

2017 Consumer Electronics Show (CES2017) Press Conference - Dr. Gill Pratt

January 04, 2017

Image not found or type unknown



Highlights

“How safe is safe enough? Society tolerates a lot of human error.
We are, after all, only human. But we expect machines to be much better. “

“What if the machine was **twice as safe** as a human-driven car and 17.5 thousand lives were lost in the US every year? Would we accept such autonomy? Historically, humans have shown nearly zero-tolerance or injury or death caused by flaws in a machine.”

“None of us in the automobile or IT industries are close to achieving true level 5 autonomy. It will take many years of machine learning and many more miles than anyone has logged of both simulated ...and real-world testing to achieve the perfection required for Level 5 autonomy.”

“Considerable research shows that the longer a driver is disengaged from the task of driving, the longer it takes to re-orient. “

“It is possible that level 3 may be as difficult to accomplish as level 4.”

“There is evidence that some drivers may deliberately test the (L2) system’s limits...essentially mis-using a device in a way it was not intended to be used.”

“Human nature, not surprisingly, remains one our biggest concerns.”

“There are indications that many drivers over-trust the (L2) system.”

“When someone over-trusts a level 2 system’s capabilities they may mentally disconnect their attention from the driving environment and wrongly assume the level 2 system is more capable than it is. We at TRI worry that over-trust may accumulate over many miles of handoff-free driving.””

“”Paradoxically, the less frequent the handoffs, the worse the tendency to over-trust may become.”

“TRI has been taking a two-track approach, simultaneously developing a system we call Guardian, designed to make human driving safer...while working on L4 and 5 systems that we call Chauffeur. The perception and planning software in Guardian and Chauffeur are basically the same. The difference is that Guardian only engages when needed, while Chauffeur is engaged, all of the time during an autonomous drive.“

In Guardian, the driver is meant to be in control of the car at all times **except** in those cases where Guardian anticipates or identifies a pending incident and briefly employs a corrective response.

As prepared for:
2017 Consumer Electronics Show
Las Vegas
Wednesday, Jan. 4, 2017
Dr. Gill Pratt, Toyota Research Institute CEO

[Bob Carter Introduction]

Thanks Bob...and thank you as well, Yui.

Good afternoon everyone and welcome again to our press conference and the one-year anniversary of the Toyota Research Institute.

TRI's job is to explore those emerging opportunities,

- to identify gaps in research
- to move in new directions ...
- to sometimes find a dead end....
- but always move quickly
- to explore great ideas.

In the one year of our existence, we have assembled an amazing team.

We have hired more than 100 people, added 50 more people from Toyota Motor Corporation... and plan to hire another one hundred people this year.

TRI's mission is focused on Artificial Intelligence and includes four goals:

- First, to greatly enhance **vehicle safety** and someday create a car incapable of causing a crash.
- Second, to greatly increase **mobility access**...for those who cannot drive
- Third, to heavily invest in **robotics** ...to move people not just across town...but in their home...from room to room.
- And finally, to accelerate **discovery in materials science** by applying techniques from artificial intelligence and machine learning.

Research is our middle name.

That's what we do, much of it... but not all of it with Stanford, in Palo Alto, California the University of Michigan, in Ann Arbor and MIT, in Cambridge, Massachusetts, where we recently celebrated the opening of our permanent facility.

Before we go any further, I must remind you that I used to be a professor...and probably always will be.

As a professor I always enjoyed the process of both doing research and giving lectures.

In fact, there will be a short test near the end of my remarks.

My remarks today reflect-on findings from a few key research projects we and our partners have been conducting this past year.

They are framed by a question, designed to offer clarity and provoke discussion on just how complicated this business of autonomous mobility, really is.

The question I'd like to discuss with you today is: ***How safe is safe enough?***

Society tolerates a lot of human error. We are, after all, "*only human.*"

But we expect machines to be much better.

Last year, there were about 35,000 fatalities on US highways...**involving** vehicles controlled by human drivers.

Every single one of those deaths is a tragedy.

What if we could create a fully autonomous car that was "*as safe, **on average***" as a human driver...would that be safe enough?

In other words, would we accept:

- 35,000 traffic fatalities a year in the US
- at the hands of a machine;
- if it resulted in greater convenience,
- less traffic,
- and less impact on the environment?

Rationally, perhaps the answer should be yes.

But emotionally, we at TRI don't think it is likely that being "as safe as a human being" will be acceptable.

However, what if the machine was **twice as safe** as a human-driven car and 17,500 lives were lost in the US every year?

Would we accept such autonomy?

Historically, humans have shown nearly zero-tolerance for injury or death caused by flaws in a machine.

And yet we know that the artificial intelligence systems on which our autonomous cars will depend are presently and unavoidably, imperfect.

"So...How safe is safe enough?"

In the very near future, this question will need an answer.

We don't yet know for sure.

Nor is it clear how that standard will be devised.

And by whom.

And will it be the same globally?

One standard that is already in place...is the SAE International J3016...revised just last September...that defines five levels of driving automation.

I want to review this standard with you because there continues to be a lot of confusion in the media about it.

All car makers are aiming to achieve level 5, where a car can drive fully autonomously under any traffic or weather condition in any place and at any time.

I need to make this perfectly clear: This is a wonderful goal.

However, none of us in the automobile...or IT industries are close to achieving true level 5 autonomy.

Collectively, our current prototype autonomous cars can handle many situations.

But there are still many others that are beyond current machine competence.

It will take many years of machine learning and many more miles than anyone has logged of both *simulated ...* and *real-world testing to achieve the perfection* required for Level 5 autonomy.

But there is good news.

SAE Level 4 autonomy is ALMOST level 5, but with a much shorter timetable for arrival.

Level 4 is fully autonomous except that it only works in a specific **Operational Design Domain**...like the MCity test facility on the campus of the University of Michigan.

Restrictions could include

- limited areas of operation...
- limited speeds,
- limited times of day
- and only when the weather is good.

When company A, or B...or T says it hopes to have autonomous vehicles on the road by early 2020s, level 4 is the technology they are probably referring to.

TRI believes it is likely that a number of manufacturers will have level 4 autonomous vehicles operating in specific locations within a decade.

Level 4 autonomy will be especially attractive and adaptable for companies offering...**Mobility as a Service**... in such forms as ride-sharing and car-sharing...and inner-city last-mile models

In fact, Mobility as a Service may well offer the best application for bringing Level 4 to market sooner, rather than later.

Moving down the ladder, Level 3 is a lot like level 4, but with an autonomous mode that at times may need to **hand-off** control to a human driver who may not be paying attention at the time.

Hand-off, of course, is the operative term ...and a difficult challenge.

In level 3, as defined by SAE, the autonomy must **ensure** that if it needs to hand-off control of the car it will give the driver sufficient warning.

Additionally... level 3 autonomy must also ensure that it will always detect any condition requiring a handoff.

This is because in level 3, the driver is not required to oversee the autonomy, and may instead fully engage in other tasks.

The term used by SAE when the vehicle's system cannot handle its dynamic driving tasks, is a ***request to intervene***.

The challenge lies in how long it takes a human driver to disengage from their texting or reading once this fallback intervention is requested...(pause) and also...whether the system can ensure...that it will never miss a situation... where a handoff is required.

Considerable research shows that the longer a driver is disengage from the task of driving, the longer it takes to re-orient.

Furthermore, at 65 miles per hour, a car travels around 100 feet every second.

This means that to give a disengaged driver 15 seconds of warning, at that speed... the system must spot trouble, about 1500 feet away or about 5 football fields ahead.

That's extremely hard to guarantee, and unlikely to be achieved soon.

Because regardless of speed, a lot can happen in 15 seconds, so ensuring at least 15 seconds of warning is very difficult.

In fact, it is possible that level 3 may be as difficult to accomplish as level 4.

This brings us to level 2, perhaps the most controversial right now because it's already here and functioning in some cars on public roads.

In level 2, a vehicle hand-off to a human driver may occur at any time with only a second or two of warning.

This means the human driver must be able to react, mentally and physically at a moment's notice.

Even more challenging is the requirement for the Level 2 human driver to always supervise the operation of the autonomy taking over control when the autonomy fails to see danger ahead.

It's sort of like tapping on the brake to disengage adaptive cruise control when we see debris in the road that the sensors do not detect.

This can and will happen in level 2 and we must never forget it.

Human nature, not surprisingly, remains one of our biggest concerns,

There are indications that many drivers, may either *under-trust* or *over-trust* a system.

When someone over-trusts a level 2 system's capabilities...they may mentally disconnect their attention from the driving environment...and wrongly assume the level 2 system is more capable than it is.

We at TRI worry that over-trust may accumulate over many miles of handoff-free driving.

Paradoxically, the less frequent the handoffs, the worse the tendency to over-trust may become.

And there is also evidence that some drivers may deliberately test the system's limits...essentially misusing a device in a way it was not intended to be used.

This is a good time to address *situational awareness and mental attention*

It turns out that maintaining awareness while engaged in monitoring tasks has been well-studied for nearly 70 years.

Research psychologists call it...the "Vigilance Decrement".

During World War Two, it became clear that radar operators looking for enemy movement became less effective as their shift wore on, even if they kept their eyes on the task.

In 1948, Norman Mackworth wrote a seminal paper called "*The breakdown of vigilance during prolonged visual search*"

The experiment he performed used a clock that only had a second hand that would occasionally and randomly jump by two seconds.

Turns out that, even if you keep your eyes on the MacWorth clock, as this graph shows, your performance at detecting two-second jumps will decrease in proportion to how long you do it.

OK, as promised, here is the 20-second test I warned you about, **earlier**.

Watch the hand of the MacWorth clock carefully.

Every time the hand bumps two seconds instead of one second, clap your hands.

OK, Here we go.

Ah, well, what was that, half the class?

That's a bit better.

Ok, so how do you think you would do at this task for two hours?

Are you likely to remain vigilant... for a possible handoff...of the Level 2 car's autonomy?

Does this body of evidence mean that level 2 is a bad idea?

Some companies have already decided the challenges may be too difficult, and have decided to skip levels 2 and 3.

As it turns out we are finding evidence that some things...texting not included...seem to reduce vigilance decrement.

We are finding that some MILD secondary tasks may actually help maintain situational awareness.

For example, long-haul truck drivers have extremely good safety records, comparatively.

How do they do it?

Perhaps because they employ mild secondary tasks that help keep them vigilant.

They talk on two-way radios and may scan the road ahead...looking for speed traps.

And I bet almost all of us have listened to the radio as a way of staying alert during a long drive.

Experts have divided opinions on whether that is a good idea or a bad one.

As Bob said earlier, the human/machine interface-and-relationship are extremely important at Toyota.

We at TRI continue to explore.

What we do know —for sure— is that as we move forward...towards the ultimate goal of full autonomy, we must strive to save as many lives as possible in the process.

Because, it will take decades to have a significant portion of the US car fleet functioning at Level 4 and above.

That's why TRI has been taking a two-track approach, simultaneously developing a system, we call *Guardian*, designed to make human driving safer... while working on Level 2 through Level 5 systems that we call *Chauffeur*.

Much of the work in hardware and software that we are developing to achieve *Chauffeur*, is also applicable to *Guardian*.

And visa-versa.

In fact, the perception and planning software in *Guardian* and *Chauffeur* are basically the same.

The difference is that *Guardian* only engages when needed, while *Chauffeur* is engaged, all of the time during an autonomous drive.

This video demonstrates testing of an early prototype, *Guardian* simulator at our temporary facility in Palo Alto.

I am happy to say we are now working on a new, highly advanced simulator which we will show-off soon... as well as a new permanent facility in Palo Alto that we are currently in the process of moving in to.

One can think of anti-lock brakes, vehicle stability control and automatic emergency braking, as early forms of Guardian.

When it arrives, it will be

- a hands-on-the-wheel,
- eyes-on-the-road
- only-when-needed system...
- merging vehicle and human
- situational awareness.

In Guardian, the driver is meant to be in control of the car at all times **except** in those cases where Guardian anticipates or identifies a pending incident and briefly employs a corrective response.

Depending on the situation, Guardian can alert the driver with visual cues and audible alarms, and if necessary influence or control speed and steering.

Like, *Yui*, our **Concept i agent** Guardian employs artificial intelligence and becomes smarter and smarter through both first-hand data-gathering experience and by intelligence shared via the cloud.

Over time, we expect Guardian's growing intelligence will allow it to sense things more clearly and quickly process and anticipate faster... and respond more accurately in a wider array of situations.

Every year cars get safer.

One reason is because every year, automakers equip vehicles with higher and higher levels of active safety.

In ever-increasing numbers, vehicles are already being entrusted...to sense a problem, choose a course of action and respond...assuming, for brief periods, control of the vehicle

And that brings me back to the Concept i.

At TRI, we think that "YUI"...the Concept i *agent* might not only be a way to engage and provide useful advice.

We think it might also be a way to promote the driver's continued situational awareness using mild secondary tasks to promote safety.

We've only begun our research to find out exactly how that would work.

Perhaps YUI could engage the driver in a conversation, that would reduce the vigilance decrement the way talking on the two-way radio or looking for speed traps seems to do with truck drivers.

We think the agent might even be more effective, because the *Yui* would be coupled to the autonomy system, which would be constantly monitoring the car's environment, inside and out merging human and vehicle situational awareness

We're not sure, but we aim to find out.

Toyota is involved in many aspects of making future cars safer and more accessible.

YUI and Concept-i is a small part of that work.

But it has the potential for being more than a helpful friend.

It may have the potential to become the kind of friend that looks out for you, and keeps you safe.

A guardian, as well as a chauffeur.

Our goal is to someday create a car that will never be responsible for causing a crash, whether it is driven by a human being or by a computer.

And Concept-i may become a key part of that plan.