

Unmanned vehicles for populated environments

Ken Wahren

Navigating an unmanned vehicle in areas with people will require a competent control system whose behaviour is assured to be safe.

Unmanned vehicles have demonstrated tremendous capabilities in remote and hostile environments, where they are used for tasks such as military reconnaissance, space exploration and subsurface ocean monitoring. However, such vehicles are not yet widespread in heavily populated environments. In these settings, there is typically a well-established infrastructure to support the operation of manned vehicles, and vehicles are extensively regulated by an authority.

Because existing regulations in these environments have evolved alongside manned vehicle technologies and their supporting infrastructure, they are generally poorly matched to the unique technical risks and opportunities presented by unmanned vehicles. To mitigate the risk introduced by this mismatch, unmanned vehicles are currently permitted to operate in populated settings only on a case-by-case basis, with special conditions enforced on their operation. In many cases, these special conditions negate the potential benefits of the vehicle being unmanned, for example, by limiting operations geographically, making stipulations for a chase vehicle or requiring that a human safety operator be on-board the vehicle at all times. Regulators and industry are therefore engaged in the challenging work of co-evolving regulatory policy and unmanned vehicle technology to accommodate unmanned vehicles in populated environments without special conditions. The implicit aim is to preserve the safety goals of existing regulations for manned vehicles while avoiding costly changes to the supporting infrastructure.

Three types of unmanned vehicle are currently at a very interesting stage of integration with populated environments: unmanned air vehicles (UAVs), unmanned surface vehicles (USVs) and unmanned ground vehicles (UGVs)—see Figures 1, 2 and 3. Although their modes of locomotion and target environments are very different, they all share the challenge of integrating with open, norm-governed environments designed for humans operating manned vehicles.



Figure 1. The Hermes 450 unmanned air vehicle, a military system with commercial prospects in civilian applications. (Source: Wikimedia Commons.)

My current work addresses two of the key concerns of regulators in every domain. First, ensuring that vehicles will avoid collisions with at least the same level of competence as that demonstrated by manned vehicles. Second, ensuring that the necessarily complex control systems of unmanned vehicles are developed to the same rigorous functional safety standards as any safety-critical embedded systems approved for manned vehicles. Each of these concerns has nuances that make the task much more challenging than it might first appear.

For example, the solution to the first concern, collision avoidance, might be construed as simply encoding behaviour written into a rule book (normative behaviour). In the case of an unmanned car, this might mean giving way at a roundabout, not driving over double lines, stopping at stop signs, obeying the speed limit and so forth. Understanding and following rules is a good start, but there is much more to being a good driver than following rules, and the road transport system actually relies heavily on people being good drivers. This poses a problem, because achieving the same level of competence as a manned vehicle becomes a poorly defined target.

Continued on next page

I am investigating how two concepts from cognitive psychology could be used to develop a competent navigation scheme. The first is prospect theory,¹ a descriptive theory of human decision making which holds that utility is affected independently by a value attached to an outcome and the probability of that outcome. It has relevance whenever an unmanned vehicle might have to choose between several unfavourable outcomes such as crossing the double lines to avoid a collision with a rogue driver who cuts off the vehicle. The second is the theory of affordances,² which holds that people can directly perceive action possibilities associated with features in an environment. For example, the handle on a cup affords holding. It is suspected that this ability might give people an edge when navigating in crowded environments.



Figure 2. An unmanned surface vehicle in action. (Source: Wikimedia Commons.)



Figure 3. Google driverless car. (Source: Wikimedia Commons.)

It is unlikely that there will be a silver bullet for the second key concern (ensuring control systems conform with rigorous safety standards), but I am examining several ways that the task of certifying complex unmanned vehicle control software can be kept as simple as possible. First, it will help if the architecture affords embedding in a safety case of the sort that typically forms the starting point for safety-critical software development according to standards for aircraft (DO-178C) or automobiles (ISO 26262). An architecture that enables a regulator to see behaviour would be an advantage. To that end, the behaviour trees formalism, which is popular in the design of computer game artificial intelligence for that very reason, is being tested as the executive layer in an unmanned vehicle control architecture. Second, it may help if the optimization and safety of action selection can be separated at some level. For example, a control system might compute the quickest route from one point to another using some fairly complex algorithms, which on their own might be difficult to certify. However, if it can be shown that the optimal navigation decisions are vetted for safety in real time by a separate component, it may be that the safety of the entire system can be argued based on the integrity of the vetting component. Although this type of vetting might be feasible for some kinds of decisions, such as those to deploy weapons by an autonomous system,³ it is likely that navigating an unmanned vehicle safely in the populated environment will require a combination of vetting and assuring bottom-up safety in the continuous-time behaviour of the control system.

Continued on next page

In summary, I am working on applying concepts from cognitive psychology to develop competent control systems for unmanned vehicles and I am examining how to assure the safe behaviour of such complex systems.

Author Information

Ken Wahren

Blue Bear Systems Research
Bedford, UK

References

1. D. Kahneman and A. Tversky, *Prospect theory: an analysis of decision under risk*, *Econometrica* **47**, pp. 263–291, 1979.
2. J. J. Gibson, *The theory of affordances*, in R. Shaw and J. Bransford (eds.), **Perceiving, Acting, and Knowing: Toward an Ecological Psychology**, Wiley, New York, 1977.
3. R. C. Arkin, *Governing lethal behavior: embedding ethics in a hybrid deliberative/reactive robot architecture. Part I: Motivation and philosophy*, **Proc. 3rd ACM/IEEE Int'l. Conf. Human-Robot Interac.**, pp. 121–128, 2008.