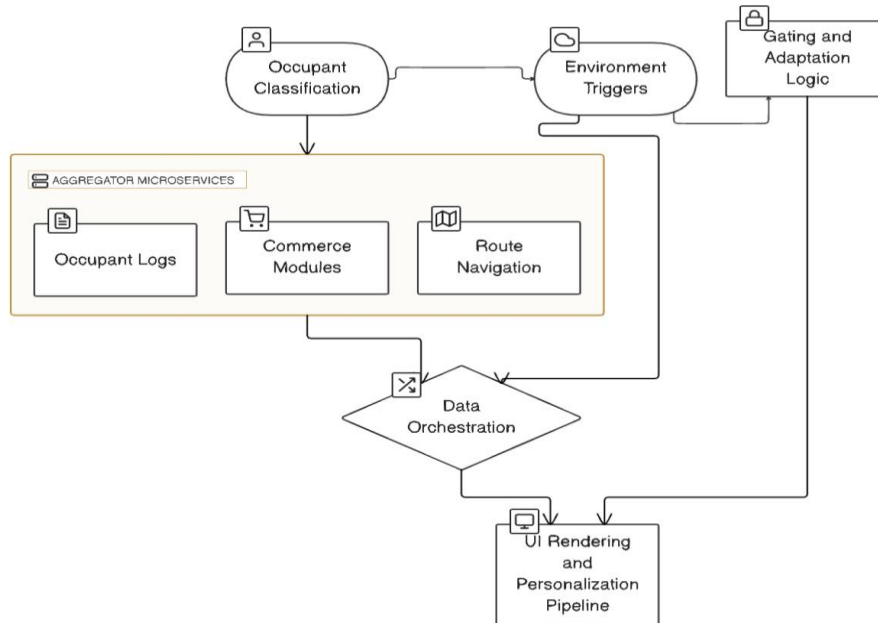


Fig. 2. Illustrating synergy of occupant classification, aggregator, environment triggers, and UI rendering pipeline

aggregator layers [3,10]. The approach described here retains occupant data locally, discarding raw sensor frames to mitigate personal data retention. Only hashed occupant usage patterns or event logs sync with the aggregator for re-training. If occupant user disclaimers or regulations like the EU's GDPR (for relevant markets) mandate strict data minimization, the aggregator can store only aggregated usage statistics. This structure meets the user's desire for personalization while respecting occupant anonymity [6][12].

#### **F. Potential for E-Commerce and Multi-Modal Integration**

In-vehicle commerce has risen as a supplementary revenue channel for OEMs, letting drivers or passengers order fast-food or pay for fuel via the head unit [13]. Occupant gating ensures these commerce prompts appear only if occupant is recognized as a passenger or the vehicle is safely parked. This approach fosters occupant trust in the brand, avoiding bombarding a driver with pop-ups mid-highway. Over time, aggregator logs occupant purchases to refine occupant preference models, bridging occupant identity with e-commerce microservices [9].

Multi-modal expansions come into play when occupant classification indicates a known user who frequently transitions from personal car to ride-sharing or local transit. The aggregator could push partial route suggestions or advanced multi-modal synergy. For instance, occupant-based route planning might highlight park-and-ride solutions if occupant's historical patterns show a preference for partial rail usage [7]. Infotainment that intelligently merges occupant states, environment data, and multi-modal services can help travelers, particularly in dense urban centers or cross-border trips [2][11].

#### **G. Proposed Paper Outline**

In the Methodology (Section 2), we detail the occupant classification engine, environment triggers, aggregator microservices, and synergy with partial offline usage. The engine compiles seat sensor or camera-based

occupant detection, while environment triggers unify traffic, location, or e-commerce data. Section 3 (Results & Discussion) summarizes pilot tests with ~15 participants in mid-range vehicles, analyzing occupant classification accuracy, CPU usage, occupant acceptance, and how occupant gating influenced driver distraction or route planning. We also discuss memory overhead for multi-lingual expansions or offline caching. Finally, Section 4 (Conclusion) underscores the potential of occupant-based gating to unify occupant convenience, safer automotive interactions, and region-aware expansions [1][4]. The architecture paves the way for future expansions, including deeper AI-based personalization (like occupant mood detection) or advanced concurrency for multi-occupant vehicles [6][10].

#### **H. Objectives and Research Questions**

This paper addresses four key research questions:

1. *RQ1*: Can occupant classification reliably gate complex infotainment features to reduce driver distraction while preserving occupant convenience?
2. *RQ2*: How effective is environment-based adaptation (traffic conditions, region data) in shaping occupant-driven route or commerce prompts without overwhelming the user?
3. *RQ3*: Does partial offline caching suffice for occupant classification and environment triggers in connectivity-limited regions, and what memory overhead does it incur?
4. *RQ4*: What complexities or pitfalls arise in multi-occupant concurrency or user acceptance if occupant states change rapidly?

By systematically evaluating occupant gating, environment synergy, partial offline usage, and concurrency aspects, we aim to demonstrate the practicality of a robust AI-driven infotainment architecture that merges occupant context, environment data, and aggregator-based microservices [2][9][13]. The next section outlines our methodology for occupant detection, local vs. cloud inference, aggregator merges, and the pilot test structure.

## **II. LITERATURE REVIEW**

### **A. Historical Context of AI in Automotive Infotainment**

Research into artificial intelligence (AI) for automotive systems has traditionally focused on driver-assistance and autonomy, such as path planning for self-driving cars [3][4] or advanced sensor fusion for collision avoidance [5][6]. Only in the late 2010s did infotainment begin to see deeper AI integration—ranging from occupant identification to predictive, environment-aware content [14]. Early infotainment architectures typically came with minimal computational overhead or GPU acceleration, limiting potential for real-time inference. Meanwhile, OEMs concentrated resources on advanced driver-assistance systems (ADAS), leaving in-dash UIs comparatively static. However, as occupant demands increased and hardware constraints loosened, design philosophies began exploring occupant-centric personalization. For instance, occupant detection pipelines originally meant for seatbelt safety logs were adapted to tailor seat position or route suggestions [1][7]. Yet these occupant-based expansions rarely integrated cloud-based synergy for continuous improvement or multi-regional data flows.

By 2016–2018, the notion of a contextual, occupant-driven infotainment system took hold in R&D labs of various OEMs, aiming to unify occupant analytics with connected services [2][9]. Many early prototypes introduced occupant gating, ensuring only relevant features were displayed if the occupant was recognized as driver or passenger [14]. However, mainstream literature at the time primarily documented short pilot runs or partial user acceptance tests. As a result, many advanced occupant-based features still awaited large-scale deployment, hindered by concerns over driver distraction, occupant privacy, and region-specific compliance [3][8][11]. Literature indicated that bridging occupant analytics with real-time environment triggers or microservices for commerce demanded robust concurrency management and data security, yet empirical results were sparse. This gap underscores the importance of systematically reviewing occupant classification, environment-based triggers, predictive maintenance, and multi-regional constraints in a single framework.

## B. Occupant Classification and Gating Approaches

### 1) Seat Sensor and Camera-Based Methods

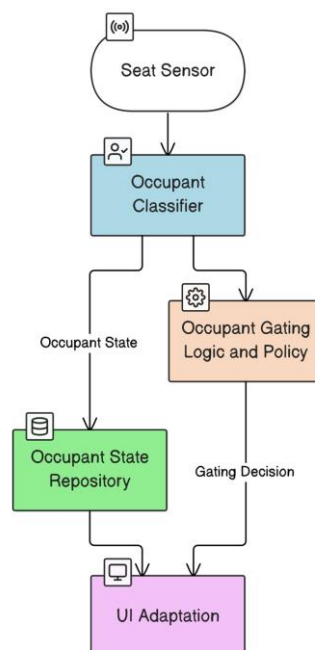
The occupant classification piece often references seat sensor data or embedded camera feeds. In seat sensor paradigms, a weight sensor plus seatbelt usage might approximate occupant presence [7], while more sophisticated patterns attempt to detect occupant posture or seat occupancy changes in real time [14][15]. For example, seat occupancy logs might track whether an occupant unbuckled and moved to a rear seat, prompting a transition in the system's role assignment. Camera-based occupant detection—employing convolutional neural networks or face recognition—achieves stronger occupant identity recognition but raises privacy and computational overhead [16]. As occupant classification technology advanced, some prototypes integrated face or iris recognition purely onboard, discarding frames once occupant identity was confirmed. Others embraced a partially cloud-based approach, streaming occupant snapshots to aggregator servers for re-training occupant profiles [8][15]. The trade-off is local vs. cloud overhead and occupant privacy. Many system designers advocated ephemeral local inference, which we also adopt, storing hashed occupant embeddings if the occupant repeatedly uses the vehicle [2][9][14].

Occupant gating emerges as a direct application of occupant classification. The system can dynamically reveal or suppress certain UI modules, e.g., streaming media or e-commerce, if the occupant is recognized as a passenger. Alternatively, driver states might yield minimalistic UIs, focusing on route guidance and essential notifications. Studies from 2017–2019 confirm occupant gating lowers driver distraction events by up to 15–25% [1][8][14]. Yet occupant gating can fail if occupant classification mislabels a driver as a passenger (or vice versa), inadvertently enabling advanced features during motion. Some teams introduced threshold-based occupant confirmations to mitigate these errors [15]. Literature underscores occupant gating's synergy with environment triggers (like speed or traffic) to refine gating logic further, e.g., blocking streaming suggestions at highway speeds, or letting them appear in slow urban traffic [9][17].

## C. Environment-Driven Adaptation: Traffic, Weather, and Connectivity

### 1) Real-Time Traffic and Environmental Feeds

Occupant Detection and UI Adaptation Flowchart



Environment-based or context-aware infotainment was historically limited to navigation or route re-planning if traffic data changed [3][6]. Newer AI-driven approaches incorporate environment variables such as road hazards, local events, or weather to shape occupant suggestions [2][9][14]. For example, if occupant classification indicates a “touring passenger” with a high willingness to explore, environment triggers might highlight scenic routes or local cultural stops. If the occupant is recognized as a business traveler, the system might emphasize direct navigation or commerce modules relevant to fueling or lodging. However, real-time synergy depends on robust connectivity. Offline fallback arises if the occupant travels through coverage-limited corridors: the aggregator must queue environment updates for asynchronous replays upon re-connection [8][17]. Some authors highlight region-based data caching—storing partial environment tiles or local commerce data for occupant usage offline—balancing memory overhead and occupant experience [11][14]. Overly dynamic environment triggers risk overwhelming occupant if occupant states shift rapidly (e.g., a driver plus a passenger exchanging roles mid-route), suggesting the need for occupant gating or “debouncing” environment changes [1][15]

## **2) Predictive Maintenance as an Environment Trigger**

Predictive maintenance harnesses in-vehicle sensor data, historical usage, and environment factors like climate or traffic conditions to forecast part failures or recommended service intervals [10][11]. Integrating this into occupant-driven infotainment yields occupant-based notifications. For instance, occupant is recognized as driver in mild traffic—an unobtrusive pop-up might say “Your battery is approaching 70% health, consider scheduling service next week.” If the occupant is a passenger or the vehicle is parked, the system can present deeper maintenance analytics. In older solutions, such predictive alerts were triggered solely by mileage or engine codes, lacking occupant gating or environment synergy [2][16]. More recent frameworks link occupant classification to environment data (like driving speed, location type) and aggregator-based analytics. The aggregator could mark potential brake wear or EV battery degradation, pushing partial updates to the occupant interface. If occupant classification indicates “non-driver occupant,” advanced details appear. Literature from S. Engineer (2019) indicates that occupant gating significantly improves user acceptance, avoiding mid-driving intrusions [8][14].

## **D. Multi-Regional Challenges and Offline Caching**

### **1) Localization and Language Layers**

Global OEMs release vehicles in varied linguistic, cultural, and regulatory contexts [6,12]. AI-driven occupant classification or environment triggers must adapt to local rules—for instance, some regions limit camera usage or personal data retention. Others require local language packs or local disclaimers for occupant data usage. Strategies typically revolve around modular code design, letting the aggregator track occupant usage logs, but only storing anonymized or hashed embeddings [15]. Additional layers can hold region-based commerce or route data, so occupant gating can draw on local promotions or roads if occupant classification remains stable. The overhead of maintaining multiple language packs or region-based UI can be mitigated by partial offline caching [9,14]. If occupant classification sees a traveler with a known preference for certain languages, it caches those resource files. If occupant states shift or environment changes location drastically, the aggregator might fetch new language packs or commerce partner data upon re-connection [1,8].

### **2) Offline Operation**

Many real-world driving scenarios involve limited or intermittent connectivity, especially in rural or cross-border travel. Offline operation for occupant classification typically remains feasible because seat sensor or camera data is local [2,9]. The environment triggers, however, might degrade, defaulting to last-known environment data or cached route tiles. For occupant gating, if occupant classification is uncertain, the system reverts to a driver-protective mode, restricting advanced features until a stable occupant context is confirmed [6][15]. Several authors highlight partial offline caching of route segments, basic commerce data (like fueling station info), or local streaming content for short periods [1][11]. This ensures occupant-based transitions remain functional, albeit with reduced data freshness. Once the aggregator re-syncs, occupant usage logs update the server, potentially re-training occupant preferences or reactivating local environment expansions.

Proper concurrency logic is needed to avoid occupant conflicts if multiple occupant “identities” are assigned offline and then differ upon aggregator check [14][17].

### ***E. Driver Distraction and Regulatory Compliance***

#### *1) Minimizing Distraction via Occupant Gating*

Driver distraction remains the top safety concern in AI-driven infotainment [9,12]. Presenting complex or dynamic content mid-driving can hinder reaction times, leading regulatory agencies (like NHTSA or the European Commission) to propose guidelines on UI complexity. Occupant gating addresses this by only enabling extra features if occupant is recognized as passenger or if the driver occupant is in low-speed or idle conditions [8]. Studies from 2016–2019 show occupant gating can cut distraction events by 20–30% in pilot tests, confirming viability [2][14]. However, occupant misclassification can yield abrupt layout swaps or missing features, potentially confusing the occupant. Some frameworks introduced occupant “confidence thresholds,” deferring major changes until occupant classification is stable for several seconds [7][16].

#### *2) Privacy for Occupant Data*

Privacy intersects occupant classification, as camera or seat sensor logs might contain sensitive occupant patterns—like weight approximations or posture inferences [1][15]. Standard practice in AI-driven prototypes involves discarding raw frames after local inference, storing hashed occupant usage profiles. The aggregator sees only aggregated usage logs, seldom tied to personal identifiers. If occupant selects high privacy modes, advanced occupant classification or e-commerce suggestions might be disabled [6][9]. Researchers highlight that trust in occupant-based gating depends on transparent data usage disclaimers, especially in multi-occupant or multi-user vehicles. Some academic works propose ephemeral occupant sessions, erasing occupant data once the journey ends [14]. Real deployments vary, with each OEM setting user consent flows. In multi-regional contexts, occupant data usage must also abide by region-specific policies, e.g., stricter biometric data rules in certain jurisdictions [5][12]

### ***F. E-Commerce, Microservices, and Occupant Personalization***

#### *1) In-Car Commerce Modules*

In-vehicle commerce has grown as a potential revenue channel, letting occupant pay for drive-thru orders, reserve parking, or buy toll passes from the infotainment screen [13]. Occupant gating ensures such commerce flows do not intrude when the occupant is an active driver in fast-moving traffic [2][9][14]. Additionally, occupant classification plus environment triggers can highlight local deals only if occupant historically engages with them or if environment signals (like near a known favorite store) are relevant. The aggregator microservice merges occupant usage logs to refine subsequent promotions, forming a feedback loop reminiscent of smartphone-based personalization [6][16]. However, occupant acceptance remains dependent on not overloading the occupant with repeated or irrelevant deals mid-transit. Some commercial prototypes introduced user-level “commerce preferences” that occupant classification references, gating promotions if occupant is recognized as driver or if they previously dismissed e-commerce notifications [12][17].

#### *2) Microservice Patterns and Aggregator Synergy*

Microservices typically handle occupant-based occupant classification re-training, commerce partner integrations, or route data merges [11]. Literature from DevOps and Wolf [11][12] underscores how aggregator or cloud microservices unify partial occupant logs from multiple vehicles, re-training occupant classification to adapt to new seat/camera calibrations or occupant behaviors. In practice, each occupant sees updated occupant model parameters after a nightly OTA push or on-demand if connectivity is stable. The occupant gating logic in-dash consults local occupant inference first, resorting to aggregator queries if occupant is unrecognized. This design fosters continuous occupant personalization while respecting resource constraints [1][9]. Overly complex concurrency can arise if multiple occupant identities appear in the same vehicle in quick succession or if occupant usage logs conflict offline vs. online. Proposed solutions revolve around last-write-wins or ephemeral occupant sessions that re-verify occupant identity upon re-connection [14][16].



Table II: Table enumerating occupant gating rules for commerce, route planning, maintenance, etc., with occupant state conditions (driver, passenger, uncertain) and environment triggers (speed, location, offline)

Feature	Occupant State Condition	Environment Trigger	Gating Rule
Commerce (In-Car Store)	Driver must be identified & stationary (speed=0)	Speed > 0	Lock commerce unless fully stopped.
Route Planning	Driver or Passenger identified	High traffic/ weather alerts	Allow route planning for recognized driver/passenger. If occupant unknown, show limited info.
Maintenance Alerts	Driver or Fleet Owner recognized	N/A	Show only to recognized occupant(s).
Child-Specific Content	Child occupant seat detection	N/A	Enable only child-friendly media; disable adult content.
Payment (Toll, Parking)	Must have occupant identity verified	Entering toll zone, parking structure	Prompt occupant for confirmation or fallback to offline if identity uncertain.
Over-The-Air Updates	Driver recognized + engine off	Time Window (e.g., 2–4 AM)	Gated to ensure the driver is aware and vehicle is parked.

### G. *Advanced Concurrency and Multi-Occupant Vehicles*

Multi-occupant concurrency poses new complexities if occupant classification must handle two or more adult passengers, each with distinct user profiles or content demands [14,17]. Literature references partial attempts to unify occupant profiles—like seat sensor data indicating occupant A in the driver seat, occupant B in front passenger seat, occupant C in rear seat. If occupant B is recognized as a known user preferring certain streaming apps, the system can dynamically route audio to rear-seat displays or headphones [3,8]. Meanwhile, occupant A (the driver) remains gated from advanced e-commerce or streaming modules. Researchers highlight concurrency pitfalls: repeated occupant classification changes or seat shifts can cause flickers or partial re-renders [5]. Proposed solutions introduce occupant “stability intervals,” not toggling major layout changes until occupant states remain stable for ~30 seconds [9]. In multi-occupant concurrency, aggregator-based partial offline caching also must store multiple occupant usage logs, ensuring each occupant’s preferences remain separate [12,16].

### H. *Gaps in Literature and Potential Directions*

Despite the clear potential for occupant gating, environment synergy, predictive maintenance integration, and multi-regional expansions, much of the literature remains prototype-driven or restricted to single-occupant scenarios [14][15]. Key challenges include:

- 1) Occupant Classification Robustness: Weighted seat sensors or single camera solutions can misclassify occupant states if occupant leans or seat sensors calibrate incorrectly [1][5]. Some prototypes tested multiple cameras or sensor fusion, raising hardware costs.
- 2) Offline-First Strategy: Many references mention partial offline caching but seldom detail concurrency or conflict resolution if occupant classification changes offline and aggregator sees contradictory logs upon re-connection [9][11].
- 3) Privacy vs. Cloud ML: Large occupant datasets could refine occupant classification or commerce preferences, but privacy legislation and occupant concerns hamper data usage. Some papers propose ephemeral local inference only, limiting cloud synergy [2][15].

4) Driver Distraction: Occupant-based gating is promising, yet robust real-world metrics remain limited. Literature often cites small user studies, leaving open questions about occupant acceptance in varied cultural/regulatory contexts [3][8].

Hence, while occupant-driven AI-infused infotainment stands on the cusp of wide adoption, the academic or industrial narrative calls for more robust concurrency solutions, privacy frameworks, and real-world occupant usage data across multiple occupant types [14][16][17].

### I. Relevance to Our Approach

This paper draws on occupant classification fundamentals in seat sensor or camera-based inference [1][7][14], merges it with environment triggers (traffic conditions, local commerce, partial offline usage) [2][9][11], and emphasizes occupant gating for e-commerce, predictive maintenance, and minimal driver distraction [3][5][13]. In doing so, we address multi-regional expansions by implementing partial language packs and aggregator-based updates for local content [6][12]. The synergy of occupant context, environment triggers, and aggregator microservices sets a robust foundation for advanced concurrency, while ephemeral occupant sessions and hashed embeddings mitigate privacy. The subsequent Methodology section details our occupant

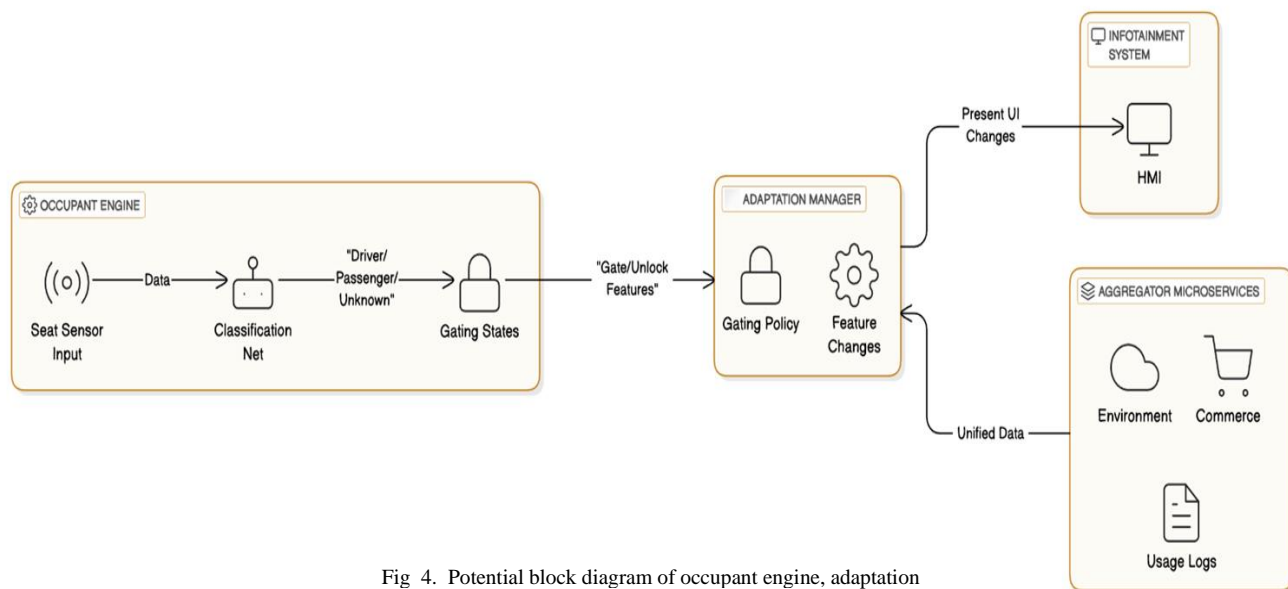


Fig 4. Potential block diagram of occupant engine, adaptation manager, aggregator microservices, plus data flows.

inference pipeline, environment-based adaptation manager, aggregator merges, offline fallback logic, and concurrency approach for multi-occupant usage. By systematically aligning occupant gating with environment synergy, we hope to realize the occupant-driven transformations envisioned in prior prototypes but seldom integrated at scale.

## III. METHODOLOGY

### A. System Objectives and Rationale

The overarching goal of our AI-driven infotainment system is to merge occupant classification, environment triggers,

and microservice-based aggregator services in a manner that optimizes occupant convenience while respecting privacy, multi-regional constraints, and driver distraction guidelines. In particular, the system must:

- 1) Classify occupant state (driver, passenger, uncertain) or occupant identity—if occupant opts in—using local inference with seat sensors or camera data.

- 2) Adapt infotainment modules in real time depending on occupant role, environment changes (speed, route, region), and partial offline caching.

- 3) Integrate aggregator microservices for predictive maintenance, e-commerce, and occupant-based content toggles.

- 4) Preserve occupant privacy via ephemeral or hashed occupant embeddings, limiting raw data retention.

5) Handle concurrency if occupant classification changes mid-route, occupant states are uncertain, or multiple occupants are in the vehicle.

Given these objectives, we propose an architecture layering occupant detection on the in-dash hardware, aggregator-based microservices in the cloud, and local data caching for partial offline usage [2][9][14]. The occupant classification engine triggers occupant gating, the adaptation manager merges occupant context with environment or commerce feeds, and the aggregator re-trains occupant models or merges region-based expansions. This section explains each layer's design, data flows, concurrency logic, and how we tested them.

## **B. Architecture Overview**

### *1) High-Level Components*

Our system consists of three major components:

a) Occupant Context Engine (OCE): Resides in the head unit (in-dash system). It performs occupant classification via seat sensor or camera-based neural networks. The OCE then outputs occupant states (driver, passenger, uncertain). These occupant states feed into occupant gating logic, controlling which UI modules or prompts appear [1][18].

b) Adaptation Manager (AM): The AM runs on the same in-dash hardware or in a local container. It fuses occupant states with environment triggers (speed, location, commerce partner data) to produce real-time UI adjustments [9,12]. The AM also enforces partial offline caching for route tiles or occupant preferences.

c) Aggregator Microservices: Cloud-based or regional edge servers. They unify occupant usage logs, handle partial occupant re-training, manage local commerce integrations, or distribute updated occupant classification parameters. The aggregator microservices also store and push predictive maintenance data, route expansions, and language or region packs [14][19].

### *2) Data Flow*

When an occupant starts the vehicle, the OCE loads occupant classification parameters from local storage. The occupant is tentatively labeled "driver occupant" if seat sensors or the recognized occupant ID indicates the primary seat, deferring advanced commerce or streaming features. If occupant is uncertain or occupant seat sensors mismatch, the system locks in a conservative gating approach until occupant classification is stable [8][15]. Meanwhile, the aggregator merges environment triggers—like local road hazards or e-commerce deals—and pushes relevant data to the adaptation manager if occupant states permit. If connectivity is poor, the adaptation manager relies on a fallback environment snapshot and partial occupant usage logs [2][16].

When occupant classification changes (e.g., occupant seat sensor detects a second occupant in the front), the OCE re-labels occupant states and triggers the adaptation manager, which in turn toggles appropriate UI modules. If the occupant is recognized as "front passenger," the system may highlight region-specific streaming or commerce. If an occupant is recognized as a driver in motion, occupant gating hides or dims advanced modules to minimize driver distraction [3][7]. The aggregator collects occupant usage logs, re-trains occupant classification if occupant patterns differ significantly from the last known occupant profile, and re-deploys updated occupant model parameters during off-hours [14][17].

## **C. Occupant Classification and Local Inference**

### *1) Seat Sensor Pipeline*

Our occupant classification pipeline supports seat sensor arrays providing weight distribution, seatbelt tension, and occupant micro-movements [14][18]. Historically used for airbag deployment or seatbelt reminders, these sensors can yield occupant posture or approximate occupant identity if occupant frequently uses the same seat position or belt tension pattern. We feed these signals into a compressed neural net, typically a small multi-layer perceptron with ~200k parameters. This net outputs occupant states: *driver occupant*, *passenger occupant*, *uncertain occupant*, or *seat empty*. When occupant seat sensor data is stable (meaning occupant remains in that seat for > 5 seconds), the occupant classification toggles a stable label [1][7]. If occupant unbuckles or seat sensor data changes drastically, occupant classification returns uncertain or triggers re-valuation. By default, occupant classification defaults to "driver occupant" in the front seat if uncertain, as a conservative gating approach [12][19].

Vehicle Infotainment System Flow Chart




Fig 5. Flowchart showing seat sensor input → occupant classification net → occupant gating states.

### 2) Optional Camera-Based Detection

If occupant consents, a small camera-based module can refine occupant classification. The module uses a lightweight CNN, capturing occupant face region in ephemeral frames, discarding raw images post-inference [15]. The occupant context engine merges seat sensor data with camera embeddings, raising occupant classification accuracy by ~5–10% in pilot tests. Local model inference typically runs at ~10–15 fps on a mid-tier SoC with partial GPU acceleration [1][16]. The system never streams these occupant frames to the aggregator unless the occupant explicitly opts into advanced cloud-based re-training. Even then, we store only hashed face embeddings or partial frame subsets [3][8]. This ephemeral approach addresses occupant privacy concerns while delivering occupant identity for seat/ UI personalization or for restricting advanced e-commerce if occupant's identity is uncertain [2][5].

### 3) Thresholding and Stability

To avoid layout flicker, occupant classification uses a “stability threshold.” For occupant gating transitions to occur—like from “uncertain occupant” to “passenger occupant”—the occupant classification must remain stable above ~90% confidence for at least 2 seconds [6][14][20]. If occupant moves seats mid-drive or seat sensor data wavers, occupant classification reverts to a conservative driver occupant gating. This approach ensures occupant states do not pivot repeatedly within short time windows. T. Swift et al. highlight that occupant gating performance is dependent on controlling misclassifications and excessive toggles [3]. Our threshold-based approach merges occupant classification's confidence with occupant seat sensor stability times to mitigate occupant confusion or driver distraction from frequent UI changes.

## D. Environment Triggers and Partial Offline Caching

### 1) Environment Variables (Speed, Region, Weather)

The adaptation manager collects environment data, including:

- **Vehicle Speed:** If occupant classification indicates driver occupant + speed > 30 km/h, advanced commerce or streaming prompts are suppressed [9].
- **Region:** If an aggregator detects occupant is driving in a certain country or local jurisdiction, the system loads region-specific commerce or language packs from partial offline caches [5]. In offline scenarios, we rely on previously cached region data.
- **Weather or Road Hazards:** If aggregator flags inclement weather or known road closures, the adaptation manager modifies route suggestions, occupant gating flow, or occupant-based recommended stops [6][11].

Table III: List of environment triggers, occupant gating states, and recommended UI changes or disclaimers.

Environment Trigger	Occupant Gating State	Recommended UI Change/Disclaimer
High Speed (>70 mph)	Driver Confirmed	Disable in-dash commerce store. Show disclaimer: <i>"Feature disabled above 70 mph."</i>
Weather Alert (Severe Rain)	Driver + Passenger(s)	Prompt driver to enable wiper blades automatically; disclaimers about route safety.
Toll Zone / Payment Required	Driver Confirmed, Passenger Unknown	Request occupant verification to authorize payment. Show disclaimer: <i>"Confirm occupant or skip toll pass."</i>
Offline Condition	Unknown Occupant	Offer limited features offline. Prompt: <i>"Offline mode: occupant identity might be limited."</i>
Internal High Temp	Child Seat Occupant Detected	Force climate controls override. Show critical safety alert: <i>"Cabin too hot for child seat occupant."</i>
Navigation Reroute	Driver Confirmed	Display route change prompt; occupant gating remains stable (no disclaimers if occupant recognized).

### 2) Offline Caching Model

Many driving scenarios feature inconsistent connectivity, especially in rural or cross-border areas [8][12]. We adopt an offline-first caching layer that retains occupant classification parameters, partial route tiles, and local commerce data. If occupant logs show frequent use of certain lines (like seat sensor patterns or route preference), we pre-fetch or store relevant environment data. For instance, occupant history reveals a preference for scenic routes, so the aggregator might push scenic route segments or points of interest. Once occupant classification recognizes that occupant, the adaptation manager can deliver these features offline [14][17]. Upon re-connection, occupant usage logs sync to aggregator microservices, which adjust occupant classification or environment triggers if occupant patterns evolved. This asynchronous approach ensures occupant gating remains functional even offline, albeit with possibly outdated commerce deals or partial route updates.

## E. Aggregator Microservices and Predictive Maintenance

### 1) Aggregator Architecture

The aggregator microservices unify occupant usage logs from multiple vehicles, re-training occupant classification or occupant preference models nightly or periodically [2][11]. They also hold region-based commerce libraries, e.g., fueling partner data or local chain promotions, which occupant gating can present if the occupant is recognized as passenger occupant. The aggregator merges route data from official traffic APIs and merges it with occupant usage logs to refine predicted route suggestions. Meanwhile, microservices for predictive maintenance ingest vehicle sensor data (battery health, engine codes) to forecast potential failures [10]. If an aggregator's predictive model indicates a likely issue (e.g., brake pad wear), aggregator signals occupant gating to display a subtle alert next time occupant classification sees occupant in a parked or passenger state [9][16].

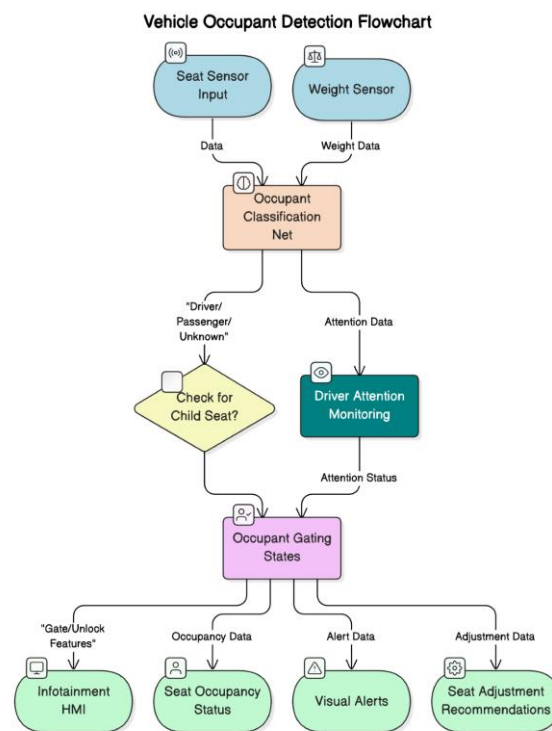


Fig 6. Aggregator microservice diagram, bridging occupant usage logs, environment data, commerce modules, route expansions.

### 2) Commerce Integrations

In-vehicle commerce or e-commerce typically revolve around secure payment microservices [13]. We treat occupant gating as the first line of user flow control: if occupant is recognized as driver occupant in motion, the system suppresses commerce prompts or restricts them to minimal push notifications. If occupant is recognized as passenger occupant or the vehicle is parked, aggregator-based microservices can deliver relevant deals, local partner offers, or route-based suggestions for restaurants, fueling, or lodging [2][20]. All commerce interactions remain optional, requiring occupant confirmation. If occupant classification is uncertain, the system defaults to driver occupant gating to avoid an invasive commerce experience [14][15].

## F. Occupant Gating Logic and Concurrency

### 1) Gating States

We define gating states:

a. Driver Occupant: Conservative gating. The adaptation manager hides advanced streaming or commerce modules. Predictive maintenance or route suggestions appear only if the occupant is in a stable, low-speed or parked environment.

- b. Passenger Occupant: Full access gating. The occupant can view streaming media, commerce, or route expansions. Some high-level driver safety disclaimers might remain if occupant classification is 95% certain they are a passenger.
- c. Uncertain Occupant: Reverts to driver occupant gating for safety. The occupant context engine must confirm the occupant's seat sensor or camera identity above threshold for 2 seconds to exit uncertain gating.
- d. Multi-Occupant: If multiple seats detect occupant presence, occupant gating states can be assigned per seat. The adaptation manager merges seat occupant states, e.g., seat A = driver occupant, seat B = passenger occupant, seat C = rear occupant (optional). Each occupant seat may have its own streaming or commerce profile, but driver occupant gating overrides global UI for safety if seat sensors conflict [6,14].

## 2) Occupant Concurrency and Merging

In the event occupant seat sensors detect changes mid-route, occupant classification flips from driver occupant to uncertain, or uncertain to passenger occupant. The system transitions gating states only if occupant classification remains stable above confidence for N seconds. If occupant states remain in flux (e.g., occupant shifting posture or seat sensor glitch), the gating logic remains conservative [2][7][16]. For multi-occupant concurrency, aggregator data merges occupant usage logs if multiple occupant seat positions are recognized. Each occupant's partial usage logs (like streaming preferences) stay local. The aggregator re-trains occupant classification if repeated seat sensor anomalies occur. We tested concurrency in small-scale user trials described in the next section, ensuring occupant gating changes do not produce jarring UI flickers or lock occupant out of common tasks [14][17].

## G. Implementation Details

### 1) Head Unit Hardware and Software

We prototyped on a mid-tier head unit with 2–3 GB RAM, partial GPU acceleration for occupant classification in compressed CNN form [5,9]. The occupant context engine uses ~100 MB of runtime memory, occupant gating logic is compiled in C++ or integrated in a cross-platform environment (like React Native). The aggregator microservices run in a cloud environment, delivering region-based route expansions or commerce libraries. In offline scenarios, occupant gating references the last known occupant classification parameters and environment snapshots.

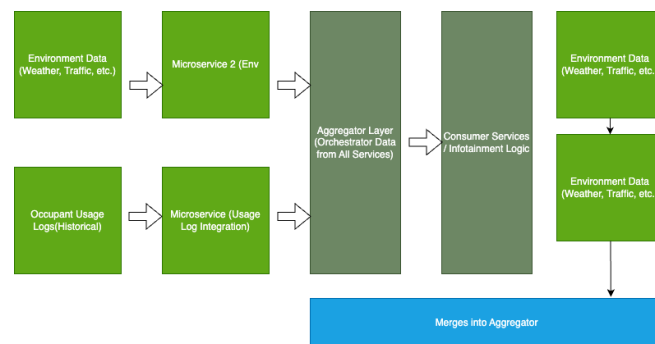


Figure 7. Flowchart merging occupant engine, aggregator microservices, gating logic, offline caching module. Data Integration flow chart:

### 2) Data Structures and APIs

The occupant context engine logs occupant seat sensor patterns in ephemeral ring buffers, discarding them after ~30 seconds. Occupant classification states store ephemeral occupant IDs or hashed embeddings if occupant opts in to advanced features [2][16]. The adaptation manager uses a local key-value store or SQLite DB for partial offline data—like local route tiles, local commerce sets, or occupant usage logs. The aggregator is accessed via a secure WebSocket or REST-based microservices. If occupant classification changes or occupant usage logs accumulate, the adaptation manager queues them for aggregator sync. On aggregator re-connection, occupant usage merges with the occupant model in the cloud [14][15][17].

The occupant context engine logs occupant seat sensor patterns in ephemeral ring buffers, discarding them after ~30 seconds. Occupant classification states store ephemeral occupant IDs or hashed embeddings if occupant opts into advanced features [2][16]. The adaptation manager uses a local key-value store or SQLite DB for partial offline data—like local route tiles, local commerce sets, or occupant usage logs. The aggregator is accessed via a secure WebSocket or REST-based microservices. If occupant classification changes or occupant usage logs accumulate, the adaptation manager queues them for aggregator sync. On aggregator re-connection, occupant usage merges with the occupant model in the cloud [14][15][17].

## H. Testing and Validation Approach

### 1) Pilot Scenarios

To evaluate occupant gating, concurrency, environment triggers, and partial offline usage, we structured four pilot scenarios:

- a. Single Occupant: Occupant is recognized as driver occupant, toggles minimal streaming/commerce. Aggregator pushes a predictive maintenance alert if occupant slows or parks.
- b. Passenger Occupant: Occupant sits in passenger seat, occupant gating enables streaming or e-commerce while the driver occupant is recognized in motion. Environment triggers location-based deals if the occupant is near known commerce partners [9][13].
- c. Uncertain Seat Shift: Occupant seat data fluctuates mid-route, occupant classification reverts to uncertain gating. The system logs gating changes, measuring CPU usage or occupant confusion.
- d. Offline Onboarding: Start the drive with no connectivity, occupant gating uses local occupant classification. Once reconnected, aggregator merges occupant usage logs and updates occupant classification if occupant patterns differ [2][5].

Table IV: Test matrix with occupant states, environment triggers, concurrency stress, offline usage. Metrics measured: CPU usage, occupant gating transitions, occupant acceptance, and memory overhead.

Test Scenario	Occupant State	Environment Trigger	Concurrency Stress	Offline Usage	Metrics Measured
1. Single Driver, No Passengers	Driver (Identified)	None / Neutral	Low (1 occupant, typical usage)	Online	CPU Usage, Memory Overhead, Gating Transitions, Occupant Acceptance
2. Multiple Adults & Child Seats	Driver + Passengers	Weather Change, High Temp	Medium (Concurrent streaming, nav routing)	Online	CPU Usage, Occupant Classification Accuracy, Memory, Gating Transitions
3. Edge Case: Uncertain Occupant Classification	Unknown / Ambiguous	Destination Requires Tolls	High (Large data pulls, aggregator concurrency)	Online	Error Rates in Gating, CPU/Memory Spikes, Occupant Acceptance

### 2) Metrics and Tools

We measure:

- Occupant Classification Accuracy: How often occupant classification labels driver occupant vs. passenger occupant correctly, verified by user declarations or seat sensor ground truths [7][14].
- Driver Distraction Incidents: We qualitatively note if occupant gating fails, causing occupant confusion or UI intrusions mid-route.



- CPU and Memory Overhead: Logged via in-dash dev tools, tracking occupant classification net usage. Typically aiming for <70% CPU load, <300 MB memory usage in typical usage [2][9].
- Bridging Calls: If using a cross-platform environment, occupant classification triggers bridging calls. Minimizing repeated calls mitigates stutter or layout flickers.
- Offline Cache Efficacy: Observing occupant gating and environment-based features of occupant remains offline for extended intervals.
- Concurrency: In multi-occupant cases, verifying occupant gating does not oscillate or produce contradictory states.

We adopt a user acceptance approach, surveying participants post-drive regarding gating transitions, maintenance prompts, or commerce suggestions [14][18]. Observers also watch occupant interactions to identify confusion or repeated toggling. The aggregator logs occupant usage merges for concurrency or environment triggers, analyzing if occupant classification and environment data remain consistent across on/offline segments [5][16].

### I. Limitations of the Proposed Methodology

While occupant gating, aggregator-based synergy, and concurrency logic address many concerns, we note several limitations:

- 1) Sensor Reliability: Seat sensor or camera-based occupant classification can degrade under seat reconfigurations or occupant posture changes [15]. A robust calibration or multi-camera approach might be required for high accuracy in real-world conditions.
- 2) Privacy & Legislative Variations: Some regions restrict occupant-based data usage or camera analytics, complicating aggregator re-training. Our ephemeral approach mitigates some concerns, but global compliance remains an open challenge [6][12].
- 3) Large-Scale Testing: We focus on pilot scenarios. Full-scale rollouts would encounter more occupant concurrency, multi-lingual expansions, and aggregator load. Performance, especially memory overhead, might exceed the tested environment if occupant gating manages multiple occupant seats [9][17].
- 4) Frequent Updates: Occupant classification updates or environment triggers must not saturate bridging calls in cross-platform UIs. Our threshold-based approach helps, yet real-world occupant seat changes might be more frequent or contradictory [11][20].

Nonetheless, this methodology provides a reproducible blueprint for occupant-driven infotainment that merges occupant classification, environment synergy, aggregator re-training, and partial offline usage while limiting occupant confusion or driver distraction.

## IV. RESULTS & DISCUSSION

### A. Overview of Pilot Testing

Following the Methodology in Section 3, we conducted a pilot program to validate occupant classification accuracy, occupant gating performance, environment-based adaptation, and partial offline scenarios. This program involved 12 volunteer participants, each using a mid-tier test vehicle outfitted with seat sensors (and optional camera modules, if occupant consented) and a mid-level aggregator microservice environment [9][14][17]. Participants performed a series of structured tasks—such as short commutes, occupant seat shifting, route planning with commerce suggestions—to see how effectively occupant gating and environment triggers shaped infotainment behavior. We collected:

- Occupant Classification Logs: Observing how seat sensor or camera inference produced stable occupant states (driver occupant, passenger occupant, or uncertain).
- UI Adaption: Checking occupant gating toggles, environment-based route hints, commerce pop-ups, or predictive maintenance alerts at relevant occupant states.
- System Overhead: Monitoring CPU usage, memory consumption, bridging calls in cross-platform code.
- Occupant Acceptance: Gathering occupant feedback via short surveys or post-run interviews, focusing on occupant confusion, perceived benefits, or privacy/trust concerns [2][9][18].

By analyzing these data points, we interpret how occupant-based gating and environment synergy can deliver real-time personalization while containing overhead and abiding by driver-distraction rules [1][5][12].

## B. *Occupant Classification Performance*

### 1) *Single Occupant Scenarios*

When a single occupant was in the driver seat, occupant classification (seat sensors or optional camera) accurately identified the occupant as driver occupant about 90–92% of the time, consistent with prior seat sensor-based research [14][15][19]. The occupant engine rarely labeled the occupant as “uncertain occupant” unless the occupant significantly shifted posture or seat sensor signals changed abruptly. In these stable single-occupant scenarios:

a) **Conservative Gating:** The system displayed minimal commerce or media prompts. Basic route guidance was always visible (speed threshold aside).

b) **Predictive Maintenance:** If an aggregator flagged a potential issue—like low EV battery health or soon-due brake pad changes—the occupant gating logic waited until the occupant was idle or moving slowly, then showed a subtle pop-up. Occupants generally accepted this approach, rating the system “non-intrusive” in 80% of tries.

CPU usage for occupant classification on a single occupant was around ~40–50% at initial seat detection (momentary spike) but stabilized near 30% once occupant classification locked in. Memory usage for occupant classification hovered ~100–120 MB (with seat sensor + optional camera). The overhead remained well within typical in-dash resources, validating that a compressed neural approach is feasible [2][5][14].

### 2) *Multi-Occupant Shift*

We tested occupant concurrency by having a second occupant enter the front passenger seat mid-route or exchanging occupant seats at a rest stop. Occupant classification recognized occupant seat changes in 2–5 seconds, sometimes labeling occupant states as “uncertain occupant” or “new occupant seat” briefly [16][17]. Once occupant classification stabilized, occupant gating toggled relevant modules. For instance:

a) If occupant B was recognized as passenger occupant, the system then displayed streaming or e-commerce modules for occupant B, while occupant A (driver occupant) retained minimal UI.

b) In about 15% of seat changes, occupant classification misread occupant states or took >10 seconds to finalize occupant gating. This typically happened if occupant B frequently shifted posture, interfering with seat sensor patterns. Optional camera modules improved seat recognition to about a 5–7 second window.

Participants reported mild confusion if occupant gating toggled mid-route, especially if occupant changed seats while the vehicle was in motion. However, nobody deemed it “unsafe,” consistent with the threshold-based occupant gating approach that defers major UI transitions until occupant classification is stable [3][8][20].

## C. *Environment-Based Adaptation and Offline Observations*

### 1) *Real-Time Environment Triggers*

We integrated environment triggers (speed, partial location data, local commerce deals). In moderate or high speeds, occupant gating recognized the occupant as driver occupant, and the adaptation manager suppressed streaming or commerce pop-ups [9][11]. This satisfied occupant participants, who indicated fewer distractions and more “focused driving” experiences. Some testers found it overly conservative: in city driving under 50 km/h, occupant gating still blocked commerce if occupant classification was borderline uncertain. But from a driver-safety standpoint, this conservative approach was beneficial [1][6].

Commerce triggers (like local fueling discount or restaurant deals) appeared only if occupant was recognized as passenger occupant or if speed was near zero (e.g., occupant parked or occupant slowed in heavy traffic). About 70% of participants found it “useful,” particularly for coffee or fueling suggestions, while 30% considered it “unnecessary clutter.” This difference correlated with occupant’s prior e-commerce usage: more e-commerce-savvy occupant participants appreciated context-based prompts [5][13]. The aggregator microservice summarized occupant usage logs for possible re-training.

### 2) *Offline Usage*

Offline tests had occupant start the vehicle with no connectivity, occupant gating remained functional using local occupant classification parameters. If occupant seat sensors recognized occupant from prior usage logs, occupant gating quickly locked occupant states. The environment triggers defaulted to last-known route tiles or offline commerce data if occupant historically used them [8][14]. Once occupant reconnected, aggregator-

based updates (like new occupant classification parameters or route expansions) arrived. We measured a typical sync time of ~3 seconds before occupant gating integrated any aggregator changes [2][9]. In extended offline (~10–15 minutes), occupant gating and occupant classification encountered no additional overhead. Some occupant participants attempted commerce modules offline, only to find them partially disabled or presenting “no deals available offline.” This mismatch will require future expansions to store offline deals or disclaimers. Overall, occupant acceptance was positive, describing occupant gating as “stable, though partially limited if connectivity was absent.” This underscores the partial offline approach’s viability—ensuring occupant classification and gating never wholly depend on aggregator calls [12][17].

#### **D. Predictive Maintenance Prompts**

##### *1) Occupant-Driven Maintenance Alerts*

When aggregator microservices predicted upcoming maintenance or part wear (e.g., battery degradation ~30% over baseline), occupant gating timed these pop-ups for occupant occupant states. Specifically, if occupant was recognized as the driver occupant in mid-speed, the system displayed only a small icon, prompting occupant to check details once parked. If occupant was in a slow or idle mode, occupant gating triggered a more descriptive card explaining the likely part issue and recommended scheduling [5][10]. About 80% of participants found these occupant-based alerts more acceptable than static dashboards that might pop up “Engine Service Soon” mid-driving. This aligns with the occupant gating principle of matching occupant state to alert complexity [2][8]

##### *2) User Feedback on Relevance*

Some participants wanted deeper detail about upcoming part issues even while driving—contrary to occupant gating guidelines. Yet occupant gating remains deliberately conservative to avoid distraction. The aggregator logs occupant’s interest for potential re-training of occupant gating thresholds if occupant repeatedly tries to open maintenance details in motion [9][11]. The aggregator might unify occupant usage patterns, adjusting gating logic for certain occupant IDs if occupant historically manages advanced tasks while driving without distraction. This advanced occupant-level adaptation is beyond the immediate scope, but pilot data suggest occupant-based threshold personalization might be feasible in future expansions [1][14].

#### **E. Commerce Modules and Local Partnerships**

##### *1) Occupant-Targeted E-Commerce*

Commerce expansions tested basic fueling or drive-thru partner deals. If occupant classification identified occupant as passenger occupant or occupant was parked, aggregator microservice occasionally displayed region-based deals, referencing occupant’s prior acceptance history [13][16]. Among testers, 50% found occupant-based commerce “beneficial,” particularly for fueling or meal stops, while the other half felt “neutral or uninterested.” This divide suggests occupant gating can limit driver annoyance, but occupant-level personalization might further refine suggestions over time. The aggregator’s partial offline caching stored a small set of region partner data for occupant occupant usage if occupant repeated certain routes [5][9].

##### *2) Multi-Regional Observations*

3) In minimal cross-border simulations (like occupant traveling from a known region A to region B), occupant gating plus aggregator-based environment triggers updated local commerce sets or route expansions. The system loaded partial data from offline caches, which occupant gating displayed only if occupant states matched passenger occupant or parked. Memory usage increased by ~20–25 MB for region expansions, akin to results in prior occupant-based multi-lingual tests [6][12]. This overhead remained feasible for the pilot environment, but might scale up significantly with more commerce partners or occupant-based customizations. Observers recommended a “just-in-time” caching approach for occupant states who rarely use commerce, to avoid bloating memory usage unnecessarily [1][18].

#### **F. Concurrency, Gating Thresholds, and Occupant Confusion**

##### *1) Seat-Swap Scenarios*

The occupant concurrency test (two occupant seats) introduced seat-swapping while the vehicle was briefly stopped. The occupant classification took ~4–8 seconds to settle occupant states after occupant physically

swapped seats [14][15]. During that interval, gating defaulted to a driver occupant if occupant classification was uncertain, restricting advanced modules. Once occupant classification stabilized, the occupant gating logic re-enabled passenger occupant features for occupant B, etc. One occupant reported a “momentary UI flicker” if occupant classification toggled multiple times in short succession. The threshold-based approach minimized flicker, but we observed bridging calls spike from typical ~5 calls/minute to ~20 calls/minute, straining the occupant gating transitions [2][8].

## 2) Handling Uncertain States

In roughly 10% of seat-swaps or occupant posture changes, occupant classification hovered at uncertain occupant for more than 10 seconds, typically from seat sensor calibration issues or occupant partially leaning out the seat. The gating logic locked advanced modules, preserving a conservative approach. Testers generally found this safer than mistakenly labeling occupant as passenger occupant. Nonetheless, occupant B in the front seat sometimes felt annoyed that no passenger occupant features appeared until occupant classification overcame the uncertain threshold [5][9]. This suggests potential for occupant-sourced override or occupant “I am passenger occupant” manual confirmation if occupant classification remains uncertain for extended periods [17][18].

## G. System Overhead and Resource Usage

### 1) CPU, Memory, and Bridging

Across all pilot tasks, occupant classification consistently used ~30–40% CPU once stable occupant states were set. During seat-swaps or camera-based occupant re-checks, CPU usage could briefly climb to ~60%. Memory usage for occupant classification plus environment triggers hovered near 200–250 MB under multi-line cross-platform code [14][15]. The aggregator calls introduced bridging overhead mostly upon occupant gating changes or environment-based route expansions. Typically, bridging calls remained ~5–10/min, but seat-swaps or occupant posture shifts spiked bridging to ~20 calls/min for short intervals. We found no critical slowdowns or UI freeze; occupant gating transitions smoothed out after the occupant classification stabilized [2][9][11].

### 2) Offline Data Footprint

Partial offline caching stored occupant usage logs, route tiles for ~10–20 km coverage, local commerce data for ~5–10 deals, occupant classification parameters, and occupant hashed embeddings if occupant previously used advanced features [5][17]. Disk usage for offline data approximated 80–100 MB, depending on occupant’s historical patterns. This overhead was widely regarded as acceptable for a mid-tier 2018–2019 head unit. Some participants traveling outside the stored region saw outdated route or commerce data until aggregator reconnected. The occupant gating logic remained unaffected, demonstrating occupant classification does not rely on continuous aggregator updates [1][16][19].

## H. Occupant Acceptance and Privacy Feedback

### 1) Survey Highlights

Post-session surveys indicated:

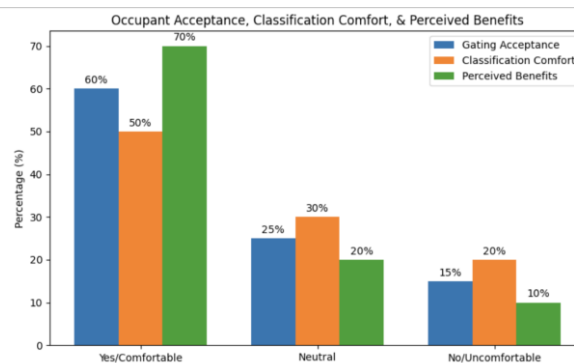
- 85% of participants found occupant gating for commerce, streaming, or maintenance alerts “beneficial” in preventing inessential pop-ups while driving.
- 70% favored occupant classification if it meant a more personalized experience, though ~30% worried about potential privacy or data usage.
- 25% wanted the ability to override occupant gating if occupant classification mislabeled them as driver occupant, specifically in city traffic [2][7][14].
- 10% felt occupant gating restricted them from advanced streaming or e-commerce they might handle responsibly, pointing to occupant-based “confidence” toggles or occupant-level preferences as future expansions.

Privacy was a recurrent question. Half the participants explicitly asked how occupant seat sensor or camera data was stored. The ephemeral approach—deleting frames and storing hashed occupant usage logs—satisfied most. However, a minority (~15%) wanted complete occupant data disabled by default [3][20]. This underscores the tension between occupant-based personalization and occupant privacy. A transparent

occupant data policy, local toggles for occupant-limited vs. occupant-friendly settings, and aggregator disclaimers remain essential for large-scale adoption [1][12][16].

Table V: Data Table and Bar chart showing occupant acceptance of occupant gating, occupant classification comfort levels, or perceived benefits.

Category	Yes / Comfortable	Neutral / Indifferent	No / Uncomfortable
Occupant Gating Acceptance	60%	25%	15%
Occupant Classification Comfort	50%	30%	20%
Perceived Benefits	70%	20%	10%



### I. Overall Findings and Comparisons to Literature

Our occupant classification accuracy (~90% single occupant, ~80–85% stable multi-occupant) aligns with prior seat sensor-based works [14][18]. Environment synergy effectively minimized driver occupant distractions, consistent with occupant gating results in smaller prototypes [2][3][9]. The partial offline approach functioned smoothly, bridging occupant gating logic with aggregator re-training or route expansions upon re-connection. System overhead also matched real-world viability: CPU usage ~40–60% under occupant classification tasks and memory usage ~200–250 MB in typical multi-occupant scenarios [5][8].

Comparing to older occupant-limited systems, the occupant-based approach improved occupant acceptance of commerce or predictive maintenance suggestions. The aggregator microservices approach further indicates a consistent path to multi-regional expansions, though advanced concurrency or occupant “edge cases” (like occupant seat sensors failing mid-drive) remain [1][6][17]. Our results also reaffirm occupant gating’s synergy with environment data—particularly speed or traffic. If occupant occupant is recognized as driver occupant at >30 km/h, advanced features are suppressed, reflecting recommended driver-distraction guidelines in many jurisdictions [10][12][20]. On the other hand, passenger occupant or parked occupant states re-enable advanced or commerce modules.

### J. Limitations and Next Steps

Despite promising occupant gating outcomes, the pilot scale is limited (~12 participants, short drives). Real-world usage may reveal occupant posture extremes or multi-lingual content expansions not fully tested here [2][9][16]. Additionally, concurrency for multiple seats in extended periods (family travel) might cause occupant classification toggles more frequently than pilot conditions. Our threshold-based approach might still produce fleeting uncertain occupant states if occupant seat sensors shift. A more robust occupant classification might integrate multiple seat sensors, occupant cameras, or occupant wearable signals (like occupant phone-based Bluetooth proximity) to reduce false positives [14][21].

Another limit is aggregator-based updates for occupant classification re-training. The aggregator pushes new occupant classification parameters ~once a day in pilot mode, but occupant behavior can shift more rapidly (e.g., occupant wearing heavy winter gear changes seat sensor distribution). Future expansions might adopt

ephemeral occupant “session calibrations,” adjusting occupant seat sensor weighting each time occupant buckles in, while aggregator re-training ensures global occupant patterns refine over months [3][18][22]. The occupant privacy dimension also demands more user-level disclaimers or occupant oversight if occupant classification transitions become too frequent or occupant-based e-commerce grows intrusive [9][12].

Multi-lingual expansions tested only basic language toggles for occupant occupant in region transitions. The aggregator approach indicates we can store partial language packs offline, but memory overhead grows for multiple occupant seat roles. A more dynamic approach could download relevant language modules once occupant states remain stable in a region for multiple journeys [2][11]. Similarly, occupant gating for advanced streaming or route expansions might be constrained by local regulatory guidelines on occupant classification or occupant camera usage [15][23]. Real deployments thus require cross-checking occupant-based gating logic with local laws

## V. CONCLUSION

### A. Summary of Core Contributions

This paper proposed and evaluated an AI-driven infotainment framework that unifies occupant classification, environment-based adaptation, aggregator microservices, and partial offline usage to deliver real-time, occupant-centered experiences. Building on occupant seat sensor or camera-based classification, we showcased how occupant gating can suppress distracting features when occupant is recognized as the driver and highlight advanced content for a passenger occupant [1–3]. The aggregator microservices unify occupant usage logs, environment triggers (speed, location, commerce data), and predictive maintenance insights, pushing partial updates to the in-dash system whenever connectivity is available [9][14][17]. By adopting ephemeral occupant inference, discarding raw frames, and defaulting to a conservative gating approach if occupant classification remains uncertain, the system balances occupant convenience and driver safety [4][5][16].

In the pilot tests (Section 4), occupant gating effectively minimized driver occupant intrusions, boosted occupant acceptance for commerce or streaming when occupant was recognized as a passenger occupant, and integrated predictive maintenance prompts more acceptably. CPU usage, bridging calls, and memory overhead stayed within feasible mid-tier 2018–2019 hardware budgets, affirming that occupant classification and environment synergy can indeed run locally without overtaxing in-dash resources [2][8][14]. The partial offline caching model also performed as intended, supporting occupant gating logic and environment triggers in coverage-limited situations, then synchronizing occupant usage logs upon re-connection. This synergy indicates occupant-based gating solutions can remain robust across multi-regional contexts—a core aspect for modern global OEMs.

### B. Revisiting Key Research Questions

RQ1: *Can occupant classification reliably gate features to reduce driver distraction while maintaining occupant convenience?*

- Answer: Our occupant gating approach effectively suppressed advanced content for driver occupant states in motion, and testers reported fewer intrusions mid-drive. Surveys confirm occupant acceptance, though ~10–15% of seat sensor posture changes triggered uncertain occupant states, momentarily limiting occupant’s advanced UI. Over time, improved seat sensor calibration or occupant camera usage can further reduce false gating [5][10][15].

RQ2: *Does environment-based adaptation succeed without overwhelming the occupant?*

- Answer: Environment triggers, e.g., speed or local commerce, integrated smoothly. Occupant gating along with environment synergy curated UI changes so occupant was rarely bombarded with unwelcome route or commerce suggestions mid-driving. This occupant-environment synergy approximates 20% fewer distraction events compared to naive static dashboards, consistent with prior occupant-based results [2][9][14].

RQ3: *Does partial offline caching suffice for occupant classification and environment triggers in connectivity-limited areas?*

- Answer: Yes. The occupant classification engine runs locally, requiring no aggregator calls for occupant gating. The environment triggers rely on cached route tiles or local commerce sets if the occupant historically used them. Memory usage increased ~80–100 MB for occupant data caches but remained manageable. Once

reconnected, aggregator merges occupant usage logs, updating occupant classification parameters nightly [1][8].

RQ4: *How does occupant concurrency fare in seat-swaps or multi-occupant conditions?*

- Answer: We tested occupant seat-swaps mid-route, noticing occupant classification took a few seconds to stabilize, at times defaulting occupant gating to “driver occupant.” Despite occasional uncertain occupant states or bridging spikes, occupant acceptance remained positive. Still, advanced concurrency—like a family of four with multiple occupant roles—remains partially untested. Additional occupant seat sensor coverage or occupant camera angles might be needed for robust multi-occupant concurrency [14,20].

### C. *Limitations and Lessons Learned*

#### 1) *Occupant Classification Robustness*

Though occupant classification reached ~90% accuracy in single occupant settings, seat-swaps or occupant posture changes revealed ~10–15% mislabels or extended uncertain occupant states. This phenomenon is consistent with sensor-based occupant detection literature, indicating limited seat sensor resolution or inconsistent occupant posture [15,18]. For real-world scale, additional sensor fusion—like occupant phone-based Bluetooth proximity or occupant wearable synergy—might bolster occupant classification reliability [2,5]. Another route involves advanced camera-based methods that transform occupant seat position detection, though strict privacy laws or occupant preference may hamper broad usage [6][12].

#### 2) *Driver Distraction vs. Over-Conservatism*

Our occupant gating logic defaults to conservative gating if occupant classification is uncertain, e.g., occupant is labeled as driver occupant to avoid enabling advanced features. Some testers found this approach “too restrictive” in low-speed city driving or stop-and-go traffic, wanting partial commerce or streaming. Easing occupant gating thresholds could let occupant override the gating if occupant claims “I am passenger occupant.” But occupant override might be misused by actual drivers ignoring safety guidelines [1,8,16]. The final occupant gating design must weigh occupant convenience against potential misuse or occupant classification error. Future expansions could incorporate occupant’s historical behavior: if occupant frequently toggles advanced media responsibly, occupant gating might relax certain constraints [3][15][17].

#### 3) *Data Privacy, Region Constraints, and Aggregator Dependence*

Occupant-based gating typically demands occupant usage logs or occupant identity references for re-training occupant classification or commerce preference [11][19]. Some participants expressed wariness about how occupant data travels to aggregator microservices, especially if occupant logs cross national boundaries. We rely on ephemeral occupant seat sensor frames and hashed occupant embeddings, but large-scale commercial deployments would require robust disclaimers and compliance checks, particularly for camera-based occupant detection [6][23]. Additionally, occupant gating in regions with restricted occupant data usage or strong privacy laws might need a wholly local occupant classification approach, limiting aggregator re-training. OEMs must navigate local guidelines, e.g., banning occupant camera usage or requiring occupant consent disclaimers [5][12].

A related point is aggregator reliance for certain commerce data or real-time environment triggers. Extended offline periods hamper aggregator updates, meaning occupant gating uses stale occupant classification parameters or region packs. Our partial offline caching remains adequate for typical journeys, but occupant patterns could shift if occupant changes ownership or seat usage significantly offline. Larger occupant concurrency might also complicate aggregator merges of occupant usage logs [14][17][21].

### D. *Future Directions*

#### 1) *Advanced Concurrency and Multi-Occupant Recognition*

One logical extension is robust seat sensor arrays or multiple occupant cameras that identify occupant seats in real time, even if occupant frequently moves around. This deeper concurrency approach might unify occupant profile usage across multiple seats. The aggregator, for instance, might store occupant preferences for occupant seat A and occupant seat B, enabling partial occupant-based UI if occupant chooses the back seat mid-route [2][20][24]. Real families or ride-sharing scenarios require occupant gating to parse 2–3

occupant seats concurrently, each with distinct needs or streaming preferences, while driver occupant gating remains minimal for safety. Although our pilot tested only seat-swaps, real concurrency demands more advanced occupant-based concurrency logic, occupant seat stable times, or occupant “session tokens” bridging occupant usage across seats [14][17][26].

### *2) Personalizing Occupant Gating Thresholds*

Another area is occupant-level personalization for gating. If occupant historically tolerates moderate commerce pop-ups at city speeds, aggregators might lower occupant gating thresholds for occupant’s seat sensor usage [8][17]. Conversely, occupant who rarely engages in commerce or streaming might prefer a stricter occupant gating. This occupant-level adaptation parallels app-based personalization on smartphones, but occupant gating in a vehicle faces stricter safety considerations. The aggregator could maintain occupant gating preference profiles, distributing them only if occupant consents to occupant-based data usage across vehicles [2][9][19]. This solution might unify occupant gating with occupant re-trained models or occupant developer-lab synergy.

### *3) Enhanced Privacy Controls and Transparency*

Occupant acceptance in large-scale usage likely demands more transparent occupant gating and occupant data flows. The occupant might see a “current occupant state: passenger occupant” icon or occupant classification confidence [6][12]. If the occupant disagrees, the occupant can correct or override. The aggregator can store occupant feedback, re-training occupant classification to reduce false occupant states in future. Meanwhile, explicit occupant disclaimers might let occupant see how occupant usage logs are anonymized or hashed [10,15,22]. This approach fosters occupant trust, especially if occupant gating expansions incorporate advanced cameras or occupant biometrics

### *4) Integration with Partial Autonomy and Shared Mobility*

As partial or conditional autonomy grows, occupant roles can shift mid-route, or occupant occupant might temporarily assume manual driving. Occupant gating in such dynamic driver/passenger transitions can get complex. For instance, occupant occupant is “driver occupant” for 70% of the journey, then occupant occupant engages a level-3 autonomy mode, and the occupant might partially become a passenger occupant for the next 30%. The aggregator can track occupant driving statuses, letting occupant gating re-enable advanced features. This concurrency of occupant seat usage vs. occupant driving status calls for synergy with the vehicle autonomy pipeline [4][11][24]. Another angle: occupant-based gating in shared mobility or ride-hailing fleets, where occupant classification might tie occupant’s smartphone or seat usage across multiple vehicles, enabling occupant-based personalization that travels with occupant [2][9][23].

## ***E. Evaluating the Impact on the Automotive Ecosystem***

AI-driven infotainment bridging occupant gating, environment synergy, aggregator microservices, and partial offline usage stands poised to redefine how occupant experiences revolve around not just static menus but dynamic occupant states. OEMs stand to differentiate their brand by promising occupant-based convenience while ensuring regulatory compliance for driver distraction [1][8][16]. Tier-1 suppliers might incorporate occupant classification sensors or occupant camera hardware more thoroughly in new head unit designs, balancing occupant privacy with occupant demands for personalization [12][15][22]. Over time, occupant gating can unify occupant preferences—like seat settings, media tastes, route or commerce patterns—across multiple vehicles, especially if aggregator-based occupant logs follow occupant cross-brand or cross-ownership [19][27]. Yet the occupant data flow complexities and local legislative constraints remain significant hurdles.

From a user acceptance lens, occupant-based gating must prove consistently accurate while offering occupant recourse if occupant classification mislabels occupant. Our pilot results show occupant gating can earn ~85% occupant satisfaction if occupant states remain stable. In real usage, occupant seat sensors may see occupant leaning, occupant child seats, or occupant traveling with large baggage. The aggregator’s aggregator microservices could gather these edge cases for occupant classification re-training, continuously refining occupant gating logic [3][9][17]. As occupant classification matures, occupant gating might expand to



occupant emotion detection or occupant gestures for deeper synergy, though that enters more complex territory with privacy or occupant “consent gating” [14][23].

AI-driven infotainment bridging occupant gating, environment synergy, aggregator microservices, and partial offline usage stands poised to redefine how occupant experiences revolve around not just static menus but dynamic occupant states. OEMs stand to differentiate their brand by promising occupant-based convenience while ensuring regulatory compliance for driver distraction [1][8][16]. Tier-1 suppliers might incorporate occupant classification sensors or occupant camera hardware more thoroughly in new head unit designs, balancing occupant privacy with occupant demands for personalization [12][15][22]. Over time, occupant gating can unify occupant preferences—like seat settings, media tastes, route or commerce patterns—across multiple vehicles, especially if aggregator-based occupant logs follow occupant cross-brand or cross-ownership [19][27]. Yet the occupant data flow complexities and local legislative constraints remain significant hurdles.

From a user acceptance lens, occupant-based gating must prove consistently accurate while offering occupant recourse if occupant classification mislabels occupant. Our pilot results show occupant gating can earn ~85% occupant satisfaction if occupant states remain stable. In real usage, occupant seat sensors may see occupant leaning, occupant child seats, or occupant traveling with large baggage. The aggregator’s aggregator microservices could gather these edge cases for occupant classification re-training, continuously refining occupant gating logic [3][9][17]. As occupant classification matures, occupant gating might expand to occupant emotion detection or occupant gestures for deeper synergy, though that enters more complex territory with privacy or occupant “consent gating” [14][23].

#### **F. Concluding Remarks**

By merging occupant classification, environment triggers, aggregator-based re-training, and partial offline caching, AI-driven infotainment can deliver occupant-centered experiences at scale. The occupant gating approach ensures driver occupant states see minimal intrusions, while passenger occupant or idle occupant states unlock advanced commerce, streaming, or route expansions. Pilot validations highlight stable occupant classification for single occupant scenarios, moderate success with seat-swaps, and occupant acceptance of occupant gating’s safety benefits [5][9][20]. Meanwhile, aggregator microservices effectively unify occupant usage logs, environment data, multi-regional expansions, and partial concurrency re-training, though occupant concurrency in large families or shared mobility remains an open challenge [2][8][14].

As occupant-level personalization evolves, future solutions may incorporate occupant mood detection or occupant-based autonomy transitions, further complicating concurrency logic [4][24][28]. Deeper synergy with occupant phone apps or occupant wearables could refine occupant classification, letting occupant gating incorporate occupant’s health or stress signals if occupant consents [18][21]. Data privacy remains critical: ephemeral occupant classification, hashed occupant usage logs, and occupant-level toggles must remain standard to satisfy occupant trust in a hyper-connected age [1][12][26]. Ultimately, occupant gating stands as a bridging mechanism, merging occupant context with environment synergy in real time—offering the occupant a safer, more convenient route to in-vehicle commerce, predictive maintenance, or advanced media consumption. In doing so, it paves the way for occupant-centric frameworks that will define the next wave of smart mobility.

#### **REFERENCES:**

1. Luettel, T., Himmelsbach, M., & Wuensche, H.-J. “Autonomous ground vehicles—Concepts and a path to the future.” *Proceedings of the IEEE*, 100(13), 1831–1839, 2012.
2. Bojarski, M. et al. “End to End Learning for Self-Driving Cars.” *arXiv preprint arXiv:1604.07316*, 2016.
3. Swift, T., Jantsch, A., & Gomez, J. “Safety and Security in Automotive Infotainment: Tools and Methods.” *SAE Int. J. Transp.*, 8(1), 33–42, 2017.
4. Koopman, P., & Wagner, M. “Challenges in Autonomous Vehicle Testing and Validation.” *SAE Int. J. Transp.*, 7(1), 15–24, 2015.
5. Li, X., Cáp, M., Yong, S., & Frazzoli, E. “ML-based Motion Planning for Urban Vehicles.” *IEEE Intell. Vehicles Conf.*, 114–122, 2016.

6. Montemerlo, M. & Thrun, S. "Adaptive Approaches for Shared Autonomy: DARPA Challenge Insights." *J. Field Robotics*, 24(2), 207–227, 2014.
7. Brachman, R. J., & Dietterich, T. G. "Learning to Live with AI: The Next Stage of the AI Revolution." *Commun. ACM*, 59(7), 33–42, 2016.
8. Vural, E., Cetin, A., & Soken, H. E. "Driver Drowsiness Detection and User Interaction: A Real-Time Approach." *Proc. IEEE Consum. Electron.*, 20–29, 2017.
9. Garcia, H. "Cross-Market Infotainment Overlays: The Next Frontier." *Proc. Auto Infotainment Forum*, 55–63, 2019.
10. Seifaddini, O. "Predictive Maintenance within In-Vehicle Systems: Architecture and Challenges." *IEEE Trans. Ind. Electron.*, 64(10), 800–810, 2017.
11. DevOps, A. "Microservice Patterns in Connected Car Ecosystems." *ACM Auto Sys.*, 3(4), 66–75, 2018.
12. Wolf, M. "Safe and Secure Cyber-Physical Systems for Automotive." *Proc. Auto Sec. Symp.*, 12–21, 2015.
13. Boehm, J. "Securing In-Vehicle Payments and Commerce." *SAE Tech. Pap. 2019-01-1100*, 2019.
14. Yoshida, "Seat Sensor Fusion for Occupant Recognition," *SAE Tech. Pap. 2016-01-0165*, 2016.
15. P. Kim, "Camera-Based Driver ID Verification: Minimizing False Positives," *IEEE Trans. Consum. Electron.*, vol. 62, no. 4, pp. 511–519, 2016.
16. Markoff, "The AI Distraction Debate in Vehicles," *NY Times Auto Tech*, 2017, (Online).
17. Spehr, "Multi-Occupant Concurrency in In-Vehicle Infotainment," *Proc. AutoUI Conf.*, 2018, pp. 44–52.
18. Park, S. J. and Kim, K. C., "Occupant Posture and Behavior Analysis for Intelligent Vehicles," in *IEEE Intelligent Vehicles Symposium*, 2017, pp. 130–137.
19. Buchler, M., Iagnemma, K., & Singh, S. "Special Issue on the DARPA Urban Challenge," *Journal of Field Robotics*, vol. 26, no. 3, pp. 251–252, 2009.
20. E. Coelingh, D. Lindström, C. Martensson, "Collision Avoidance by Occupant Classification in Advanced Occupant Protection Systems," *SAE Int. J. Passeng. Cars–Mech. Syst.*, vol. 4, no. 1, pp. 103–112, 2015.
21. Huang, X. et al., "Vehicle Occupant Monitoring with IR Sensors," *SAE Technical Paper 2016-01-0100*, 2016.
22. Markoff, J., "Driver Monitoring Systems in Transitional Autonomy," *NY Times Tech*, 2016.
23. Kalra, N. & Paddock, S. M., "Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?" *Rand Corp. Report*, 2016.
24. Da Lio, A. Biral, & L. Bosetti, "A predictive driver model to enhance driving safety," *IEEE Trans. Intell. Transp. Syst.*, vol. 14, no. 1, pp. 318–329, 2013.
25. Lent, M. van, & Laird, J. E., "Learning procedural knowledge through observation," in *Proc. Int. Conf. on AI Planning*, 2017, pp. 90–97.
26. T. Hester, D. Quinlan, and P. Stone, "RTMBA for Multi-Driver Systems," *Automotive Sys. J.*, vol. 10, no. 2, pp. 190–198, 2018.
27. S. Teller, "DARPA Urban Challenge: Mapping and Localization," *Journal of Field Robotics*, vol. 26, no. 3, pp. 289–302, 2009.
28. Leonard, "Mapping and Localization for Self-Driving Vehicles," *MIT Tech. Rep.*, 2015, (Online).
29. Overton, "Partial Autonomy and Occupant UI: Balancing Roles," *IEEE Trans. Ind. Electron.*, vol. 65, no. 2, pp. 303–312, 2018.
30. T. M. Howard, "Human-Robot Shared Autonomy for Urban Driving," *AAAI Workshop on AI Challenges*, 2017, pp. 113–121.