OUR APPROACH TO SAFETY

Pony.ai Safety Report

Dec. 2020

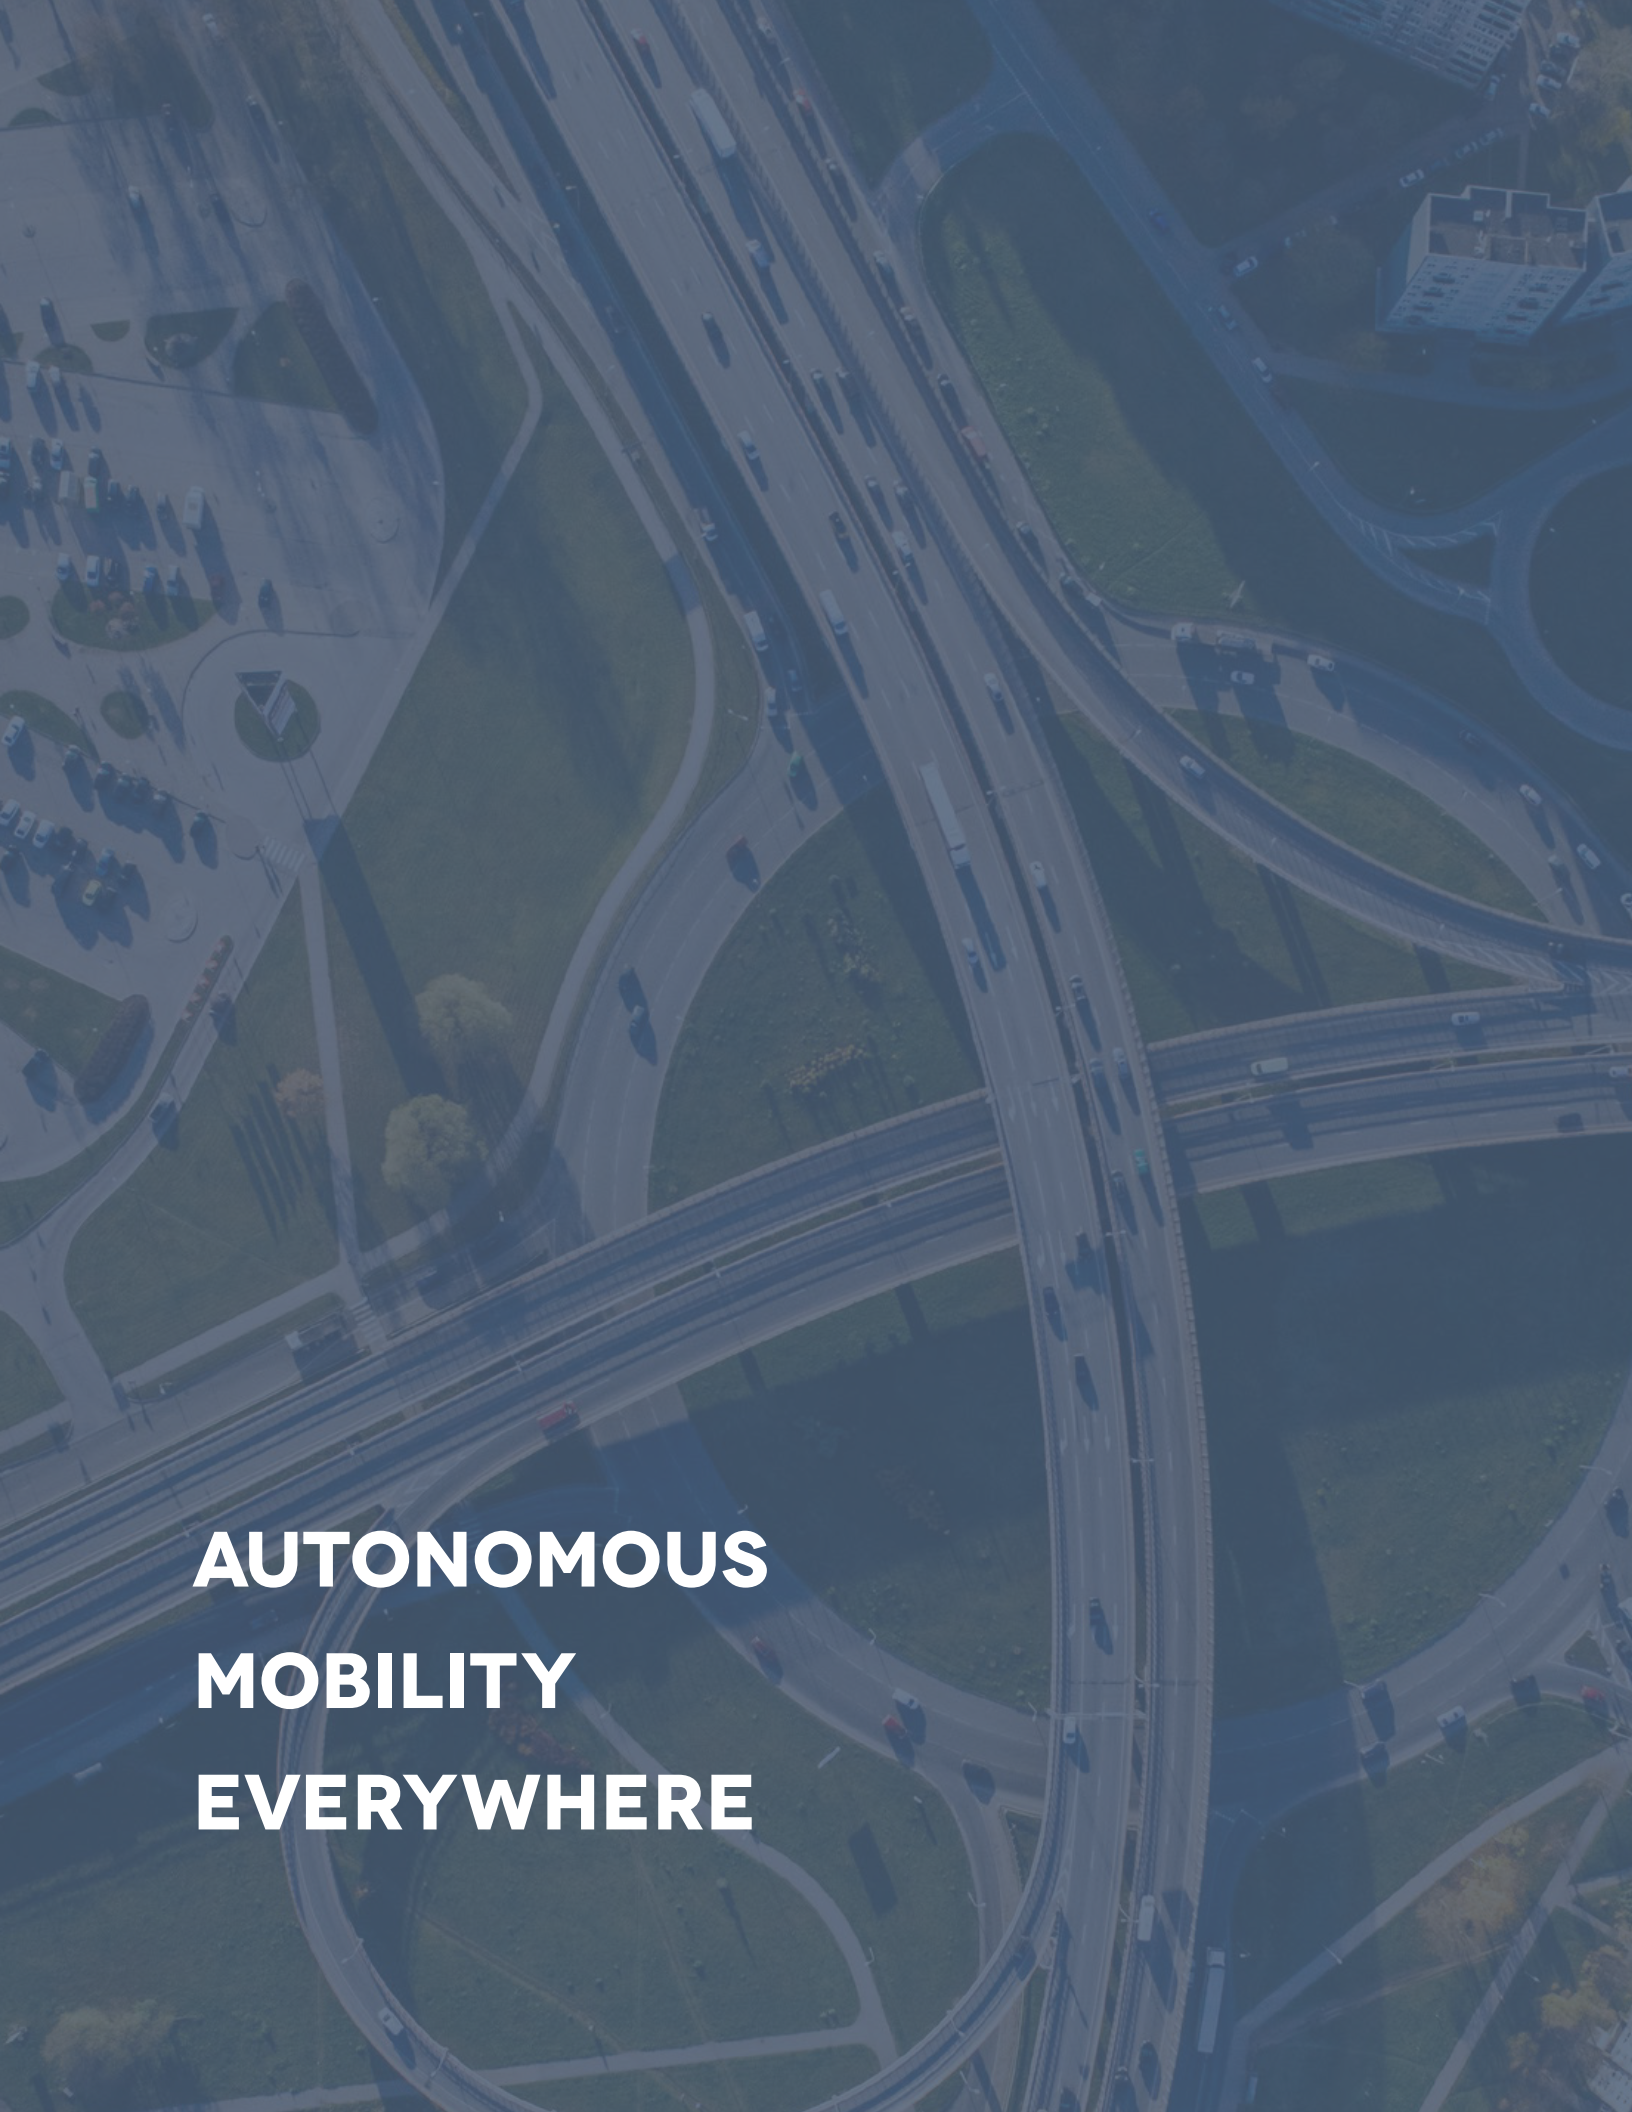

AUTONOMOUS MOBILITY EVERYWHERE

INTRODUCTION:
LETTER FROM THE FOUNDERS

Pony.ai, Inc. ("Pony.ai") is pursuing an ambitious vision for autonomous mobility. We aim to bring safe, sustainable, and accessible mobility to the entire world. We believe that autonomous technology can make our roads exponentially safer for travelers, which is a primary goal for our Company. Safe and reliable systems are essential to realizing the benefits of this technology. Careful and thoughtful development applying detail-focused engineering and demanding testing and evaluation is essential to safe deployment of autonomous vehicles. Thus, a rigorous safety culture and adherence to best practices are the foundation of everything that we do here at Pony.ai. We intend to be one of the first companies to bring autonomous vehicle ("AV") technology to a diverse range of users, populations and markets. We plan to achieve that goal through a corporate culture and safety focus that will establish Pony.ai as a long-term leader in pioneering safe, reliable, and sustainable autonomous transportation.

Our multi-location approach to developing our technology promotes and enhances safe operations over a greater range of driving environments and conditions. Testing and validation in different environments, currently various locations in the United States and China—where traffic conditions, laws, and driving cultures may be vastly different, gives us unique vantage points from which to explore and enhance AV performance. In particular, our approach optimizes identification and resolution of rare "corner case" driving scenarios that present some of the toughest challenges to the safe deployment of AVs at scale. The lessons we learn on the streets of China can inform improvements to our AVs in California, and vice versa. Develop-ing, testing, and refining autonomous systems and vehicles in such vastly different environments is not easy, but that is the point. We have chosen to develop our AVs in the most diverse and challenging environments because it bolsters the safety and reliability of our systems in the long run.

To achieve our goal of safe and reliable autonomous transportation, we have built a world-class team at Pony.ai that is hyper-focused on rigorous engineering and we are partnering with industry-leading vehicle and component manufacturers. We combine groundbreaking innovation with tried-and-true engineering and safety principles, methods, and practices. At Pony.ai, we are convinced that autonomous vehicles will bring about a revolution in mobility that will make transportation systems safer, more efficient, and more accessible for all. We believe the best course to that future—indeed, the only justifiable course—is through a cautious, rigorous and "no shortcuts" approach and philosophy. That core, pervasive belief guides all of our autonomous mobility research, development, testing, and deployment.

James Peng
Co-founder & CEO

Tiancheng Lou
Co-founder & CTO

CONTENTS

01 THE PONY.AI AUTONOMOUS DRIVING SYSTEM

- Sensors
- Software
- Hardware & Vehicle Integration
- Onboard Monitoring System

Sensors

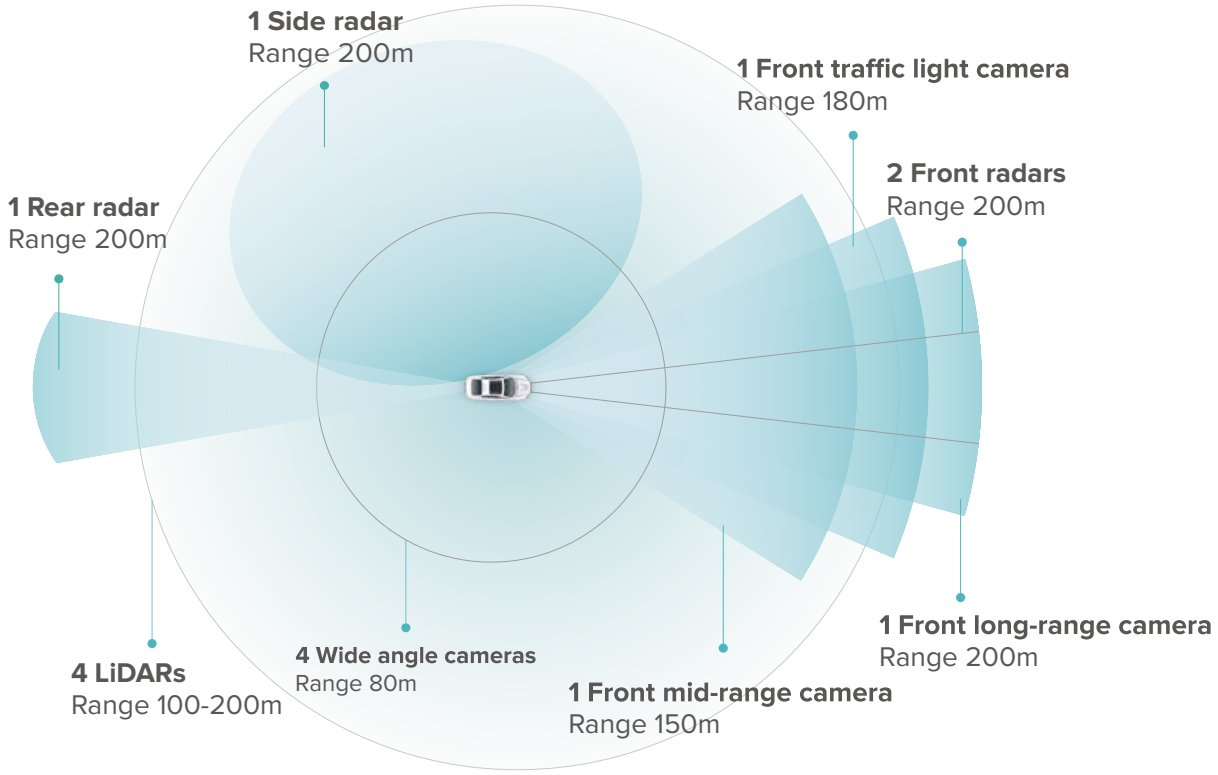

In order to accurately and precisely perceive and understand the environment surrounding our AVs, the Pony.ai autonomous driving system ("ADS") employs a multi-sensor approach. Two aspects of our sensor suite are essential to our safety approach. The first aspect is diversity: by using cameras, radars, and LiDARs together in our Perception sensor suite, we can leverage the important strengths of each type of sensor to cover, collectively, the range of necessary perception required for safe autonomous driving. Cameras provide excellent data for object classification and color information, e.g., what is that object ahead of us? At the same time, radars and LiDARs can be more helpful for location and motion analysis, e.g., how far away is the object and where is it moving? By combining these relative strengths, we can design a suite of complementary capabilities that gives our vehicle a complete picture of the world around it.

The second essential aspect is robustness. To be fully effective, a sensor suite must be reliable. Our sensors and components are engineered for the rigors of the road and the real world. Moreover, we have designed and incorporated to our ADS a robust active monitoring system that keeps constant watch on the status of each component. If a sensor is not performing, our system identifies the problem before it creates a safety risk.

In addition to Perception sensors, we also use two other types of sensors in our system, a high-accuracy Global Navigation Satellite System (GNSS) and an Inertial Measurement Unit (IMU). These sensors work together with our high-definition maps and localization module to ensure accurate positioning of our autonomous vehicles.

Software

Perception

The Perception module enables us to see and understand the world around our autonomous vehicles. We do this by fusing and processing the data collected by our sensors and high-definition map. The Perception module is responsible for object segmentation, detection, classification, and scene (i.e. event) understanding. In short, Perception identifies and communicates what is happening around the vehicle. Accordingly, the fidelity and accuracy of our Perception module is critical to the safe operation of our autonomous vehicles. To ensure performance, we leverage a hybrid solution that combines the best of deep learning and the heuristics approach to process, refine, and use the data collected by our sensors and high-definition map. This hybrid approach allows high-confidence and accurate detection, classifications and tracking in dense and complex environments. In addition, we use sensor fusion to optimize data usage, which in turn increases the reliability of Perception overall.

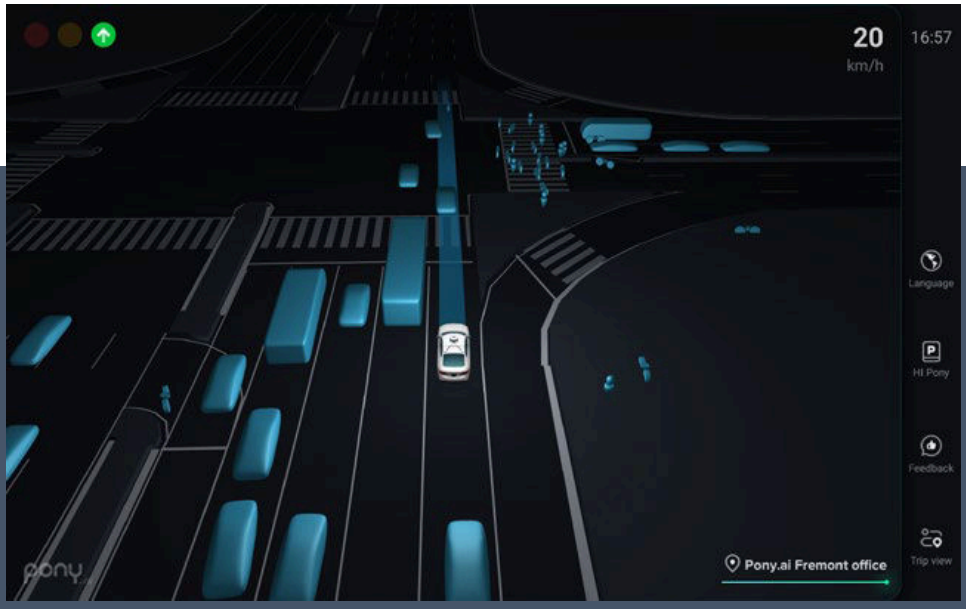

Prediction

The Prediction module works to project how other vehicles, pedestrians, and objects (together, "road agents") may move or behave based on several inputs, including Perception output, raw sensor data, and data regarding previous decisions made by the road agent (i.e. sequence of events). The output of the Prediction module for a certain road agent is a series of predicted trajectories, each with an assigned probability of occurrence.

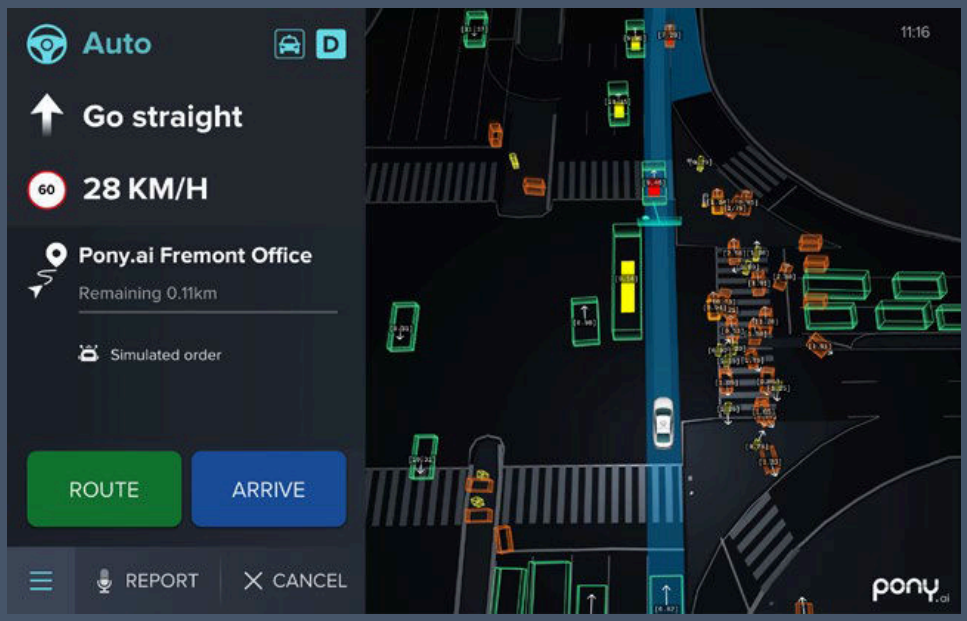

Similar to the methodology used in Perception, Pony.ai's Prediction module uses a mix of deep learning and heuristics to enable rapid learning and adaptation. This is particularly important for autonomous driving systems because of the sheer amount of road data generated on any given day across the entire global fleet. We are able to fully leverage the breadth and variety of these road scenarios to refine and improve Prediction capabilities at a speed that far outpaces that of a human driver.

Localization and Mapping

Accurate localization and mapping are in many ways the foundation of a reliable self-driving system. However, ordinary mapping technology is not sufficient to meet the high demands of self-driving technology. What is required is centimeter-level position accuracy of the autonomous vehicle and surrounding road agents, and lane-level understanding of the surrounding world to handle daily and corner-case driving conditions and events.

At Pony.ai, our engineering expertise and the multi-sensor approach, which captures a rich dataset, allows us to create a state-of-the-art localization solution and high-definition maps. Our localization module currently supports precise and accurate x, y, z, roll, pitch, yaw, velocity, and the acceleration and angular velocity . In addition, the multi-sensor approach ensures safety because if the localization information from one system becomes unavailable, our vehicle can still use localization information generated by other sources.

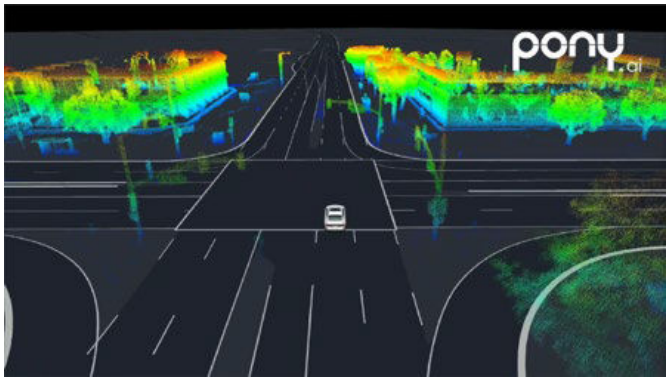

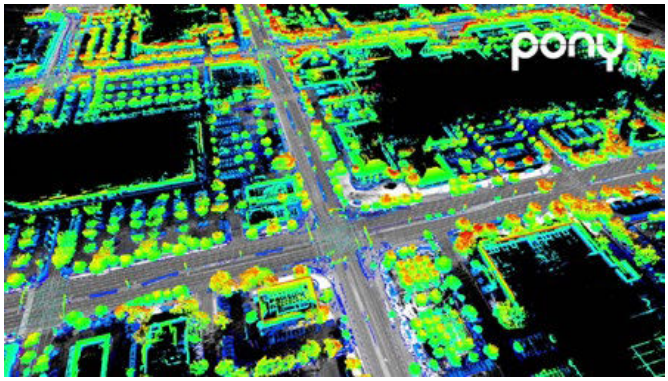

Planning and Control

The Planning and Control module is responsible for planning and executing safe, correct, and effective road maneuvers based on input from the Perception, Prediction, and Localization modules. This module also combines machine learning models to ensure safe general on-road operation and appropriate responses to outlier conditions, actions, or events. Our control module performs precise autonomous vehicle control maneuvers at a centimeter-level accuracy.

Hardware & Vehicle Integration

Computation System

One of the unique challenges of autonomous driving technology is the sheer volume of data that needs to be taken in and processed. It is simultaneously an Artificial Intelligence (AI) and Big Data challenge. The onboard computation system is responsible for processing the data gathered from the sensors and running our proprietary algorithms in real-time to enable our vehicles to drive autonomously. In designing and configuring our computing system, we optimize for performance, reliability, and resource efficiency.

The central computation system utilizes the best available industrial heterogeneous computing architecture, including central processing unit (CPU), graphics processing unit (GPU), field programmable gate array (FPGA), microcontroller unit (MCU) and high throughput storage. It collects sensor data at a high bandwidth, provides compute power and acceleration for software algorithms, and eventually sends control commands to the vehicle interfaces.

Each piece of the computation system represents the best available automotive grade components and is validated by well-defined compliance tests. The Pony.ai hardware architecture emphasizes redundancy to ensure safety and reliability. Different processors cross-check and function as back-up systems for one another. For example, certain algorithms running on GPU can fall back to the CPU if an error occurs. In addition, our computation system is supported by a dual input redundant power system that does not affect the base vehicle's original OEM-designed power system. If the main power system fails, the back-up power system will seamlessly engage and ensure continuous power to the computation system and, therefore, the continued operation of the autonomous system as a whole.

Vehicle Integration

The Pony.ai system's final layer is its integration into base vehicles. We have built a reliable interface between the autonomous driving software stack and the vehicle system to ensure that the vehicle platform can accurately execute control commands. The solution is built on an automotive grade hardware and software toolchain and includes mature prognostic / diagnostic management to monitor both the autonomous system and the base vehicle system.

Onboard Monitoring System

Onboard System

The Pony.ai Onboard System schedules, runs, and monitors all of the separate modules that comprise the total autonomous driving system. The Onboard System also implements a unified application programming interface (API) for module communications, which maintains a stable data flow from the upstream sensors all the way to the downstream planning and control modules.

The Onboard System is a highly optimized system, which in turn plays a key role in ensuring safety and performance. For instance, in the data flow, camera data is stored only in the GPU. Elimination of data transfer between the main memory and GPU memory reduces overall latency, which enables our system to respond faster to obstacles on the road. In addition, the Onboard System closely monitors any potential malfunctions of the autonomous system and its components. For instance, vehicle operators are alerted with visual and audio signals when any sensor, controller, or module is in an unhealthy state or seeing abnormal latencies. In addition, autonomous mode cannot be engaged if the vehicle is in manual mode when such issues are detected by the Onboard System.

Vehicle Operator User Interface (UI)

Our Vehicle Operator UI provides a 3D graphical user interface between the vehicle operator and the onboard system. It displays the obstacles around our vehicle, and shows the state (autonomous/manual, speed, heading, etc.), routes, and planned trajectory of our vehicle. Operators can also use touch-screen mechanisms to provide commands to the Vehicle Operator UI.

Vehicle Operator UI features a visual alert system. If the onboard system or any of its components is in an abnormal state, a red pop-up alert notifies the operator to disengage the autonomous system. In addition to visual features, the Vehicle Operator UI also provides an audio interface, which allows each component to tell the operators safety-critical decisions, and remind safety operators to disengage in case of potentially dangerous road conditions, e.g., the vehicle behind us is fast approaching and risks a collision with our vehicle. In addition to the rigorous training our vehicle operators undergo, the Vehicle Operator UI functions as an additional layer to help preempt and mitigate on-road safety risks.

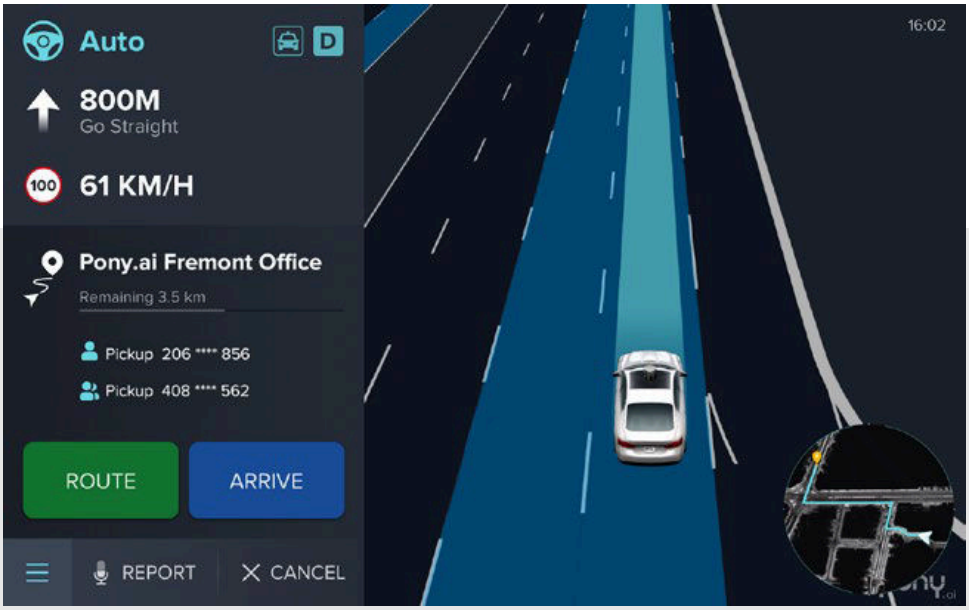

PonyDash for vehicle operater

02 DEVELOPING, TESTING, AND VALIDATING THE PONY.AI SYSTEM

- Multi-Market Approach
- Rigorous Engineering
- Fallback Solutions and Additional Safety Layers

SAFETY FIRST

"Safety first"—it is important to say, but actions are more important than words. In every action we take and every line of code we write, Pony.ai puts safety at the heart of what we do.

This starts with an appreciation of the challenges in developing autonomous vehicles and how we approach testing these systems. Most of what a human driver encounters in a vehicle on their average commute is relatively simple, but real safety is a "long tail" problem—unusual circumstances and unexpected challenges make up a relatively small portion of what a vehicle encounters, but they sometimes represent the most critical safety risks, and thus the most challenging problems for an autonomous driving system. Successfully navigating rare but dangerous conditions or circumstances is a must for any safe autonomous driving system.

A second challenge is validation—how to confirm that the system can safely operate in a challenging road environment. Pony.ai's response to validation challenges combines the innovation of new technology and software with well-established safety and engineering principles. We take pride in the transformative potential of our technology, and we marry it to established best practices used by vehicle safety engineers to ensure the safety of vehicle occupants and other road users.

Multi-Market / Location Approach

Many companies, from century-old car manufacturers to newer technology companies like Pony.ai, are working to develop fully autonomous vehicles, but few are designing, developing and testing vehicles globally. Pony.ai is currently testing in multiple demanding environments: the streets and highways of different locations in California, and multiple locations in China, including Guangzhou, a center of Chinese economic growth and innovation, Beijing, China's political and cultural capital, and Shanghai, its automotive capital and major economic hub.

While there are important business advantages to this multi-market approach, perhaps more important are the safety advantages. One of the biggest hurdles to deployment of automated vehicles are "corner cases" – relatively rare occurrences (often the result of a combination of multiple factors at once) that present potential safety risks that challenge the driving system's decision making. These cases may challenge the decision-making capacity of human drivers as well, but humans often intuitively work their way through unanticipated difficult circumstances, even those they have not previously experienced. Safe autonomous vehicles require similar problem-solving skills. Therefore, development and testing in challenging environments where corner cases are more likely to occur is essential.

California and China are extremely challenging, but decidedly different, environments for motorists, human or automated. They have different car cultures, different expectations from drivers, passengers and vulnerable road users, different road designs and traffic laws, and different written and unwritten rules of the road. Training our autonomous vehicles under these conditions allows us to test and refine our systems under dramatically greater exposure to corner case challenges.

In addition, Pony.ai leverages lessons learned in one testing environment to benefit the system's performance everywhere. When we encounter a challenging scenario in one testing location, we design a solution that not only solves the challenge for that specific location but also works in similar scenarios in other testing locations. If a solution developed for one location is ineffective in another, we continue to refine the solution until it is more broadly applicable. This helps us create an AV system that can handle the most difficult corner-case scenarios across a diverse range of road and traffic conditions in multiple locations as safely as possible. Our ultimate goal is to create an AV system that can operate safely in many different markets and road scenarios. Our multi-market approach is key for us to reach that goal.

Rigorous Engineering

The Pony.ai safety concept is first and foremost built on a foundation of safe and responsible development practices. Our team collectively has expertise in truly world-class engineering and product development. We apply a systems engineering approach to ensure every component is considered separately and as part of a whole. In our test and validation methods, we also apply existing industry safety standards such as ISO 26262, but adapt its framework to better test and validate autonomous driving systems. In addition, we employ various tried and proven safety analysis tools such as FMEA (Failure Mode and Effects Analysis) and FTA (Fault Tree Analysis) to validate the system and inform development areas.

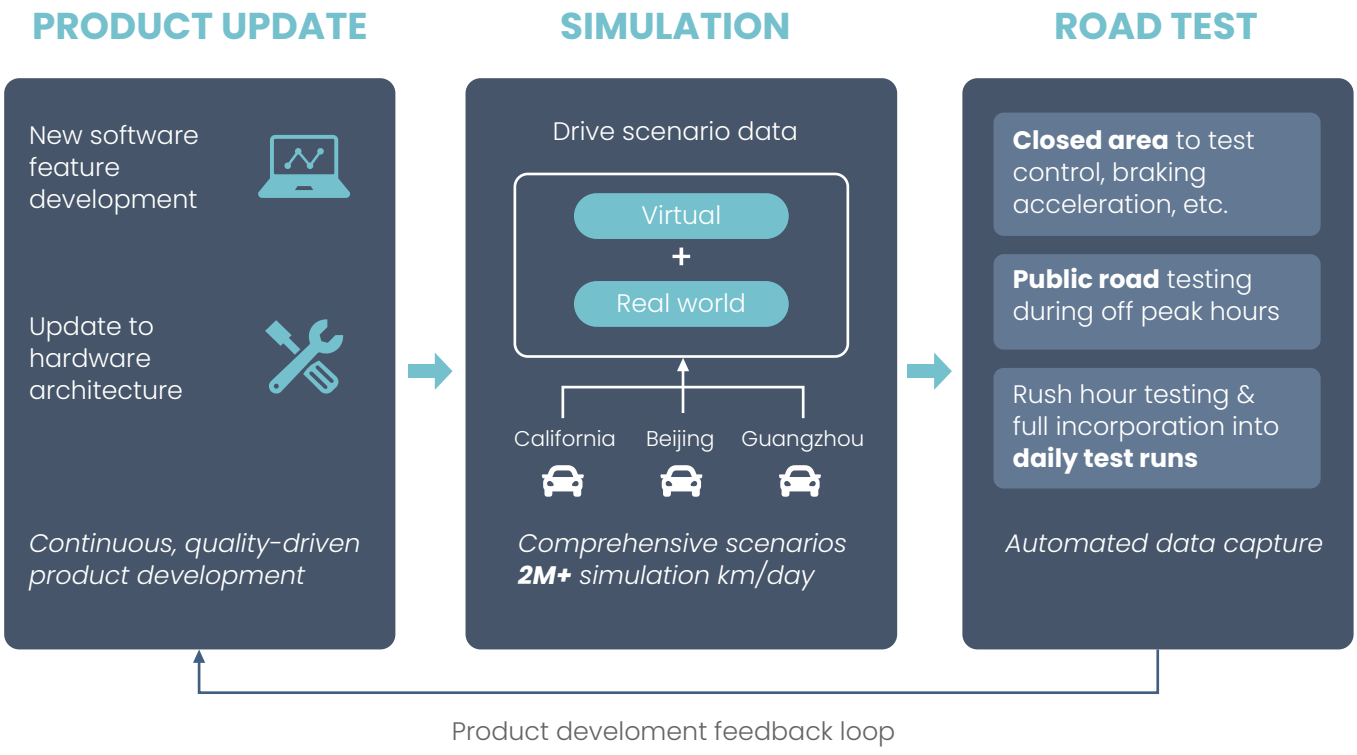

Software Development

Pony.ai's overall development process embeds safety throughout. We begin with a rigorous and carefully designed development workflow. Software development at this level of complexity demands a rigorous development process with multiple safety checks at each step of the way. Each time we make a change to our software, it is in response to an issue we've identified for improvement during our testing or to solve a challenge in expanding our vehicle's safe operating domain. We use data from our on-road testing, including feedback from onboard safety engineers, to identify potential risks to the vehicle's safe and dependable operation.

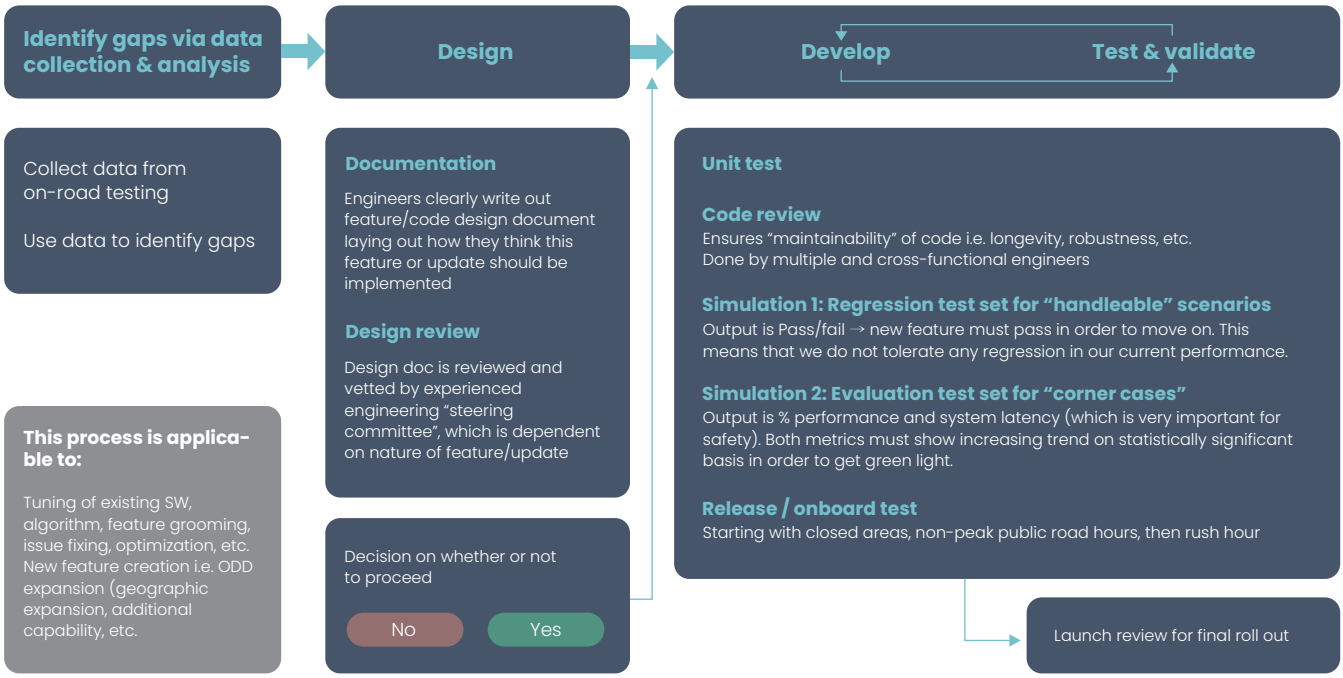

A key initial step in addressing such identified issues is documenting how we intend to fill it. Each time one of our engineers starts working on a new feature or parameter adjustment, they begin by documenting how the new element or feature is intended to improve the system. This design document is an important reference both in establishing the problem we are working to solve and in outlining how the new code is intended to solve that problem—it initiates the process with a clear purpose, and subsequently serves as a reminder of that purpose and goal through every following step. That design document is then reviewed by a team of experienced engineers on one of several feature-specific steering committees to ensure that it is both technically sound and consistent with our broader approach. The relevant steering committee must approve a design document before detailed development work begins.

Once code is written, it is subjected to a multi-element validation process before it is authorized for installation on a test vehicle. First, experienced engineers cross-check to ensure that it is maintainable, i.e., that it is not a quick fix that could cause more challenges later but rather can remain a long-standing feature in our system.

We also perform two kinds of simulation tests that leverage an ever-growing database of test scenarios. The first is a regression test to ensure any code to be deployed in our system will not diminish the system's ability to handle any foreseeable scenarios or degrade the high level performance. The second test uses a series of corner-case scenarios to challenge the ability of the system with the proposed revisions to handle those rare but important occurrences. We are looking for improvements in two areas: the system's ability to safely navigate corner cases, and to process those cases more quickly. Excessive latency can be especially prob-lematic from a safety standpoint. Therefore, we focus on improving not only the system's ability to get it right, but to do so at speed. In order to move forward, any software change must demonstrate significant improvement on overall performance and on latency.

Once the new code has passed both those tests, it is ready to deploy in the physical world. This process happens slowly, carefully, and through an established ramp-up from testing on closed courses, to lightly traveled streets, and gradually moving into more challenging traffic scenarios. We review data from these test scenarios rigorously before the code is approved for final roll-out.

Data Recording and Retention

Data is the fuel for autonomous driving technology and algorithm development. All development and feature improvement begin with real-world data collected through our on-road testing operations.

All sensor data and internal system data (i.e., how every driving decision was made) is collected and processed through the Pony.ai data pipeline, which helps to identify and categorize the types of data we collect and then informs how the data is stored and used. This type of systematic data combing is incredibly important because it identifies valuable road scenarios and provides our system with robust datasets to train our system. As described above, this recorded data is also used for several types of regression tests that we run in simulation, which allows better validation of our system. In addition, our stored data allows us to perform event reconstruction analysis and determine the root-cause of any issues our vehicles encounter on the road.

Hardware Development

Similar to our approach to software development, our hardware development processes leverage existing best practices while also balancing rapid iteration and unique needs of autonomous driving technology development. Our hardware development concept is based on the SAE J1938 industry standard and the V-model to ensure the highest quality output. We have developed an internal design flow that is both agile and meets rigorous performance needs: requirement collection, design, review, validation and release.

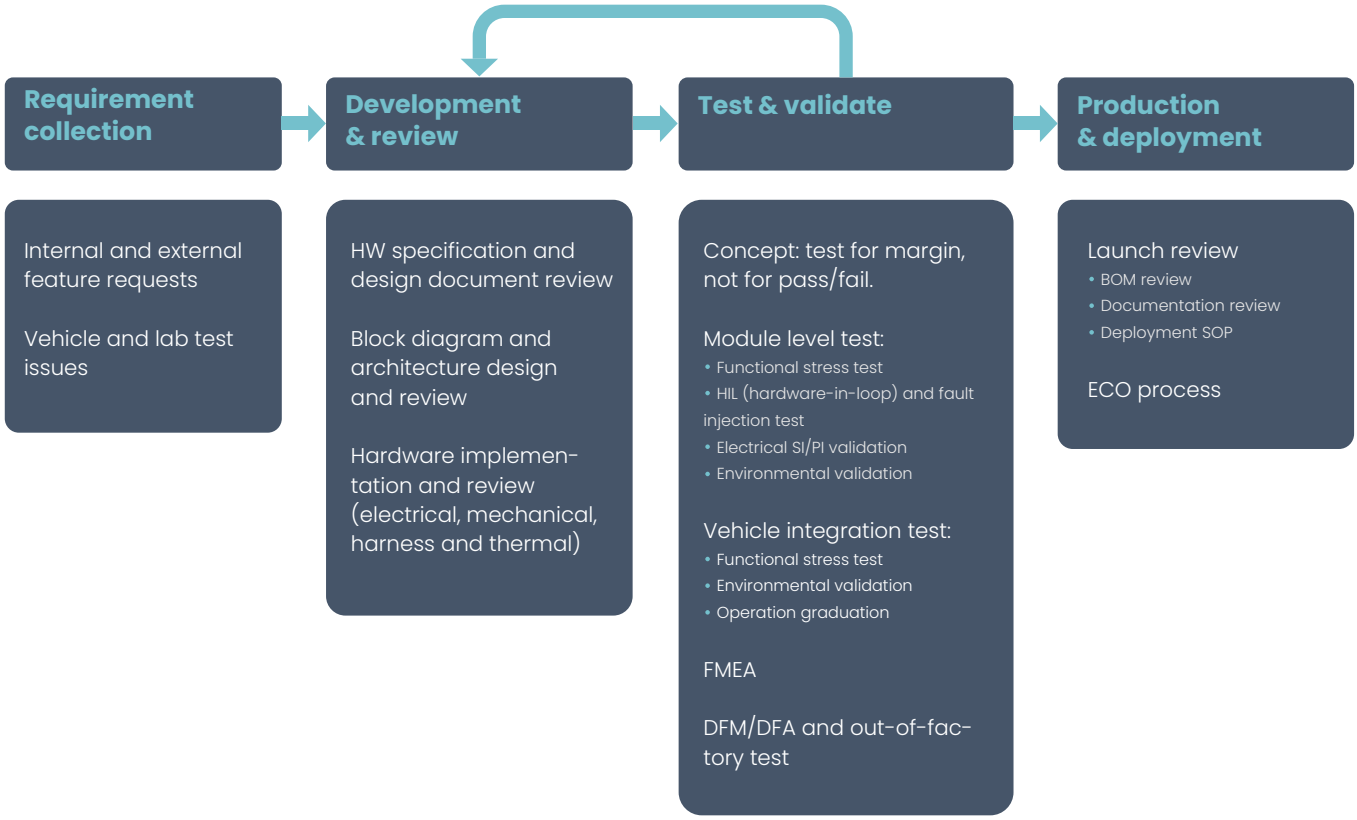

Operational Design Domain (ODD)

Another important concept in safe autonomous vehicle design is the Operational Design Domain (ODD), which encompasses physical surroundings, road types, traffic conditions, weather, time of day—a host of parameters that describes the conditions under which our vehicles can safely operate autonomously. To responsibly test and ultimately deploy fleets of autonomous vehicles at scale, understanding the capabilities and limits of the system is vital.

The key is in how we approach ODD selection. Pony.ai will only deploy its autonomous vehicles in environments where the vehicles' driving behavior meets Pony.ai's safety and performance standards. Before we put our autonomous system in any kind of driving environment, we have performed testing and validation to ensure that the hardware and software are up to the task, and we maintain constant vigilance to ensure that every decision the vehicle makes holds up in the real world.

Pony.ai's current ODD is designed to ensure our vehicles are prepared to navigate a wide-range of driving challenges, including:

- Nearly every type of roadway condition, from major multi-lane arterials to congested city streets and tight-quarters parking lots
- A wide range of traffic conditions, from wide-open streets to heavy congestion
- A wide range of weather conditions including but not limited to wind, rain, 24-hour light conditions, and extreme temperatures
- With regard to the full range of possible road users – including every type of vehicle, pedestrians, cyclists, as well as a broad array of potential objects and obstacles – from a wild turkey wandering across a country road to a runaway shopping cart rolling across an urban parking lot

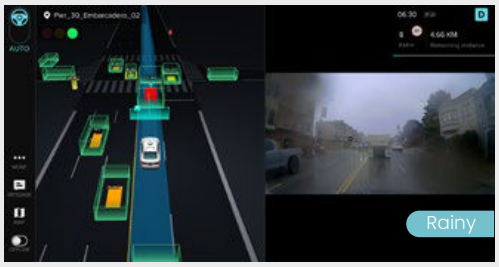

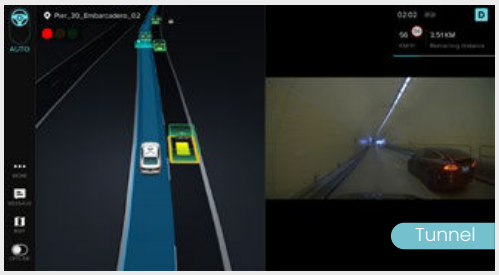

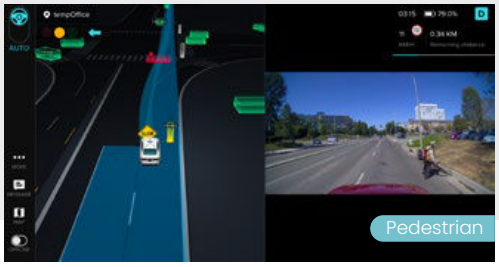

Our vehicles operate in specific geographic locations in the communities where we test. We carefully consider any expansion of those geographic boundaries, ensuring that our system has demonstrated safe operation in similar conditions and is prepared to comply with all local traffic laws and regulations. When we go forward with testing in new locations, we rigorously monitor performance in the new environment and adjust the ODD as necessary to ensure a wide margin for safety.

Fallback Solutions and Additional Safety Layers

Independent Monitoring System and Minimal Risk Condition

An important feature of the Pony.ai system is a proprietary complete and independent monitoring system that checks the real-time health and status of each module. When the vehicle is running, the monitoring system is active across each module of the hardware and software systems and can alert the system or onboard personnel of any abnormalities. These personnel are trained to understand when to take the vehicle out of self-driving mode. In addition, when the monitoring system detects an autonomous vehicle system failure that will render the system unable to safely control the vehicle, the system will automatically enter minimal risk condition (MRC) mode by safely pulling the vehicle over to the side of the road and coming to a complete stop, including automated turn signals and hazard lights if necessary.

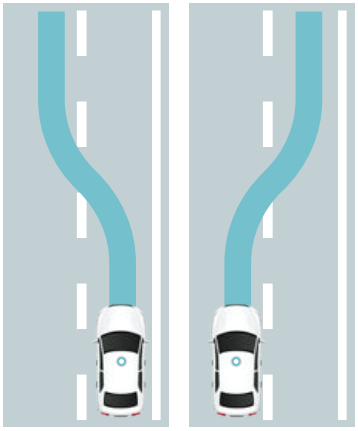

Behavior: safe pull-over

Safe state: vehicle drive to the closest safe operational area

Vehicle Operators and Onboard Safety Engineers

Vehicle Operators and Onboard Safety Engineers are tasked with monitoring the vehicle's surrounding environment and onboard system to ensure that the system is safely performing as intended. The role of these onboard personnel is particularly important when we roll out new features or test relatively bigger updates to the system. Our goal is to develop a vehicle that operates more safely and reliably than the best human driver, but developing a vehicle that humans can rely on means human beings must be at the heart of the process.

When we perform on-road testing, we want our Operators to look for challenging scenarios because that is the most valuable feedback as we work to develop a safe and reliable system. They are evaluating the vehicle's performance with an engineering eye—they understand how the hardware and software are designed to work and can identify issues that only experienced engineers can see. They are also evaluating the system from an ordinary driver's perspective, looking out for actions that might surprise human drivers or make the occupant's experience uncomfortable.

To ensure that Vehicle Operators are adequately prepared for their roles, we have designed a multi-week training program that every Operator is required to undergo before independent testing is allowed.

Crashworthiness & Post-Crash ADS Behavior

Pony.ai is fortunate to count some of the world's most experienced and established automotive original equipment manufacturers (OEMs) as strategic partners. We rely on these base vehicles' safety systems as the foundation into which our system is integrated. All of these base vehicles meet applicable Federal Motor Vehicle Safety Standards (FMVSS) and are certified by our partners and suppliers. In addition to these assurances, we also test and validate each base vehicle's automated driver assistance (i.e. advanced driver-assistance systems (ADAS)) functions through our vehicle engineering development and system integration process.

One structure we currently build in-house is the sensor housing structure that holds our perception sensors in place and sit on top of the vehicle. We design these housing structures based on best practices in crashworthiness and ruggedness. Structural integrity under extreme load is validated according to existing automotive industry standards.

In addition, our operators are trained in post-crash procedures in the event of any public road collision, which include protocols to interact with the other vehicle and its occupants (if any) and work with law enforcement or other responders.

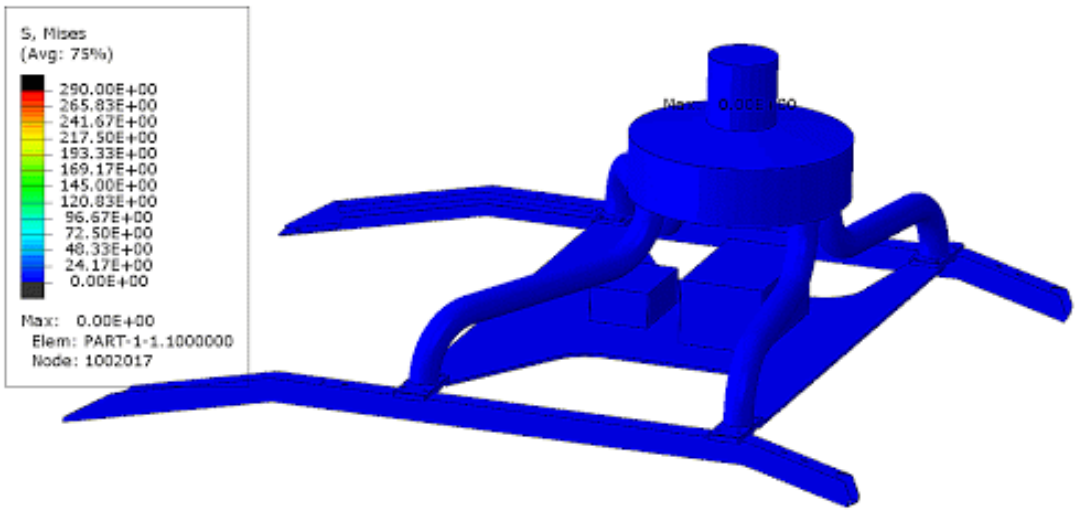

Vehicle Cybersecurity

Much of autonomous driving technology development is about defining and creating new playbooks where none existed previously. Vehicle cybersecurity is no exception and is a central and critical safety layer, from how data is protected in the data center to the real-world data that our vehicles are processing in real-time.

While we continue to evolve and iterate our concept of vehicle cybersecurity to make autonomous vehicles safer for everyone, we have always leveraged existing standards ISO 21434 and SAE J3061 to serve as a foundation on which to build a more tailored and protective cybersecurity concept. Additionally, we have developed robust procedures and proprietary systems to protect against cybersecurity threats.

As our autonomous fleet scales and our technology develops, we continue to refine our cybersecurity features to make our protections even more robust. We recognize that cybersecurity is an essential safety layer that will be critical to scaling and expanding our operations.

03 COMMUNITY ENGAGEMENT

We take great pride in being one of the first companies to allow members of the public in multiple markets to experience fully autonomous vehicles for the first time. However, we also recognize that being first comes with a great responsibility to ensure that the public understands and accepts this technology. The road to widespread adoption of autonomous vehicles is a long one, and consumer acceptance is a critical step along the way.

In all of the locations where we have run pilots, we have engaged deeply with the local authorities and communities to make them aware of our presence and the services we offer. We often organize "tech talk" events that are open to the public and our pilot participants to introduce our technology and what to expect when riding in one of our autonomous vehicles. For example, in the months leading up to the launch of our public-facing pilot in Irvine, CA, we conducted multiple such sessions with the local community. In addition, after our public-facing pilots have launched, we stay close to our user base to garner feedback that becomes valuable input to our future product development and refinement.

CONCLUSION

Pony.ai specializes in developing the systems that enable a vehicle to operate autonomously—i.e., to drive safely without any human intervention. We are working to expand the universe of safe operation, to deploy the system in more places, under more conditions, so that we can grow the safety and mobility benefits that autonomous driving technology promises.

We are deeply committed to a vision of wide-scale deployment of autonomous vehicles, and passionate about their promise and benefits to society. At the same time, we are humbled by the magnitude of this goal, both in its technical complexity as well as its societal signifi-cance. Since day one, the first tenet of our company culture has always been "Safety First," and we embody this spirit into every-thing that we do.

If you would like to learn more about Pony.ai, please visit us at *www.pony.ai*.

*This report illustrates the safety of Pony.ai's self-driving system for road testing in California.