

Artificial Passenger

¹Shivam M. Tayade, ²Krunal M. Ingle, ³Sumit A. Tirpude and ⁴Kaustubh U.Pathak,

^{1,2,3,4}Department of ENTC Engineering, Mauli College of Engineering & Technology, Shegaon, Maharashtra, India

Abstract— An artificial passenger (AP) is a device that would be used in a motor vehicle to make sure that the driver stays awake. IBM has developed a prototype that holds a conversation with a driver, telling jokes and asking questions intended to determine whether the driver can respond alertly enough. Assuming the IBM approach, an artificial passenger would use a microphone for the driver and a speech generator and the vehicle's audio speakers to converse with the driver. The conversation would be based on a personalized profile of the driver. A camera could be used to evaluate the driver's "facial state" and a voice analyzer to evaluate whether the driver was becoming drowsy. If a driver seemed to display too much fatigue, the artificial passenger might be programmed to open all the windows, sound a buzzer, increase background music volume, or even spray the driver with ice water.

Keywords: *Telematic, microphone, voice analyzer.*

I. INTRODUCTION

The AP is an artificial intelligence-based companion that will be resident in software and chips embedded in the automobile dashboard. The heart of the system is a conversation planner that holds a profile of you, including details of your interests and profession. When activated, the AP uses the profile to cook up provocative questions such "Who was the first person you dated?" via a speech generator and in-car speakers. A microphone picks up your answer and breaks it down into separate words with speech-recognition software. A camera built into the dashboard also tracks your lip movements to improve the accuracy of the speech recognition. A voice analyzer then looks for signs of tiredness by checking to see if the answer matches your profile. Slow responses and a lack of intonation are signs of fatigue. If you reply quickly and clearly, the system judges you to be alert and tells the conversation planner to continue the line of questioning. If your response is slow or doesn't make sense, the voice analyzer assumes you are dropping off and acts to get your attention. The system, according to its inventors, does not go through a suite of rote questions demanding rote answers. Rather, it knows your tastes and will even, if you wish, make certain you never miss Paul Harvey again. This is from the patent application: "An even further object of the present invention is to provide a natural dialog car system that understands content of tapes, books, and radio programs and extracts and reproduces appropriate phrases from those materials while it is talking with a driver. For example, a system can find out if someone is singing on a channel of a radio station. The system will state, "And now you will hear a wonderful song!" or detect that there is news and state, "Do you know what happened now—hear the following and play some news." The system also includes a recognition system to detect who is speaking over the radio and alert the driver if the person speaking is one the driver wishes to hear." Just because you can express the rules of grammar in software doesn't mean a driver is going to use them. The AP is ready for that possibility: It provides for a natural dialog car system directed to human factor engineering for example, people using different strategies to talk (for instance, short vs. elaborate responses). In this manner, the

individual is guided to talk in a certain way so as to make the system work—e.g., "Sorry, I didn't get it. Could you say it briefly?" Here, the system defines a narrow topic of the user reply (answer or question) via an association of classes of relevant words via decision trees. The system builds a reply sentence asking what are most probable word sequences that could follow the user's reply.

II. WHY ARTIFICIAL PASSENGER

IBM received a patent in May for a sleep prevention system for use in automobiles that is, according to the patent application, "capable of keeping a driver awake while driving during a long trip or one that extends into the late evening. The system carries on a conversation with the driver on various topics utilizing a natural dialog car system." Additionally, the application said, "The natural dialog car system analyzes a driver's answer and the contents of the answer together with his voice patterns to determine if he is alert while driving. The system warns the driver or changes the topic of conversation if the system determines that the driver is about to fall asleep. The system may also detect whether a driver is affected by alcohol or drugs." If the system thinks your attention is flagging, it might try to perk you up with a joke, though most of us probably think an IBM engineer's idea of a real thigh slapper is actually a signal to change the channel: "The stock market just fell 500 points! Oh, I am sorry—I was joking." Alternatively, the system might abruptly change radio stations for you, sound a buzzer, or summarily roll down the window. If those don't do the trick, the Artificial Passenger (AP) is ready with a more drastic measure a spray of icy water in your face.

III. FUNCTIONS OF ARTIFICIAL PASSENGER

A. Voice Control Interface

One of the ways to address driver safety concerns is to develop an efficient system that relies on voice instead of hands to control Telematics devices. It has been shown in various experiments that well designed voice control interfaces can reduce a driver's distraction compared with manual control situations. One of the ways to reduce a driver's cognitive workload is to allow the driver to speak naturally when interacting with a car system (e.g. when playing voice games, issuing commands via voice). It is difficult for a driver to remember a complex command menu (e.g. recalling specific syntax, such as "What is the distance to JFK?" or "Or how far is JFK?" or "How long to drive to JFK?" etc.). This fact led to the development of Conversational Interactivity for Telematics (CIT) speech systems at IBM Research. CIT speech systems can significantly improve a driver-vehicle relationship and contribute to driving safety. But the development of Full fledged Natural Language Understanding (NLU) for CIT is a difficult problem that typically requires significant computer resources that are usually not available in local computer processors that car manufacturers provide for their cars. To address this, NLU components should be located on a server that is accessed by cars remotely or NLU should be downsized to run on local computer devices (that are typically based on

embedded chips). Some car manufacturers see advantages in using upgraded NLU and speech processing on the client in the car, since remote connections to servers are not available everywhere, can have delays, and are not robust. Our department is developing a “quasi-NLU” component - a “reduced” variant of NLU that can be run in CPU systems with relatively limited resources. It extends concepts described in the paper. In our approach, possible variants for speaking commands are kept in special grammar files (one file for each topic or application). When the system gets a voice response, it searches through files (starting with the most relevant topic). If it finds an appropriate command in some file, it executes the command. Otherwise the system executes other options that are defined by a Dialog Manager (DM). The DM component is a rule based sub-system that can interact with the car and external systems (such as weather forecast services, e-mail systems, telephone directories, etc.) and a driver to reduce task complexity for the NLU system.

The following are examples of conversations between a driver and DM that illustrate some of tasks that an advanced DM should be able to perform:

1. Ask questions (via a text to speech module) to resolve ambiguities:

- (Driver) Please, plot a course to Yorktown
- (DM) Within Massachusetts?
- (Driver) No, in New York

2. Fill in missing information and remove ambiguous references from context:

- (Driver) What is the weather forecast for today?
- (DM) Partly cloudy, 50% chance of rain
- (Driver) What about Ossining?
- (DM) Partly sunny, 10% chance of rain (The DM assumes that the driver means Yorktown, NY, from the earlier conversational context. Also, when the driver asks the inexplicit question “What about Ossining?” it assumes that the driver is still asking about weather.)

3. Manage failure and provide contextual, failure- dependent help and actions

- (Driver) When will we get there?
- (DM) Sorry, what did you say?
- (Driver) I asked when will we get there. The problem of instantaneous data collection could be dealt systematically by creating a learning transformation system (LT).

Examples of LT tasks are as follows:

- Monitor driver and passenger actions in the car’s internal and external environments across a network;
- Extract and record the Driver Safety Manager relevant data in databases;
- Generate and learn patterns from stored data;
- Learn from this data how Safety Driver Manager components and driver behavior could be improved and adjusted to improve Driver Safety Manager performance and improve driving safety.

B. Embedded speech recognition

Car computers are usually not very powerful due to cost considerations. The growing necessity of the conversational interface demands significant advances in processing power on

the one hand, and speech and natural language technologies on the other. In particular, there is significant need for a low-resource speech recognition system that is robust, accurate, and efficient. An example of a low-resource system that is executed by a 50 DMIPS processor, augmented by 1 MB or less of DRAM can be found. In what follows we give a brief description of the IBM embedded speech recognition system that is based on the paper Logically a speech system is divided into three primary modules: the front-end, the labeler and the decoder. When processing speech, the computational workload is divided approximately equally among these modules. In this system the front-end computes standard 13- dimensional mel-frequency cepstral coefficients (MFCC) from 16-bit PCM sampled at 11.025 KHz. Front-End Processing Speech samples are partitioned into overlapping frames of 25 ms duration with a frame shift of 15ms. A 15 ms frame-shift instead of the standard 10 ms frame-shift was chosen since it reduces the Overall computational load significantly without affecting the recognition accuracy. Each frame of speech is windowed with a Hamming window and represented by a 13 dimensional MFCC vector. We empirically observed that noise sources, such as car noise, have significant energy in the low frequencies and speech energy is mainly concentrated in frequencies above 200 Hz. The 24 triangular mel-filters are therefore placed in the frequency range [200Hz – 5500 Hz], with center frequencies equally spaced in the corresponding mel-frequency scale. Discarding the low frequencies in this way improves the robustness of the system to noise. The front-end also performs adaptive mean removal and adaptive energy normalization to reduce the effects of channel and high variability in the signal levels respectively. The labeler computes first and second differences of the 13 dimensional cepstral vectors, and concatenates these with the original elements to yield a 39- dimensional feature vector. The labeler then computes the log likelihood of each feature vector according to observation densities associated with the states of the system’s HMMs. This computation yields a ranked list of the top 100 HMM states. Likelihoods are inferred based upon the rank of each HMM state by a table lookup. The sequence of rank likelihoods is then forwarded to the decoder. The system uses the familiar phonetically-based, hidden Markov model (HMM) approach. The acoustic model comprises context-dependent sub-phone classes (all phones). The context for a given phone is composed of only one phone to its left and one phone to its right. The allophones are identified by growing a decision tree using the context-tagged training feature vectors and specifying the terminal nodes of the tree as the relevant instances of these classes. Each allophone is modeled by a single-state Hidden Markov Model with a self loop and a forward transition. The training feature vectors are poured down the decision tree and the vectors that collect at each leaf are modeled by a Gaussian Mixture Model (GMM), with diagonal covariance matrices to give an initial acoustic model. Starting with these initial sets of GMMs several iterations of the standard Baum-Welch EM training procedure are run to obtain the final baseline model. In our system, the output distributions on the state transitions are expressed in terms of the rank of the HMM state instead of in terms of the feature vector and the GMM modeling the leaf. The rank of an HMM state is obtained by computing the likelihood of the acoustic vector using the GMM at each state, and then ranking the states on the basis of their likelihoods. The decoder implements a synchronous Viterbi search over its active vocabulary, which may be changed dynamically. Words are represented as sequences of context dependent phonemes, with each phoneme modeled as a three-state HMM. The observation

densities associated with each HMM state are conditioned upon one phone of left context and one phone of right context only. A discriminative training procedure was applied to estimate the parameters of these phones. MMI training attempts to simultaneously maximize the likelihood of the training data given the sequence of models corresponding to the correct transcription, and minimize the likelihood of the training data given all possible sequences of models allowed by the grammar describing the task. In 2001, speech evaluation experiments yields improvement from 20% to 40% relatively depending on testing conditions (e.g. 7.6% error rate for 0 speed and 10.1% for 60 mph).



Figure 1: Embedded speech recognition indicator



Figure 2: Embedded Speech Recognition device

C. Driver Drowsiness Prevention

Fatigue causes more than 240,000 vehicular accidents every year. Currently, drivers who are alone in a vehicle have access only to media such as music and radio news which they listen to passively. Often these do not provide sufficient stimulation to assure wakefulness. Ideally, drivers should be presented with external stimuli that are interactive to improve their alertness. Driving, however, occupies the driver's eyes and hands, thereby limiting most current interactive options. Among the efforts presented in this general direction, the invention suggests fighting drowsiness by detecting drowsiness via speech biometrics and, if needed, by increasing arousal via speech interactivity. When the patent was granted in May 22, 2001, it received favorable worldwide media attention. It became clear from the numerous press articles and interviews on TV, newspaper and radio that Artificial Passenger was perceived as having the potential to dramatically increase the safety of drivers who are highly fatigued. It is a common experience for drivers to talk to other people while they are driving to keep themselves awake. The purpose of Artificial Passenger part of the CIT project at IBM is to provide a higher level of interaction with a driver than current media, such as CD players or radio stations, can offer. This is envisioned as a series of interactive modules within Artificial Passenger, that increase driver awareness and help to determine if the driver is losing focus. This can include both conversational dialog and interactive games, using voice only. The scenarios for Artificial Passenger currently include: quiz games, reading jokes, asking questions, and interactive books. In the Artificial Passenger (ArtPas) paradigm, the awareness-state of the driver will be monitored, and the content will be modified accordingly. Drivers evidencing fatigue, for example,

will be presented with more stimulating content than drivers who appear to be alert. This could enhance the driver experience, and may contribute to safety. Most well known emotion researchers agree that arousal (high, low) and valence (positive, negative) are the two fundamental dimensions of emotion. Arousal reflects the level of stimulation of the person as measured by physiological aspects such as heart rate, cortical activation, and respiration. For someone to be sleepy or fall asleep, they have to have a very low level of arousal. There is a lot of research into what factors increase psychological arousal since this can result in higher levels of attention, information retention and memory. We know that movement, human voices and faces (especially if larger than life), and scary images (fires, snakes) increase arousal levels. We also know that speaking and laughing create higher arousal levels than sitting quietly. Arousal levels can be measured fairly easily with a biometric glove (from MIT), which glows when arousal levels are higher (reacts to galvanic skin responses such as temperature and humidity). The following is a typical scenario involving Artificial Passenger. Imagine, driver "Joe" returning home after an extended business trip during which he had spent many late nights. His head starts to nod ... ArtPas: Hey Joe, what did you get your daughter for her birthday? Joe (startled): It's not her birthday!

ArtPas: You seem a little tired. Want to play a game?

Joe: Yes.

ArtPas: You were a wiz at "Name that Tune" last time. I was impressed. Want to try your hand at trivia?

Joe: OK.

ArtPas: Pick a category: Hollywood Stars, Magic Moments or Hall of Fame?

Joe: Hall of Fame.

ArtPas: I bet you are really good at this. Do you want the 100, 500 or 1000 dollar level?

Joe: 500

ArtPas: I see. Hedging your bets are you? By the time Joe has answered a few questions and has been engaged with the dynamics of the game, his activation level has gone way up. Sleep is receding to the edges of his mind. If Joe loses his concentration on the game (e.g. does not respond to the questions which Artificial Passenger asks) the system will Activate a physical stimulus (e.g. verbal alarm). The Artificial Passenger can detect that a driver does not respond because his concentration is on the road and will not distract the driver with questions. On longer trips the Artificial Passenger can also tie into a car navigation system and direct the driver to a local motel or hotel.



Figure 3: Condition Sensors Device

D. Workload Manager

In this section we provide a brief analysis of the design of the workload management that is a key component of driver

Safety Manager An object of the workload manager is to determine a moment-to-moment analysis of the user's cognitive workload. It accomplishes this by collecting data about user conditions, monitoring local and remote events and prioritizing message delivery. There is rapid growth in the use of sensory technology in cars. These sensors allow for the monitoring of driver actions (e.g. application of brakes, changing lanes), provide information about local events (e.g. heavy rain), and provide information about driver characteristics (e.g. speaking speed, eyelid status). There is also growing amount of distracting information that may be presented to the driver (e.g. phone rings, radio, music, e-mail etc.) and actions that a driver can perform in cars via voice control. The relationship between a driver and a car should be consistent with the information from sensors. The workload manager should be designed in such a way that it can integrate sensor information and rules on when and if distracting information is delivered. This can be designed as a "workload representational surface". One axis of the surface would represent stress on the vehicle and another, orthogonally distinct axis, would represent stress on the driver. Values on each axis could conceivably run from zero to one. Maximum load would be represented by the position where there is both maximum vehicle stress and maximum driver stress, beyond which there would be "overload". The workload manager is closely related to the event manager that detects when to trigger actions and/or make decisions about potential actions. The system uses a set of rules for starting and stopping the interactions (or interventions). It controls interruption of a dialog between the driver and the car dashboard (for example, interrupting a conversation to deliver an urgent message about traffic conditions on an expected driver route). It can use answers from the driver and/or data from the workload manager relating to driver conditions, like computing how often the driver answered correctly and the length of delays in answers, etc. It interprets the status of a driver's alertness, based on his/her answers as well as on information from the workload manager. It will make decisions on whether the driver needs additional stimuli and on what types of stimuli should be provided (e.g. verbal stimuli via speech applications or physical stimuli such as a bright light, loud noise, etc.) and whether to suggest to a driver to stop for rest. The system permits the use and testing of different statistical models for interpreting driver answers and information about driver conditions. The driver workload manager is connected to a driving risk evaluator that is an important component of the Safety Driver Manager. The goal of the Safety Driver Manager is to evaluate the potential risk of traffic accident by producing measurements related to stresses on the driver and/or vehicle, the driver's cognitive workload, environmental factors, etc. The important input to the workload manager is provided by the situation manager whose task is to recognize critical situations. It receives as input various media (audio, video, car sensor data, network data, GPS, biometrics) and as output it produces a list of situations. Situations could be simple, complex or abstract. Simple situations could include, for instance: a dog locked in a car; a baby in a car; another car approaching; the driver's eyes are closed; car windows are closed; a key is left on a car seat; it is hot in a car; there are no people in a car; a car is located in New York City; a driver has diabetes; a driver is on the way home. Complex situations could include, for example: a dog locked is locked in a car AND it is hot in a car AND car windows are closed; a baby is in a car AND there are no people in a car; another car is approaching AND the driver is looking in the opposite direction; a key is left on a car seat AND a driver is in the

midst of locking a car; the driver is diabetic AND has not taken a medicine for 4 hours. Abstract situations could include, for example: Goals: get to work, to cleaners, to a movie. Driver preferences: typical routes, music to play, restaurants, shops. Driver History accidents, illness, visits. Situation information can be used by different modules such as workload, dialog and event managers; by systems that learns driver behavioral patterns, provide driver distraction detection, and prioritize message delivery. For example, when the workload manager performs a moment-to-moment analysis of the driver's cognitive workload, it may well deal with such complex situations as the following: Driver speaks over the phone AND the car moves with high speed AND the car changes lanes; driver asks for a stock quotation AND presses brakes AND it is raining outside; another car approaches on the left AND the driver is playing a voice interactive game. The dialog manager may well at times require uncertainty resolution involving complex situations, as exemplified in the following verbal query by a driver: "How do I get to Spring Valley Rd?" Here, the uncertainty resides in the lack of an expressed (geographical) state or municipality. The uncertainty can be resolved through situation recognition; for example, the car may be in New York State already (that is defined via GPS) and it may be known that the driver rarely visits other states. The concept associated with learning driver behavioral patterns can be facilitated by a particular driver's repeated routines, which provides a good opportunity for the system's "learning" habitual patterns and goals. So, for instance, the system could assist in determining whether drivers are going to pick up their kids in time by, perhaps, reordering a path from the cleaners, the mall, the grocery store, etc.

IV. WORKING OF ARTIFICIAL PASSENGER

The AP is an artificial intelligence-based companion that will be resident in software and chips embedded in the automobile dashboard. The heart of the system is a conversation planner that holds a profile of you, including details of your interests and profession. When activated, the AP uses the profile to cook up provocative questions such as, "Who was the first person you dated?" via a speech generator and in-car speakers. A microphone picks up your answer and breaks it down into separate words with speech-recognition software. A camera built into the dashboard also tracks your lip movements to improve the accuracy of the speech recognition. A voice analyzer then looks for signs of tiredness by checking to see if the answer matches your profile. Slow responses and a lack of intonation are signs of fatigue. If you reply quickly and clearly, the system judges you to be alert and tells the conversation planner to continue the line of questioning. If your response is slow or doesn't make sense, the voice analyzer assumes you are dropping off and acts to get your attention. The system, according to its inventors, does not go through a suite of rote questions demanding rote answers. Rather, it knows your tastes and will even, if you wish, make certain you never miss Paul Harvey again. This is from the patent application: "An even further object of the present invention is to provide a natural dialog car system that understands content of tapes, books, and radio programs and extracts and reproduces appropriate phrases from those materials while it is talking with a driver. For example, a system can find out if someone is singing on a channel of a radio station. The system will state, "And now you will hear a wonderful song!" or detect that there is news and state, "Do you know what happened now—hear the following—and play some news." The system also includes a recognition system to detect who is speaking over the radio and alert the driver if the person speaking is one the

driver wishes to hear.”Just because you can express the rules of grammar in software doesn’t mean a driver is going to use them. The AP is ready for that possibility: “It provides for a natural dialog car system directed to human factor engineering—for example, people using different strategies to talk (for instance, short vs. elaborate responses). In this manner, the individual is guided to talk in a certain way so as to make the system work—e.g., “Sorry, I didn’t get it. Could you say it briefly?” Here, the system defines a narrow topic of the user reply (answer or question) via an association of classes of relevant words via decision trees. The system builds a reply sentence asking what are most probable word sequences that could follow the user’s reply.” Driver fatigue causes at least 100,000 crashes, 1,500 fatalities, and 71,000 injuries annually, according to estimates prepared by the National Highway Traffic Safety Administration, which estimated further that the annual cost to the economy due to property damage and lost productivity is at least \$12.5 billion. The Federal Highway Administration, the American Trucking Association, and Liberty Mutual co-sponsored a study in 1999 that subjected nine volunteer truck drivers to a protracted laboratory simulation of over-the-road driving. Researchers filmed the drivers during the simulation, and other instruments measured heart function, eye movements, and other physiological responses. “A majority of the off-road accidents observed during the driving simulations were preceded by eye closures of one-half second to as long as 2 to 3 seconds,” Stern said. A normal human blink lasts 0.2 to 0.3 second. Stern said he believes that by the time long eye closures are detected, it’s too late to prevent danger. “To be of much use,” he said, “alert systems must detect early signs of fatigue, since the onset of sleep is too late to take corrective action.” Stern and other researchers are attempting to pinpoint various irregularities in eye movements that signal oncoming mental lapses—sudden and unexpected short interruptions in mental performance that usually occur much earlier in the transition to sleep. “Our research suggests that we can make predictions about various aspects of driver performance based on what we glean from the movements of a driver’s eyes,” Stern said, “and that a system can eventually be developed to capture this data and use it to alert people when their driving has become significantly impaired by fatigue.” He said such a system might be ready for testing in 2004.



Figure 4: Eye Tracker



Figure 5: Camera for detection of Lips Movement
Determination of eye blinking and closing

V. FEATURES OF ARTIFICIAL PASSENGER

Telematics combines automakers, consumers, service industries. And computer companies. Which is how IBM fits in. It knows a thing or two about computers, data management, and connecting compute devices together. So people at its Thomas J. Watson Research Center (Hawthorne, NY) are working to bring the technology to a car or truck near you in the not-too-distant future.

A. Conversational Telematic

IBM’s Artificial Passenger is like having a butler in your car—someone who looks after you, takes care of your every need, is bent on providing service, and has enough intelligence to anticipate your needs. This voice-actuated telematics system helps you perform certain actions within your car hands-free: turn on the radio, switch stations, adjust HVAC, make a cell phone call, and more. It provides uniform access to devices and networked services in and outside your car. It reports car conditions and external hazards with minimal distraction. Plus, it helps you stay awake with some form of entertainment when it detects you’re getting drowsy. In time, the Artificial Passenger technology will go beyond simple command-and-control. Interactivity will be key. So will natural sounding dialog. For starters, it won’t be repetitive (“Sorry your door is open, sorry your door is open . . .”). It will ask for corrections if it determines it misunderstood you. The amount of information it provides will be based on its “assessment of the driver’s cognitive load” (i.e., the situation). It can learn your habits, such as how you adjust your seat. Parts of this technology are 12 to 18 months away from broad implementation.

B. Improving Speech Recognition

You’re driving at 70 mph, it’s raining hard, a truck is passing, the car radio is blasting, and the A/C is on. Such noisy environments are a challenge to speech recognition systems, including the Artificial Passenger. IBM’s Audio Visual Speech Recognition (AVSR) cuts through the noise. It reads lips moment speech recognition. Cameras focused on the driver’s mouth do the lip reading; IBM’s Embedded via Voice does the speech recognition. In places with moderate noise, where conventional speech recognition has a 1% error rate, the error rate of AVSR is less than 1%. In places roughly ten times noisier, speech the recognition has about a 2% error rate; AVSR’s is still pretty good (1% error rate). When ambient noise is just as loud as the driver talking, speech recognition loses about 10% of words; AVSR, 3%. Not great, but certainly usable.

C. Analyzing Data

The sensors and embedded controllers intoday’s cars collect a wealth of data. The next step is to have them “phone home,” transmitting that wealth back to those who can use those data.

Making sense of that detailed data is hardly a trivial matter, though especially when divining transient problems or analyzing data about the vehicle's operation over time. IBM's Automated Analysis Initiative is a data management system for identifying failure trends and predicting specific vehicle failures before they happen. The system comprises capturing, retrieving, storing, and analyzing vehicle data; exploring data to identify features and trends; developing and testing reusable analytics and evaluating as well as deriving corrective measures. This sort of technology has helped Peugeot diagnose and repair 90% of its cars within four hours, and 80% of its cars within a day (versus days). An Internet-1999based diagnostics server reads the car data to determine the root cause of a problem or lead the technician through a series of tests. The server also takes a "snapshot" of the data and repair steps. Should the problem reappear, the system has the fix readily available.

D. Sharing Data

Collecting dynamic and event-driven data is one problem. Another is ensuring data security, integrity, and regulatory compliance while sharing that data. For instance, vehicle identifiers, locations, and diagnostics data from a fleet of vehicles can be used by a variety of interested, and sometimes competitive, parties.

VI. APPLICATIONS

- Artificial passenger
- Artificial passenger is broadly used to prevent accidents.
- In any problem it alerts the vehicles near by so the driver there can become alert.
- Prevents the driver falling asleep during long and solo trip.
- Opens and closes the doors and windows of the car automatically.
- It is also used for the entertainment.
- Also used in cabins in airplanes, trains, boats etc.

Future Implementation

Will provide us with shortest time routing based on road conditions changing because of weather and traffic,

information about the cars on the route, destination requirement.

CONCLUSIONS

We suggested that such important issues related to a driver safety, such as controlling Telematics devices and drowsiness can be addressed by a special speech interface. This interface requires interactions with workload, dialog, event, privacy, situation and other modules. We showed that basic speech interactions can be done in a low-resource embedded processor and allows a development of a useful local component of Safety Driver Manager. The reduction of conventional speech processes to low – resources processing was done by reducing a signal processing and decoding load in such a way that it did not significantly affect decoding accuracy and by the development of quasi-NLU principles. We observed that an important application like Artificial Passenger can be sufficiently entertaining for a driver with relatively little dialog complexity requirements – playing simple voice games with a vocabulary containing a few words. Successful implementation of Safety Driver Manager would allow use of various services in cars (like reading e-mail, navigation, downloading music titles etc.) without compromising a driver safety. Providing new services in a car environment is important to make the driver comfortable and it can be a significant source of revenues for Telematics. We expect that novel ideas in this paper regarding the use of speech and distributive user interface in Telematics will have a significant impact on driver safety and they will be the subject of intensive research and development in forthcoming years at IBM and other laboratories.

References

- [1] L.R. Bahl
- [2] Lawrence R. Rabiner, A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition.
- [3] Kanevsky D. Telematics: Artificial Passenger and beyond, *Human Factors and Voice Interactive Systems*, Signals and Communications Technology Series.