A perception-based approach to robotics and artificial life

Rupert Young

Intelligent robots that operate convincingly within the dynamic chaos of the real world are far from being a reality. For artificial systems to behave anything like living systems, and avoid being mired in complexity, we need to fundamentally rethink the standard view of what behaviour is all about[1]. Living systems employ novel and parsimonious methods which resolve and dissolve many of the seemingly intractable problems faced by conventional robotics[1].

Generally, standard robotics, including robot arms and self-driving cars, has been viewed as an engineering problem of the geometric manipulation of objects within three-dimensional space, which requires definition of the equations of motion, within the robotic system, that govern the physics of those objects and worlds.

Moving a robotic arm normally requires predicting in advance what joint angles are required to form a particular pose, for which we need equations to compute those angles, along with parametric knowledge such as mass, gravity and length of the limbs. Self-driving cars require detailed mapping of the environment, and algorithms to plan and move objects through its virtual worlds. The trajectories and models are continually adjusted by checking that the real, perceived world matches the predictions made by the models.

That might seem like a valid approach initially, but the problem becomes increasingly complicated as the environments become more complex, such as in the real world.

Humans and animals work in a very different way. Rather than trying to predict and compute the actions required to achieve a task, people vary their actions in order to perceive the world as they want. In other words, behaviour is a process of controlling perceptual input.

For example, catching a baseball is not about the computation of trajectories and intercept points, but merely about keeping the speed of the ball on the retina constant. Operation of a robot arm can be achieved by a multitude of simple controllers[2], controlling perceptual inputs such as relationships between joint angles, rather than by complex kinematics computations. As parametric knowledge is not required, parameters such limb length can be changed without any adverse effect on the robotic control system.

When driving, we turn the steering wheel to maintain our perception of the car between the white lines. We do not turn the wheel to a specific angle or by a specific amount, but until we perceive the car between the lines.

There are many factors that can affect the heading of the car; wheel balance, tyre pressure, rain, road surface and especially wind, but we do not need to know anything about them as we simply counteract their combined effects on the perceived position of the car. So, with the perceptual control theory (PCT) approach there is no need to model the transfer function between, say, the steering wheel and the heading of the car. In fact, we could go as far to say that there is no transfer function that could be modelled anyway. At least not in practice, unless you are some sort of Laplacian demon that knows the entire state and dynamics of the universe.

However, you may say, in some cases transfer functions can be defined and robotic systems can be developed in that way. Well, yes, that is true in simulated environments. In a simulation the transfer function between the steering wheel and the heading can easily be defined, and a car can be controlled perfectly by this approach. That is because with simulations we are acting as if we are Laplacian demons, as it us who are defining the universe in which the simulation is running.

It is similar with controlled environments such as the factory floor or the laboratory, as the uncertainty is limited and managed, allowing us to define some relatively modest models. However, in the real world we cannot do this, we cannot be Laplacian demons, and there are no transfer functions to model. This is why robots have been stuck in predictable, structured environments and why the modelling approach is not viable in the real world. The fact that the conventional approach works in simulation has misled researchers to think that it is valid in the real world.

The perceptual control system is a negative feedback process where the error between the goal and current perception drives the output action. By not requiring knowledge of the relationship between input and output, perceptual control systems overcome the complexity problem associated with predictive modelling. Although it seems as if we do prediction, it is only in the sense that we can set perceptual goals in advance of acting, not in the sense of predicting output by way of internal simulations of the physical world. The conventional approach may work in restricted scenarios, but even then perceptual control shows that conventional approach is unnecessarily complex.

There is another way that living systems overcome complexity which is with the perception and control of high-level invariants. This is achieved with a hierarchical architecture of perceptual control systems. The higher up the hierarchy the more complex the perceptions, and more psychological in the sense that they do not correspond directly to physical properties or objects of the real world.

This is directly relevant to self-driving cars and represents a good reason to be wary of the ongoing hype around the imminent arrival of autonomous vehicles on our roads. There is more to driving than low-level object manipulation. It also requires dynamic response to novel circumstances as well as psychological interpretations of the intentions of other control systems (that's other drivers to you and me). So far, self-driving cars have largely been restricted to test environments. Unless the systems are able to incorporate understanding and control of high-level psychological perceptions; then that is where they will remain.

Any task can be decomposed into a set of perceptual goals. Although the basic process is simple, there is a lot more going on in an everyday task than might be appreciated[3].

For decades the computational approach has been the predominant paradigm in the field of robotics and artificial life (ALife). Although this has been successful in the restricted environments of automation, it has largely been a failure in dynamic environments. If this deadlock is to be broken, it is vital that students are exposed to other potential methodologies rather than being saddled with the status quo. It should not be construed from any apparent simplicity of an approach that it is less intellectually challenging. On the contrary, it may actually require more creative ideas than the, seemingly, more obvious traditional approach.

The perceptual control approach does provide a different way of thinking about behaviour, with a simple process that is applicable to all types and levels of behaviour. Demonstrating that complex internal models of the world are not required and that perception is the goal of behaviour rather than actions has profound implications not only for our understanding of ourselves as human beings, but also for behavioural sciences and robotics. The resultant control architecture is parsimonious and computationally lightweight, making it ideal for implementing in artificial systems.

For Robotics and artificial life the benefits are significant:

Simplicity – The basic unit of perceptual control systems has a simple operation, of negative feedback control, avoiding the complexity of the conventional computational and modelling approach.

Universality – The basic perceptual control process is common to all levels and types of behaviour.

Scalability – The arrangement of basic control units into an interdependent hierarchy results in a highly scalable architecture of simple units.

Purpose – The structure of the basic process of negative feedback bringing perceptual input into line with a desired state provides the system with innate, internal motivation, a necessary characteristic of purposeful systems.

Adaptivity – The structure of the basic perceptual control unit is inherently adaptive. A disturbance to the perceptual input results in error, which in turn results in output that acts upon the perception, automatically cancelling out the effects of the disturbance. This functionality enables a system to control and maintain its goal despite wide-ranging disturbances in dynamic, chaotic and unpredictable environments.

Autonomy – As the goals of the system and the means by which they can be achieved are themselves embodied within a perceptual control system, it can be said to be truly autonomous.

Model-free – As output is not predicted, world models are not required. Therefore, associated problems of model acquisition and complexity are non-issues. The case for predictive world models is based upon the invalid premise that output is a function of input.

Artificial life framework – An approach to building artificial intelligent systems based upon the principles of perceptual control is an ideal candidate for a framework for artificial life. It addresses some of the key ALife challenges of defining a simple, common process and a scalable, dynamical hierarchy, and of explaining how complex and intelligent behaviour emerges from the simplicity.

ALife research has led to some interesting insights into self-organization and how simplicity gives rise to complex behaviour; however, it is recognized that the current ALife methodologies are not sufficient for fully explaining life in terms of general behaviour.

The process of perceptual control holds the promise of furnishing an organizing principle of behaviour that differentiates living from non-living systems, providing the missing 'stuff' of artificial life systems.

The biological, evolutionary path represents the only known example that has produced intelligent beings. It is highly questionable whether other approaches are likely to be fruitful, given the context in which that intelligence has emerged. For a few decades it has been recognized that an approach to developing artificial intelligent systems will require situatedness, embodiment and dynamism. What has not been recognized, until now, is that it will also require perceptual control.

1. Young, R., *Artif. Life*, 2017, **23**, 2; http://goo.gl/n8QAUx
2. Young, R. Dynamic visual robot arm control, 2016; https://youtu.be/jmwH0AZtGG4
3. Young, R., Effing control, 2017; https://youtu.be/fSH6OKS0dTM

Rupert Young
e-mail: rupert@perceptualrobots.com