

60 Seconds with ...

Prof. Philip Koopman



Phil is an internationally recognised expert on Autonomous Vehicle (AV) safety whose work in that area spans over 25 years. He is also actively involved with AV policy and standards as well as more general embedded system design and software quality.

His pioneering research work includes software robustness testing and run-time monitoring of autonomous systems to identify how they break and how to fix them. He has extensive experience in software safety and software quality across numerous transportation, industrial, and defence application domains including conventional automotive software and hardware systems. He was the principal technical contributor to the UL 4600 standard for autonomous system safety issued in 2020.

He is a faculty member of the Carnegie Mellon University Electrical and Computer Engineering department, where he teaches software skills for mission-critical systems. In 2018 he was awarded the highly selective IEEE-SSIT Carl Barus Award for outstanding service in the public interest for his work and in 2022, he was named to the National Safety Council's Mobility Safety Advisory Group. He is the author of the book *How Safe is Safe Enough: Measuring and Predicting Autonomous Vehicle Safety* (2022).

What first attracted you to working in the field of System Safety?

It has been more collecting pieces over time than a conscious plan to work in the field of System Safety. As a teenager I taught bicycle safety to school kids as a public service project. After college, wearing a personal radiation dosimeter 24 hours a day in a sealed environment during my Navy submarine service certainly focused my attention on safety.

I got an opportunity to be the dependability/safety expert on an early autonomous vehicle project at Carnegie Mellon University in the mid-1990s by being the right person at the right place at the right time. Since then I've been fortunate to have opportunities to work with safety in various domains including rail, industrial controls, commercial buildings, consumer electronics, power systems, and conventional automotive by doing industry design reviews. I have also done some product defect litigation expert work.

What aspect of your career are you most proud of?

For me it has always been about finding a way to make a difference. Most recently, I think leading the development of the ANSI/UL 4600 autonomous vehicle safety standard has pushed the conversation forward on autonomous vehicle standards and safety cases in both the technical and public policy arenas.

The professional recognition I've received that is most important to me is the IEEE SSIT Carl Barus award for my work on computer-based system defects in automotive applications. Receiving that award is inherently bittersweet because it is only given to people who persevere amidst controversy, typically involving public safety.

What advice would you give to yourself age 12?

If it is worth doing, it is worth taking your best shot.

What excites (or frightens!) you the most about the future of Autonomous Vehicles?

The deployment governance model, especially in the US, is dominated by big money pressing for quick results. This pressure has done short-term harm via needlessly increasing the risk of individual loss events. But just as importantly, long-term harm is being done in the areas of undermining regulatory authority, exclusion of legitimate stakeholders from deployment decisions, and eventual disillusionment with the technology when hype collapses.

What's your most favourite quote or motto?

Illegitimi non carborundum.

If you could learn to do anything, what would it be?

The next thing I need to learn to continue on my path. Every day is a good one if I'm learning.

What's the best piece of advice you've ever been given?

Find a way to share your talents.

In a world of conflict, austerity and climate change, which work of art or fiction best sums up your current outlook in life?

I posed that question as a prompt to "DALL·E 2" (a new AI system that can create realistic images and art from a description in natural language), and it generated this "work of art" on the first try.

Here is an "AI" (really it is not intelligent at all) that generated a complete nonsense response to a reasonable question. But the nonsense is captured in a format that people are accustomed to accepting as deep wisdom, and might even pass if you don't look long enough to try to figure out what language it is written in. (Google translate says this is not a language it knows. But that is an "AI" too I suppose).

I think there is a cautionary message here for those who argue that machine learning will inevitably be a superior replacement for human judgment in any particular application.

