Transformational Autonomy and Personal Transportation: Synergies and Differences between Cars and Planes

Kenneth H. Goodrich[1]
NASA Langley Research Center, Hampton, VA, 23681

Jim Nickolaou[2]
General Motors Technical Center, Warren, MI, 48090

Mark D. Moore[3]
NASA Langley Research Center, Hampton, VA, 23681

Highly automated cars have undergone tremendous investment and progress over the past ten years with speculation about fully-driverless cars within the foreseeable, or even near future, becoming common. If a driverless future is realized, what might be the impact on personal aviation? Would self-piloting airplanes be a relatively simple spin-off, possibly making travel by personal aircraft also commonplace? What if the technology for completely removing human drivers turns out to be further in the future rather than sooner; would such a delay suggest that transformational personal aviation is also somewhere over the horizon or can transformation be achieved with less than full automation? This paper presents a preliminary exploration of these questions by comparing the operational, functional, and implementation requirements and constraints of cars and small aircraft for on-demand mobility. In general, we predict that the mission management and perception requirements of self-piloting aircraft differ significantly from self-driving cars and requires the development of aviation specific technologies. We also predict that the highly-reliable control and system automation technology developed for conditionally and highly automated cars can have a significant beneficial effect on personal aviation, even if full automation is not immediately feasible.

I.   Introduction

Over the past ten years, the projected technical feasibility of practical, highly-automated cars with the capacity to autonomously operate under all situations and conditions that can be managed by human drivers has gone from largely unthinkable to something that is considered likely in ten to twenty years by many experts[1]. While this capability would transform how we own and operate cars on the ground, it is also important to consider how the underlying technologies and capabilities might be applied to other domains such as aviation. This paper explores the applicability of driverless car technologies to small, passenger-carrying aircraft that are commonly described as general aviation aircraft. These aircraft are typically operated by the travelers themselves (i.e., private piloting) or are flown by a professional pilot in the context of on-demand charter or air taxi-operations. Specifically, the scope of aircraft and operations considered in this paper is nominally up to 6 passengers (i.e., certification Level-2 per the FAA's recently proposed update to small aircraft certification standards[2]) and personal or small commercial operations (i.e., Part 135, on-demand). The terms personal aircraft, airplanes, and aviation are used in this paper to denote this size and type of operations.

As becoming and remaining a safe, proficient private pilot requires an on-going commitment of time and money beyond what is practical for most people and the direct and indirect impacts of needing professional pilots for air-taxi operations reduces the availability, flexibility, and cost-competitiveness of these operations[3], self-piloting personal aircraft could transform the utility of general aviation as a practical means of high-speed, on-demand mobility for many people. Further, when synergistically combined with other emerging technologies such as distributed-electric

---

[1] Aerospace Engineer, Dynamic Systems and Control Branch, MS 308, Senior Member AIAA.
[2] Sr. Technical Lead, Autonomous Driving Advanced Development.
[3] Aerospace Engineer, Aeronautics Systems Analysis Branch, MS 442, Senior Member AIAA

propulsion (DEP), the community and environmental impacts of small aircraft use are dramatically improved compared to current GA airplanes and cars[3]. Also, DEP may enable advanced configurations and operations such as quiet, vertical take-off and landing (VTOL) aircraft that are practical for relatively short-range trips across otherwise congested urban and suburban areas with limited space for additional road infrastructure[4]. While not discussed in detail in this paper, it should be considered that networked, self-driving cars (or even current cars coordinated by ride-sharing applications) strongly complement personal aircraft to provide door-to-door transportation capabilities without the complexity and potential penalties of multi-mode or convertible vehicles that must drive and fly.

This paper reviews the status, operational requirements, and emerging technologies underlying self-driving cars and compares these to what is likely required to enable self-piloting GA airplanes. The paper also considers what the aviation impacts might be if self-driving car technology does not fully eliminate the need for human oversight and possible intervention in all situations. Rather than self-piloting airplanes, this level of technology might enable simplified piloting with a level of training similar to learning to drive today. We first review the current and emerging future capabilities of automated cars using an automation-level taxonomy developed by the automotive industry. We then present a generalized decision and control process hierarchy applicable to all transportation vehicles. Next, we use this hierarchy to compare the domain-specific requirements of self-driving cars and self-piloting aircraft. This is followed by discussion of the application of conditional self-driving and operator assist technologies in the aviation domain. Finally, the paper concludes with a summary of the synergies and differences between self-driving cars and airplanes and identifies common needs and opportunities.

II. Overview of Current and Future Automated Cars

The capabilities of automated on-road vehicles can be characterized according to a scale developed by the Society of Automotive Engineers (SAE) in reference[5]. As summarized in figure 1, the SAE taxonomy consists of 6 levels: Level-0 (no automation); Level-1 (driver assistance); Level-2 (partial automation); Level-3 (conditional automation); Level-4 (high automation); and Level-5 (full automation). It is important to recognize that these levels only address the, "dynamic driving task", which is summarized in the SAE report as, "all of the real-time functions required to operate a vehicle in on-road traffic, excluding the selection of destinations and waypoints (i.e., navigation or route

SAE level	Name	Narrative Definition	Execution of Steering and Acceleration/ Deceleration	Monitoring of Driving Environment	Fallback Performance of Dynamic Driving Task	System Capability (Driving Modes)
Human driver monitors the driving environment						
0	No Automation	the full-time performance by the *human driver* of all aspects of the *dynamic driving task*, even when enhanced by warning or intervention systems	Human driver	Human driver	Human driver	n/a
1	Driver Assistance	the *driving mode*-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	Human driver and system	Human driver	Human driver	Some driving modes
2	Partial Automation	the *driving mode*-specific execution by one or more driver assistance systems of both steering and acceleration/ deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	System	Human driver	Human driver	Some driving modes
Automated driving system ("system") monitors the driving environment						
3	Conditional Automation	the *driving mode*-specific performance by an *automated driving system* of all aspects of the dynamic driving task with the expectation that the *human driver* will respond appropriately to a *request to intervene*	System	System	Human driver	Some driving modes
4	High Automation	the *driving mode*-specific performance by an automated driving system of all aspects of the *dynamic driving task*, even if a *human driver* does not respond appropriately to a *request to intervene*	System	System	System	Some driving modes
5	Full Automation	the full-time performance by an *automated driving system* of all aspects of the *dynamic driving task* under all roadway and environmental conditions that can be managed by a *human driver*	System	System	System	All driving modes

Figure 1 SAE On-Road Automation Taxonomy

Figure 2 Example of Driver Monitoring Diligence

planning)." The dynamic driving task nominally encompasses object and event detection, recognition, and classification; object and event response; maneuver planning; steering, turning, lane keeping, and lane changing; acceleration and deceleration; and enhancing conspicuity (lighting, signaling and gesturing, etc.). Performing these functions generally encompasses perceiving and maintaining an assured, clear distance ahead for basic safety and planning and performing maneuvers as needed to follow the roadway, obey traffic control signals, and manage lane positioning to optimize flow within traffic and facilitate upcoming turns and other actions specified in a route or mission plan. While route planning is not considered part of the dynamic driving task, it is an essential function for automated vehicles and will be considered in more detail in following sections.

It should also be noted that the automation levels are evaluated within specific "driving modes" or use cases. Use cases can include strictly low-speed operations such as parking or driving in traffic jams. For the purposes of this paper, our primary interest is on-road driving modes applicable to normal driving at posted speeds or speeds appropriate for conditions.

Distinguishing characteristics automation Levels-0 and 1 (manual and assisted driving) include low-level, embedded automation that is often taken for granted, for example anti-lock brakes, electronic stability control, electronic engine controls, and even in a sense, automatic transmissions. A more overt example of Level-1 automation is cruise control, including adaptive cruise control capable of maintaining a set time or distance behind leading traffic. At Level-2 (partial automation) the driver can engage integrated speed and steering automation that works in unison to perform most of the physical activities of driving including lane-following and speed management considering lane curvature and traffic immediately ahead. This level of automation is similar to coupled autopilots used on airplanes for decades in that the automation can follow programmed flight legs or routes and relieve the operator from performing a high-attention, perceptual-motor control task for long-periods of time but has insufficient situation awareness to relieve the operator from full-time monitoring to ensure the route is hazard free and to determine that aspects of the autopilot's operation (e.g., active modes and targets) are safe and appropriate for current conditions. With Level-2 automation, the driver is expected to continuously monitor the overall situation (e.g., figure 2) and manage or deactivate the automation as needed. While physical workload and fatigue may be reduced, the monitoring, management, and intervention responsibilities retained by the driver means that safety depends to a very high degree on their vigilance, judgment, and manual driving skills. Vehicles with Level-2 automation are now becoming commercially available, with Tesla Motor's "Autopilot" being an example[6].

The step from Level-2 to Level-3 automation entails a significant increase in vehicle capabilities and perhaps more notably, responsibility. In particular, in an allowed set of driving modes or uses cases, the automation becomes responsible for monitoring the dynamic driving task and assessing whether it can safely continue self-driving or is sufficiently uncertain that handing control back to the human driver is required. Further, a request to hand-back control must be done with sufficient lead-time that the driver, who can be immersed in non-driving activities, can safely take back control. Recent research suggests this lead-time will be on the order of 5 seconds for highway operations[7]. Currently, no commercially available cars achieve Level-3 for anything other than low-speed use cases such as automated parking.

At Level-4 (high automation), the automation is expected to monitor conditions and autonomously react to maintain safety, without human oversight or intervention, should conditions change such that it has, or more conservatively is at risk of, exceeding its design conditions. At Level-4, if complete trips can be planned within the supported set of use cases, it would be acceptable to initiate these trips without a human driver on-board. Depending on the design conservatism of the automation, the probability of trip completion may be reduced as compared to when a human back-up driver is present, but this would nominally result in reduced operational reliability and utility but would not affect safety. Initially, Level-4 automation for at-speed driving will likely be limited to simpler use cases such as freeway driving in non-hazardous weather. As shown in figure 3 from Reference 1, according to a survey conducted

American Institute of Aeronautics and Astronautics

Q9: Automated freeway driving under most conditions, with the ability to ensure safety even if the driver is incapacitated?
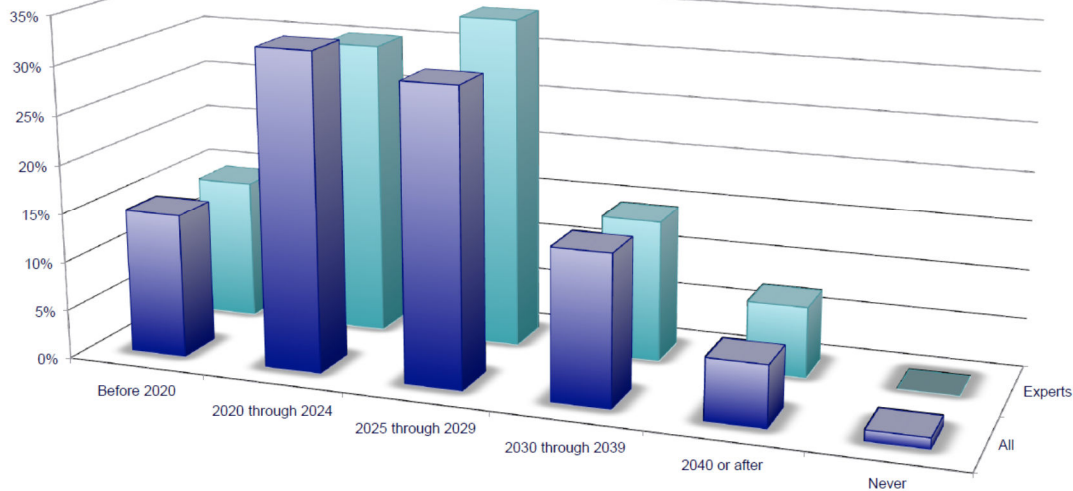
Figure 3 Automated Vehicle Forecast Survey for Level-4 Freeway Driving, Reference 1

at the 2014 International Automated Vehicle Symposium, a majority of experts predict Level-4 capabilities for freeway use will available within 15 years.

The final level of the scale, Level-5 (full automation), is defined as automation which has equivalent operational breadth and capability as human driver in all conditions. If no back-up driver is required, a manual driver interface (e.g., steering wheel, accelerator and brake pedals) becomes unnecessary and can be minimized or done away with, as reflected in concepts being explored or developed by companies like Mercedes-Benz as shown in figure 4[8] and Google[9]. It should be noted that Level-5 implies performing any operation a human driver can which is a very demanding requirement and probably unnecessary for the majority of daily operations. It is likely that predictions of fully-automated or autonomous cars mostly refer to broadly capable Level-4 cars.

## III.   Conceptual Decision and Control Hierarchy for Transport Vehicles

Before evaluating potential technology crossover from automated cars to personal aircraft, let us first consider a conceptual framework through which to assess and compare transportation-oriented vehicle operations. In this framework, vehicle operation is decomposed into a hierarchy of three simultaneous decision and control layers or loops as shown in figure 5. This hierarchy is similar to the mission, behavior, motion (planning), and (motion) control layers described in reference[10] with a degree of simplification obtained by merging the behavior selection and motion-planning layers into a single, maneuver planning loop. As presented, the framework is intended to be agnostic as to vehicle type and whether human or automation agents are used to perform the required tasks.

Figure 4 Mercedes-Benz Autonomous Concept Vehicle

American Institute of Aeronautics and Astronautics

Figure 5 Diagram of Mission, Maneuver, and Motion Loops Hierarchy
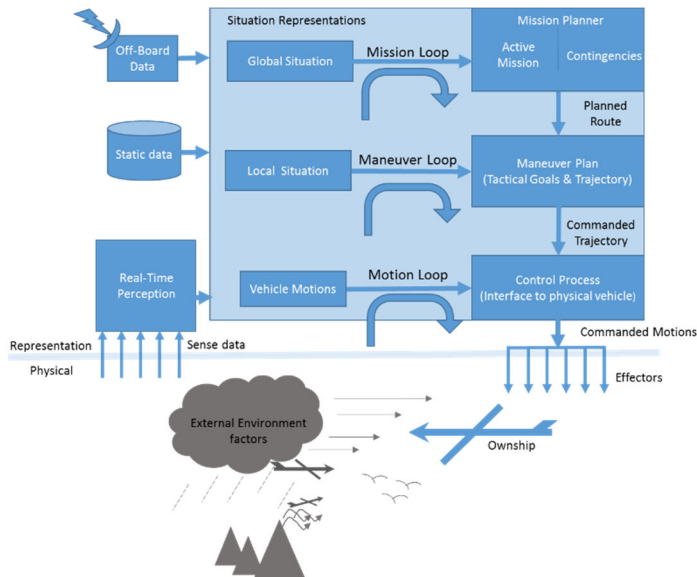
A. Mission-loop

The mission-loop is the outer-most process and involves formulating a long-term sequence of planned actions leading from the vehicle's present location to the desired destination. In addition to a nominal course of action and route, mission-management may also involve contingency plans and constraints anticipating, for example, the possibility of expected resources becoming unavailable (e.g., roads, airports); conditions such as weather being meaningfully different than forecast; or changes in the vehicle status that require a precautionary stop. Contingency planning is particularly important for vehicles like ships and aircraft that can travel long distances over which conditions may change and could otherwise be an unacceptable distance from a "safe harbor" if only the most direct route from to the destination is considered. If it becomes necessary to execute a contingency plan, the active destination and route plan are appropriately updated.

Several other functions within the mission loop, which admittedly might be placed elsewhere in the hierarchy, are communication with any external, controlling authority; load planning and actual loading; and assuring that the vehicle is operationally sound prior to initiating a trip. These functions are included under mission management because the tasks involved should either be completed prior to initiation of a trip (loading and pre-operational inspections) or are usually tolerant of significant latency (e.g., external communications relative to the route or clearance).

B. Maneuver-loop

The maneuver loop, or middle layer, is another planning process. It constructs awareness of the real-time, local situation (using inputs from a perception system briefly described below fused with stored information) and blends this with desired actions specified in the mission plan to determine both the vehicle's active, near-term goals and a corresponding planned trajectory sequence or maneuver to support the goals. This trajectory sequence must be consistent with the kinematic and dynamic capabilities of the vehicle in its current operating condition and situation and actionable by the motion loop. Depending on the sophistication and needed reliability of the automation, the maneuver loop might contain processes that maintain real-time models of the vehicles current operating capabilities to minimize any mismatch of planned versus achievable maneuvers. This modeling can also help facilitate fault detection and contribute to determining when a precautionary stop may be advisable (i.e., execute a contingency plan rather than continuing).

Although architecturally distinct from it, the maneuver-loop is functionally coupled to a real-time perception system. The perception system must sense and perceive relevant elements in the local environment with sufficient resolution (spatial and temporal) that the maneuver loop has the level-one situation awareness (SA)[11] inputs needed to characterize and project the overall situation as needed to identify and assess possible future states, hazards, constraints, goals, and actions.

C. Motion-loop

The motion-loop, the inner-most process of the hierarchy, manages vehicle's control effectors in response to physical measurements (e.g., accelerations, closure rates) and estimates of its dynamic state such that it accurately performs the planned maneuver sequences. Like the maneuver-loop, it is a core aspect of the dynamic driving task. Unlike the other two loops, it deals primarily with continuous kinematic and dynamic processes rather than decision making. In addition to controlling motions of the vehicle, other system management functions controlling the actual operation of vehicle systems reside in this loop. Minimum update rates within inner-loop processes are dictated by the dynamics of the controlled elements, but in general, the primary motion control loops should approximate

continuous systems. In addition to a high update rate, the motion-loop has very demanding safety requirements as loss of function essentially equates to loss of control.

IV. Comparison of Fully-Automated Car Capabilities and Personal Aviation Requirements

In this section, the potential contributions and limitations of fully-automated (FA) driving systems and the underlying capabilities needed for personal aircraft automation are considered using the hierarchy described in the previous section. This discussion is also relevant to the capabilities of Level-4 on-road vehicles with a breadth of use cases sufficient to complete many, but not all, trips a human driver could perform. For simplicity, we use the term FA to also include reference to these broadly capable Level-4 automobiles.

A. Mission Loop Comparison

Some of the most significant differences between automobile and aircraft capabilities needed for FA operations become apparent when considering the requirements and criticality of the mission loop. On the automotive side, implementations of the mission management process are sufficiently well understood that even basic requirements for it are not included in the SAE taxonomy for on-road automation. The relative maturity of this area for on-road applications has probably been experienced by most readers through reliance on smartphone navigation applications for turn-by-turn directions in an unfamiliar city.

While automatic route and waypoint generation is commonly used in flight planning, using it as the only source of pre- and in-flight planning information would be considered reckless and does not satisfy statutory preflight information requirements specified in the applicable FAA regulations (i.e., FAR 91.103). This regulation requires that as a minimum for cross country flights, the responsible agent (nominally the "pilot in command" but for FA aircraft this might be an automated agent or possibly an off-board human dispatcher) be familiar with: current and forecast weather; fuel requirements and alternative airports available if the planned flight cannot be completed; known traffic delays; runway lengths at airports of intended use; and a reliable means of determining takeoff and landing runway requirements considering expected values of airport elevation and runway slope, aircraft gross weight, and wind and temperature.

Several significant differences between road and flight operations account for the differences in acceptable mission planning rigor. Cars operate exclusively on a network of roads engineered to maintain a nominal level of risk and demands on the driver in virtually all conditions, except for the presence of infrequent weather hazards. In general these hazards are, with some relatively rare exceptions, easily observed and reasonably foreseeable. Further, the performance of modern cars is largely unaffected by non-hazardous weather so there is little need to consider the effects of, for example, winds along the route on attributes like safety, ride-quality, navigation, enroute time estimates. Finally, if for some unforeseen reason, the planned route does becomes unacceptably risky due to, for example, heavy rain, while the operational impact may be large (e.g., a missed appointment), the net safety impact is relatively small as simply stopping and waiting for conditions to improve is often an option. Going from a cruise condition to a reasonably safe, rest state on the side of the road typically takes well under a minute.
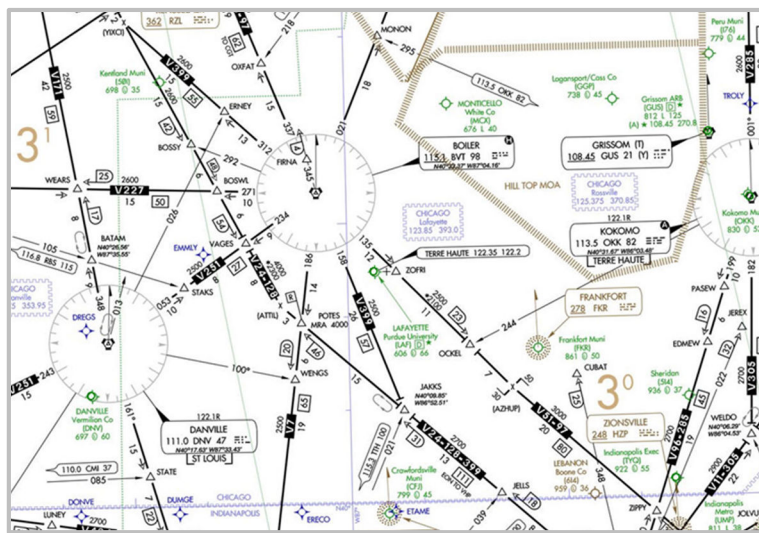

Figure 6 Aeronautical Chart Depicting Airways Over Indiana

In comparison, even though aircraft can operate along the aerial equivalent of roads (e.g., airways, figure 6) which are engineered to minimize some risks such as terrain hazards and radio navigation and communication blockages, the risks, piloting demands, and vehicle performance relative to the ground, are highly dependent on dynamic factors, primarily related to weather, and require more detailed trip planning than car trips. Factors included in the planning process typically include, to expand on the intent of FAR 91.103:

- Avoidance of hazardous weather
- Avoidance or minimization of uncomfortable weather (e.g., light or moderate turbulence)
- Stability and reliability of forecast conditions over the time and route of flight

- Altitude constraints relating to terrain and weather
- Forecast weather conditions at potential contingency airports along the route of flight
- Unique, local weather phenomena such as significant terrain and wind interactions
- Extended flight over inhospitable terrain (e.g., weather, mountains, wilderness) and possible survival equipment (e.g., inflatable vest for extended over-water flight)
- The effect of winds on airplane performance including speed, fuel-burn, and climb gradients relative to the ground
- Required and desired fuel reverses and fuel loading for potential contingencies
- Time of day considerations such as flight after sunset
- Routing constraints preferences dictated by the National Airspace System (NAS) and air traffic control (ATC).
- Potential for pilot fatigue considering pre-, in-, and post-flight activities

Monitoring of these factors and updating the mission plan should occur at regular intervals throughout a flight.

While automated and network-enabled flight planning tools, on the ground and in-flight, have made flight planning and mission monitoring much easier and integrated over the past ten years, little effort has been made to fully automate the process. While automotive route-planning is largely a rule or procedure based process, flight planning requires significant knowledge-based reasoning which has historically been more difficult to automate. Flight planning is a risk and uncertainty management process. Making informed flight plans and go/no-go or continue-divert decisions requires practical awareness of the risks and uncertainties involved and judgment as to the desirability of the trip and possible alternatives. These decisions typically involve factors not readily accessible to an automated system and require some degree of informed interaction with the traveler as might be done today by a charter pilot and his clients.

As mentioned earlier, in addition to route planning, the mission-loop includes responsibility for ensuring the fitness of the vehicle for the intended trip as well as proper loading. The relative importance and required detail of pre-flight verification of an airplane's airworthiness compared to a car should be obvious to most readers. While automated systems are likely to require fail-safe, built in checks, to assure safety, many basic vehicle properties, not normally observable by the automation, may also need to be checked prior to operation. In general, the inspection requirements of aircraft are much more demanding than cars. For example, as long as the sensors are not obscured (and these could be heated or otherwise protected) operating a FA car with a layer of ice on the body is likely to have minimal impact on performance and safety. In contrast, an airplane with a thin layer of ice or frost on the wings is not airworthy. Fully automating the variety of relatively simple but important visual and mechanical checks required to assure an aircraft is airworthy goes well beyond automotive requirements and would add significant cost and complexity.

The factors involved in ensuring proper loading of an aircraft versus a car may be unfamiliar to some readers but again, the relatively importance compared to ground operations is easy to intuitively appreciate. In addition to the importance of respecting gross weight limits, the location of an aircraft's center-of-gravity (CG) must fall within specified limits or its stability and control characteristics will be hazardously degraded. Currently, human pilots (or other crew members) are responsible for ensuring proper loading and use weight and balance charts or automated applications to determine the placement of passengers, cargo, and fuel is such that the CG remains within design limits throughout a flight. While the load verification process could be automated, like fully-automating the preflight with sufficient reliability that human involvement or understanding is no longer required even for monitoring purposes, the costs, complexity, and potential risks of doing so are not being addressed by FA cars.

The final function mentioned earlier within the mission-loop is communication. Unlike cars, aircraft often have to coordinate their plans and actions with off-board, traffic management agents such as ground, tower, terminal, and enroute controllers. The majority of this communication is done today via voice exchanges over sometimes noisy analog radios. While conceptually automating such communication, and more significantly the negotiation supported by it, may be feasible, it is a significant technical and regulatory challenge and is not an area being advanced in the automotive domain. The automotive community is developing and deploying Dedicated Short-Range Communication (DSRC) technologies and standards to facilitate vehicle-to-vehicle (V2V) and infrastructure-to-vehicle (I2V) communications for brief, safety messages between vehicles and infrastructure (e.g., traffic lights also transmitting their state). As described in reference[12], communication distances are targeted at hundreds of meters to 1 km, making it of limited utility for in-flight communication.

B. Maneuver Loop Comparison

Successfully automating the maneuver-loop is one of the major challenges and will be one of the major breakthroughs of FA cars. While on-road operation greatly simplifies mission planning as described in the last section, it also requires precise maneuvering relative to roadways, other cars, and a complex milieu of other static and dynamic situation

elements. Further, unlike many aviation systems, road markings, supporting signage, and rules of the road were developed without any consideration of support from technological systems.

The primary automotive tasks of interest in this discussion are on and near-road object and event detection, recognition, and classification followed by, from an own-vehicle perspective, identifying specific, short-term operational goals and constraints and ultimately developing a planned trajectory that maximizes progress in support of active goals while honoring constraints. In this discussion, it is important to recognize that how FA cars organize and integrate tasks within the maneuver loop is quite different than human drivers and achieving equivalent or better performance compared to human drivers does not necessarily require perception and comprehension that's equivalent to human capabilities or easily generalizable to other applications. Having different strengths and limitations compared to humans, practical automation agents typically do not replicate how humans perform any given task or process.

Considering first the requisite perception requirements that provide the foundation on which higher-level SA depends (e.g., situation analysis and planning), in the automotive domain, near vehicle awareness and response ranges from 0 to approximately two seconds ahead (e.g., ~68 yards at 70 mph), mid-level awareness to drive down the road safely and successfully extends from two to eight seconds ahead (68-273 yards at 70 mph) and further than eight seconds ahead is considered far awareness (i.e., >273 yards). To help aid the SA process, FA car prototypes currently under development typically use detailed databases of the road environment to reduce the need for and guide the perception[13]. Far awareness, at least awareness of what static features to expect, is largely obtained from stored data. This is sufficient, for example to anticipate turns without having to "see" them first. Real-time, sensor-based perception is mostly limited to within ~200 yards and primarily uses LIDAR and RADAR sensors. This maximum range is not determined so much by direct sensor range limitations, but rather the difficulty of reliably interpreting the lane-relative positions of more distant features when the road itself may curve and slope[14]. While vision-based systems also have to the potential to support longer-range perception, the image processing and computational challenges currently limit their application. Finally, in addition to reducing the need for long-range sensing for far awareness, on-board data is used to prime the perception system as to what it should be detecting. This significantly aids the object detection and identification process. It is also an important aspect of the high accuracy and reliability positioning needed for guidance and navigation by enabling precise localization relative to mapped reference features.

To compare these automotive capabilities to the requirements of personal aircraft, let us first consider the types of situation elements and hazards that need to be perceived. In up-and-away flight (i.e., well above local terrain), primary factors include other air traffic; birds singularly and in flocks; and weather phenomena. Unlike the "solid" hazards discussed so far, weather typically involves diffuse phenomena and determining interior hazard levels and boundaries is a challenge very different from the automotive domain. Examples of desired local weather perception capabilities include determining the location, height, and development stage of nearby clouds sufficient to assess potential hazards such as turbulence and hail, and to remain clear of them when appropriate; characteristic indicators of potential clear-air turbulence (e.g., lenticular clouds); rain shafts relative to cumulonimbus clouds and storm systems; and ice accumulation types and rates. Considering flight near or approaching the ground such as during approach and landing, in addition to the earlier items, low-altitude man-made threats like newly erected, uncharted towers or power-lines are considerations as are tall trees. Considering the runway environment itself during landing and takeoff, it becomes desirable to sense the runway outline as a double check on navigation and data base accuracy and integrity. At night, runway, taxiway, and other surface lights may be available to assist this process and convey additional information. In addition to the surface configuration, verifying that the runway is free of other aircraft, ground vehicles, people, and wildlife, is needed. Similar to on-road operations, the ability to see surface conditions such as large cracks, potholes, puddles, ice, snow, foreign objects, etc., is needed during surface operations.

The examples in this brief discussion are not intended to be complete, but rather to highlight how the perception challenges involved in FA personal aviation include a range of features very different than road vehicles. So far, nothing has been said about corresponding detection ranges and suggesting specific range requirements is beyond the scope of this paper. As a rough guide, the requirements for visual flight rules (VFR) provide an initial indication of what the perception range might be, at least for the detection of aircraft-like objects. The lowest visibility in which VFR operations by airplanes are allowed is one statute-mile. However, this distance is associated with a small portion of the NAS (i.e., class G airspace) with limited utility to enroute operations. For other airspace classes, the visibility requirement is 3 miles for altitude below 10,000'. Considering approach and landing flight phases, current instrument flight rules (IFR) suggest that perception of the runway environment from a minimum distance of ½ mile is needed. A ½ mile is the minimum visual perception distance associated with the final, visual segment of instrument approaches outside of highly specialized category II and III instrument landing system (ILS) operations. These specialized operations require a high degree of control and monitoring of the runway environment that is typically only available at large commercial airports and probably understates aircraft-based perception requirements in the absence of these

controls. Considering perception ranges associated with avoiding hazardous weather, a standard guideline[15]is to adjust the flight trajectory to avoid active thunderstorms by at least 20 miles. While off-board sources such as data linked Next-Generation Radar (i.e., NEXRAD) can help with this task, the total latency of the data presented in the cockpit can exceed 15 minutes, make it unsuitable for tactical weather avoidance[16].

Continuing within the maneuver loop, while the flight perception process has unique and challenging requirements beyond what is needed for FA cars, once the local situation elements are perceived and characterized, the goal and maneuver planning process is expected to be similar to and generally less demanding than on-road operations. Differences include much more forgiving tolerances (e.g., enroute airways are +/- 4 miles wide), other traffic is less common and usually much further away, and much of the airspace infrastructure (e.g., airways, instrument procedures, etc.) is defined by technological means, making it directly accessible to the automation. Compared to surface on-road operations, flight-based maneuver planning does introduce additional considerations such as the ability to maneuver vertically in addition to laterally and the need to anticipate a variety of wind conditions and effects. While the unique requirements and of flight will require additional, specialized development, the crossover of fundamental maneuver planning concepts and algorithms from FA cars is anticipated to be high.

C. Motion Loop Comparison

The motion loop controls the physical vehicle such that it follows the reference trajectory and action sequences commanded by the maneuver loop. Compared to the other loops, it involves more conventional feedback control elements, with reference signals derived relatively directly from parameter measurements (e.g., accelerometer and rate sensors). As such, the low-level elements and technologies of the motion loop between cars and aircraft will be relatively similar and most differences between the two applications will arise from differences inherent in the complexity of controlling the two types of vehicles and differences in required reliability and precision.

A preliminary comparison of the basic control requirements of cars and aircraft can be done with relatively simple kinematic and dynamic relationships providing insight into basic architectural requirements of the primary control-loops. For most routine driving situations on good roads and at posted speeds, the car control task can be modeled by two single-input, single-output control tasks: 1) a lateral control task using the steering wheel (really, the steering angle of the front tires) to control the turn radius (which is equivalent to lateral acceleration at constant speed), and 2) a longitudinal control task managed through longitudinal acceleration inputs made via engine power and brake commands and occasional mode changes using the transmission (e.g., forward and reverse). Except during very aggressive maneuvering which FA cars may not perform, the steering and speed control tasks are largely independent.

Considering a typical fixed-wing airplane, the control tasks consists of managing a coupled, multi-input, multi-output system. Pitch (elevator), roll (aileron), and yaw (rudder) moment effectors are used in conjunction with a longitudinal-power to manage airspeed, vertical speed, and turn rate during most phases of flight and airspeed, vertical speed, lateral translation, pitch angle, and heading angle during landing. Cross coupling of an airplane to control inputs is typically much greater than a car, and for example, the process of maneuvering an airplane quickly to a new heading involves making coordinated elevator, aileron, rudder, and throttle changes. Simply comparing the number of control inputs and controlled responses suggest that the airplane control task requires 2-2.5 times the resources (e.g., sensors, computation, actuators) as a car. So, while the technology will be essentially the same, based on task complexity considerations alone (i.e., exclusive of reliability and precision differences), the control loop of a personal airplane should be expected to be 2-2.5 times more complex than a car.

Considering the other attributes of precision and reliability and for the moment restricting the discussion to the primary motion control elements, complete loss of a degree of control in either domains is likely to result in loss of vehicle and high risk to the occupants and persons and property nearby. Considering this common requirement, FA cars are likely to require low-cost, secure, and redundant control elements, including redundant computing, sensing, actuation, and networking, that have high applicability to personal aircraft. Control precision requirements are also expected to be reasonably similar between the two domains. Since the size of the automotive industry is currently several orders of magnitude larger than the small aircraft industry, economies of scale and investment from FA car component elements are likely to be very beneficial to parallel aircraft applications.

The discussion above only addresses path control elements and the motion-loop includes other real-time systems management functions as well. Briefly considering these other systems operation and management functions, it is likely that compared to cars, aircraft will have a higher number of fail-operational versus fail-safe or fail-off systems. For example to avoid a single-point of failure and provide redundancy, the fuel supply of aircraft is often stored in multiple fuel tanks with a valve allowing independent access. While having separate tanks is in part driven by space and CG considerations, it also provides protection against, for example, a leak causing loss of all fuel or a loose fuel cap allowing rain water to contaminate all fuel on board, thus providing a fail-operational capability. Similar to the

path-related control elements in the motion loop, the more demanding reliability requirements of an FA aircraft systems are likely to add expense, but the basic technology from FA cars will be transferable.

D. Summary of Comparison

Briefly summarizing the discussion of this section, the mission-loop capabilities of FA cars are relatively simple compared to the needs of FA aircraft. Automating the knowledge-based flight-planning and other mission layer capabilities needed for FA flight (preflight inspection, load verification, ATC communication) requires significant, independent technology development. Within the maneuver-loop, the sensor and perception requirements for flight have important differences from road applications and enabling FA personal aircraft will require significant, unique technology development with an emphasis on characterizing local, but diffuse weather elements and hazards. The general operation of the goal and maneuver planning tasks are sufficiently similar that much of the fundamental technologies of FA cars will probably be extensible to aircraft. The development of low-cost, safety-critical automotive systems will likely be highly applicable to safety-critical trajectory control and other system management functions. That said, the greater functional complexity and required reliability of flight systems will generally result in more expensive sub-systems than cars, even when developed using similar components.

V. Personal Aircraft with Less Than Full Automation

From the comparisons in the previous section, even if FA cars become feasible and available in the near future, it is likely that significant additional technology development will be needed to enable FA aircraft. Another possibility is that for the next few decades, cars only achieve conditional and high-automation (e.g., Levels 3 and 4) for a limited set of enroute use cases such as highway driving, requiring licensed and fit drivers to complete most trips. This situation would inhibit many of the more transformational concepts suggested for self-driving cars such as dramatically reducing private ownership given the ability to summon unoccupied cars on demand. This section briefly explores the impact of these two possibilities on the potential capabilities, implementation, and usefulness of less-than fully automated personal aircraft.

At the start of this discussion, it is important to recognize that even if FA aircraft cannot be achieved for decades, 200 mph personal aircraft that are much easier to use compared to today (i.e., simplified training and reduced workload combined with increased safety and utility) could still be transformational if they are affordable and accessible to travelers who regularly need to travel distances between approximately 120 and 600 miles. These regional trips are poorly served by cars (too slow) and hub-and-spoke airlines (limited service, fixed schedules, high-cost for on-demand travel, and limited speed advantage going between spoke cities). Today, just under 470,000 people, or about 0.2% of the adult population of the United States, hold pilot's licenses (excluding student certificates). This compares to 210 million licensed automobile drivers. Clearly, there is plenty of opportunity to grow the pilot population if the costs and benefits of personal aviation are made more favorable.

Currently, becoming a proficient instrument rated pilot requires a significant investment of time and money, roughly $30,000 dollars and 150 or more hours of flight training and practice time. Remaining current and proficient regularly entails additional training time and money. Quite simply, outside of a career or serious hobby, earning a pilot's certificate is not justifiable for most people. But, there is plenty of opportunity to make it much more accessible by appropriately using technology to simplify the flight tasks and training without necessarily having to reduce them to zero through FA. But, it should be understood that this more lightly trained "operator" will be dependent on technology to perform entire tasks and assist others that a current-day pilot would be expected to perform. This dependence needs to be fully considered in assessing reliability requirements for the automation.

In addition to improving workload and training, improving safety is also an important consideration in improving the costs and benefits of personal aviation. Figure 7 shows a summary from the National Transportation Safety Board (NTSB) of the defining events for fatal and non-fatal personal flying accidents in 2012[17]. This figure highlights events and associated operations that are prime targets for simplification and mitigation through application of crossover technologies from FA cars. It is interesting to note that most of the defining events in figure 7 reflect break-downs in portions of the maneuver and motion loops that tend to be well-matched to the capabilities of emerging automotive automation. This observation suggests safety can be significantly improved by using this technology to perform, assist, or monitor the status of these mid- and lower-level tasks.

Considering the likely possibility that FA cars do not lead directly to FA airplanes and/or that cars with only limited high or conditional automation are available for some time, how might the underlying technologies benefit personal aircraft? Based on the previous section, for either outcome, the pilot will likely retain overall responsibility for mission planning and monitoring, with the assistance of technology that makes relevant information readily available, automates performance calculations, and helps monitor and update the mission status and plans in-flight. It is also
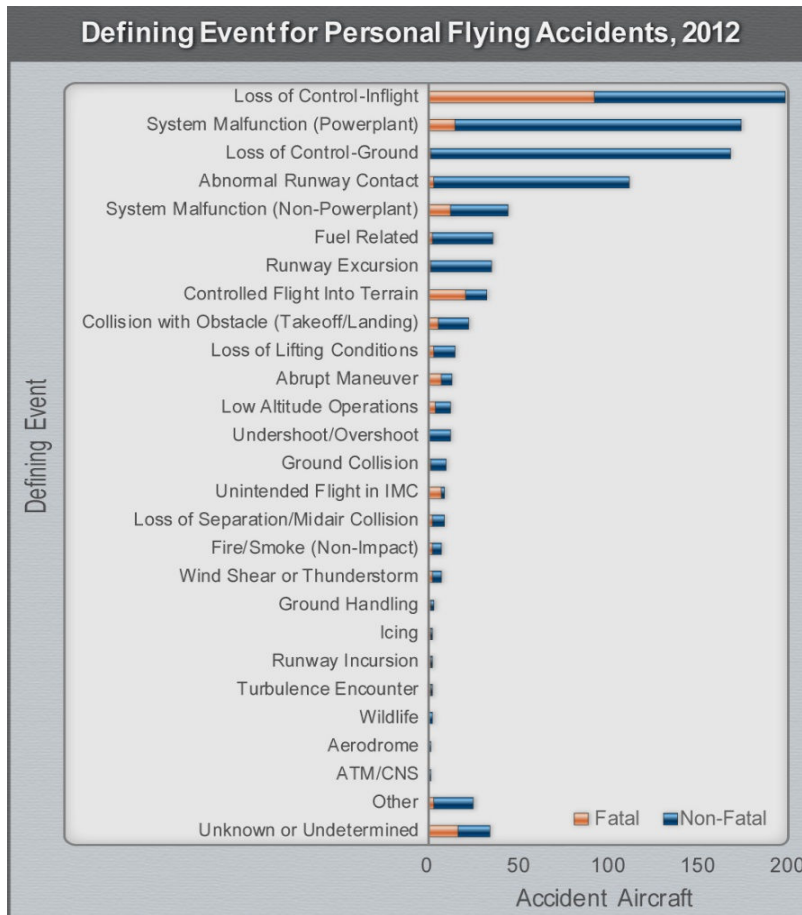
Figure 7 NTSB Summary of Defining Events for Personal Flying Accidents, 2012

likely that the pilot will need to be trained, and engaged, in monitoring the local weather (largely visually) and making, via a suitable interface, tactical changes to the route of flight to avoid local weather hazards that the automation may not be capable of reliably perceiving or evaluating.

Speculating about the impact of automotive automation not progressing to FA is less certain because the specific future, technical challenges and gaps are of course unknown. However, considering the general strength of automation applied to highly structured or deterministic tasks and brittleness in the face of unexpected or novel situations, the most likely scenarios involve cars with highly reliable motion loop automation combined with maneuver loop implementations that, while very capable overall, have limitations mostly likely arising from perception limitations in unusual situations. Human engagement and interaction will be needed in these situations to verify or augment the automation's awareness to ensure that active goals and resulting planned paths and actions are appropriate. The impact on technology transfer to personal aircraft is not particularly different than the FA case since it is already likely that human knowledge and oversight, particularly within tasks of the mission and perception portion of the maneuver loops, will be needed to compliment the automation for either outcome. Despite that probability, the potential of full-time, safety-critical, motion-loop automation, leveraging high-volume and low-cost automotive components, creates the opportunity to fundamentally change the role of the human from a pilot trained to perform all functions within the three loops to that of an operator focused on mission and risk management while relying on the automation to perform high-bandwidth, motion control and system management tasks. As learning manual flight handling and systems management skills throughout the full range of normal and abnormal operations is a significant portion of the current flight training curriculum, requires regular practice to retain proficiency, creates significant workload in critical flight situations, and performance lapses are leading contributors to small aircraft fatalities (e.g., figure 7), crossover automotive automation technology has the potential to dramatically simplify the overall piloting task while improving safety.

Realizing this potential will require significant human-machine integration research including fundamental functional-allocation considerations, new human-machine interaction and interface paradigms[18], and new training methods. In addition, aircraft and pilot certification standards would need to be dramatically revised to reflect the modified roles and responsibilities of both agents.

VI. Conclusion

We have reviewed the emerging capabilities of automated cars and compared these to high-level requirements underlying personal aircraft that are safe and easy to use. The results of this comparison suggest that the technology of cars with full automation is necessary but not sufficient to enable fully automated aircraft. Challenging aviation specific development is likely to be needed in the areas of risk awareness and management along with perception

systems capable of characterizing relevant local weather elements and hazards. While fully-automated cars probably will not directly enable fully-automated aircraft, their constituent technologies may be critical to improving the safety and utility of aircraft. These technologies have the potential to enable aircraft to assist and back-up the pilot in all phases of flight and eliminate or simplify the need to learn, remain proficient at, and perform many lower-level tasks such as manual control and systems operation in both normal and non-normal situations. This fundamental change to the piloting skill set and responsibilities has the potential to dramatically simplify training and in-flight workload while also improving safety. However, realizing these benefits requires significant research in the area of human-automation integration along with significant changes to the regulatory requirements and standards for both aircraft and pilot certification.

References

[1] Underwood, Steven. E., "Automated Vehicles Forecast, Vehicle Symposium Opinion Survey", Automated Vehicles Symposium, San Francisco, CA, 2014. Retrieved from: goo.gl/zzQYBF

[2] Federal Aviation Administration14 CFR Parts 21, 23, 35, 43, 91, 121, and 135 [Docket No.: FAA-2015-1621; Notice No. 16-01] Revision of Airworthiness Standards for Normal, Utility, Acrobatic, and Commuter Category Airplanes. Retrieved from: federalregister.gov/a/2016-05493

[3] Moore, Mark D., et al, "High-Speed Mobility through On-Demand Aviation", 2013 Aviation Technology, Integration, and Operations Conference, Los Angeles, CA, August 12-14, 2013,

[4] Antcliff, Kevin, R., Moore, Mark, D., Goodrich, Kenneth H., "Silicon Valley as an Early Adopter for On-Demand Civil VTOL Operations", AIAA Aviation 2016 Conference, Washington D.C., June 13-17, 2016

[5] SAE International, "Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems, Surface Vehicle Information Report J3016, January, 2014

[6] Kessler, Aaron, M., "Riding Down the Highway, With Tesla's Code at the Wheel", New York Times, October 16, 2015, pp B1

[7] Ito, T., Takata, A., and Oosawa, K., "Time Required for Take-over from Automated to Manual Driving," SAE Technical Paper 2016-01-0158, 2016

[8] Wood, Molly, "Inside the F 015, Mercedes's Self-Driving Car.", New York Times, March 15, 2015, pp B4

[9] Gomes, Lee, "Google Self-Driving Car Will Be Ready Soon for Some, in Decades for Others." IEEE Spectrum, March 18, 2016. Retrieved from: spectrum.ieee.org/cars-that-think/transportation/self-driving/google-selfdriving-car-will-be-ready-soon-for-some-in-decades-for-others-for-others

[10] Urmson, Chris, et al, "Autonomous Driving in Urban Environments: Boss and the Urban Challenge." Journal of Field Robotics, Volume 25, Issue 8, August 2008, pp 425-466

[11] Endsley, Mica, R., "Automation and Situation Awareness." In Automation and Human Performance: Theory and Applications, Lawrence Erlbaum, Mahwah, NJ, 1996

[12] Kenney, John, B., "Dedicated Short-Range Communications (DSRC) Standards in the United States." Proceedings of the IEEE, Volume 99, No. 7, July 2001, pp1162-1182

[13] Wolcott, Ryan W., Eustice, Ryan M., "Visual Localization within LIDAR Maps for Automated Urban Driving." 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, Chicago, IL, September 14-18, 2014

[14] Gohring, Daniel, Wang, Miao, Schnumacher, Michael, Ganjineh, Tinosch, "Radar/Lidar Sensor Fusion for Car-Following on Highways." 5th International Conference on Automation, Robotics and Applications, Wellington, New Zealand, Dec 6-8, 2011

[15] Federal Aviation Administration, "Thunderstorms." Advisory Circular Number 00-24C, February 19, 2013, Retrieved from: www.faa.gov/documentlibrary/media/advisory_circular/ac%2000-24c.pdf

[16] National Transportation Safety Board, "Safety Alert, In-Cockpit NEXRAD Mosaic Imagery." Safety Alert 017, June 2012, revised December 2015. Retrieved from: www.ntsb.gov/safety/safety-alerts/Documents/SA_017.pdf

[17] National Transportation Safety Board, "Summary of US Civil Aviation Accidents for Calendar Year 2012." Retrieved from: www.ntsb.gov/investigations/data/Pages/2012%20Aviation%20Accidents%20Summary.aspx

[18] Schutte, Paul, C., et al, "The Naturalistic Flight Deck System: An Integrated System Concept for Improved Single-Pilot Operations." NASA TM-2007-215090, December 2007