Workshop report

Cognitive systems: Applications, requirements and capabilities with a focus on Aerospace and Automotive

Ecole des Mines, Paris
20 March 2008

Patrick Courtney, Bill Sharpe and Erik Hollnagel

Executive summary

A significant research programme in cognitive systems is now underway. This focuses on the developing and technology and the necessary scientific understand to provide significant levels of autonomy and decision making into computer-based systems. Active research approaches in the area range broadly, from traditional rule-based AI, through to connectionist, dynamical and emergent systems and include embodied systems combining computing and robotic systems.

One major practical motivation for the development of cognitive systems is to overcome the problems faced by traditional computer systems in dealing robustly with the uncertainties and changing demands that characterise the real world. Potential applications cited span a very broad range and have included care-giver robots, and easier-to-use interfaces. In order to make link with the developing discipline of cognitive systems, a new level of multidisciplinary dialogue on the centre ground is needed to build the concepts and community of Cognitive Systems. The objective on this activity is to identify stepping stones between applications and research into cognitive systems.

The domains of aerospace and automotive are examined with a view to identifying how the issues of autonomy and decision making are addressed, and trends which call for increased autonomy. Two specific areas are selected for further discussion: context-aware detection; and the coordination of multiple cognitive agents using contracts. Other possible areas are briefly indicated and potential next steps are suggested.

1. Introduction

A significant research programme in cognitive systems is now underway. This focuses on the developing and technology and the necessary scientific understand to provide significant levels of autonomy and decision making into computer-based systems. Active research approaches in the area range broadly, from traditional rule-based AI, through to connectionist, dynamical and emergent systems and include embodied systems combining computing and robotic systems.

One major practical motivation for the development of cognitive systems is to overcome the problems faced by traditional computer systems in dealing robustly with the uncertainties and changing demands that characterise the real world. Potential applications cited span a very broad range and have included care-giver robots, and easier-to-use interfaces. In order to make link with the developing discipline of cognitive systems, a new level of multidisciplinary dialogue on the centre ground is needed to build the concepts and community of Cognitive Systems. The objective on this activity is to identify stepping stones between applications and research into cognitive systems.

1.1 Background and motivation

The automotive and aerospace sectors are selected as the focus of the discussion because they already have long experience in the development of mechanical platforms upon which the current embodied paradigm relies. They have some history in enhancing these platforms to improve some of the same characteristics of usability and robustness that is the focus for cognitive system community. These sectors are already commercially well developed, not least in comparison to care-giving or domestic automation.

The target audience of this report is primarily researchers in cognitive systems interested in making a link to applications. It can be seen as part of the 'learning journeys' activity that the eucognition network has been driving as a precursor to a roadmap for the domain which is currently still missing.

1.2 Methodology

A number of experts in the area of automotive and aerospace systems with a potential interest in cognitive systems were identified and invited to a workshop organised and funded by the eucognition network[1]. This report is based on the discussions that took place at the workshop.

1.3 Structure of report

The report starts by examining domains of aerospace and automotive with a view to identifying how the issues of autonomy and decision making are addressed, and trends which call for increased autonomy. Two specific areas are selected for further discussion: context-aware detection; and the coordination of multiple cognitive agents using contracts. Other possible areas are briefly indicated and potential next steps are suggested.

[1] www. eucognition.org

2. Position of cognitive systems technology with the aerospace & automotive domains

2.1 Aircraft as a cognitive system

Within the modern commercial aerospace environment one can discern a number of distinct participants – the main ones being the **pilot** working with the aircrew to operate the individual aircraft; and the **air traffic controller** operating within the air traffic control system, the main task of which is to ensure separation of aircraft and provide additional information to the pilot.

The final responsibility for the safety of the flight rests with the pilot. Over time the pilot has been provided with a number of aids which assist in this responsibility. Watt and Wiener control loops were introduced in the 1930s, while Kalman filters appeared in the 1960s. By the 1980s the control of multiple loops had became an important issue and whilst further supporting the pilot, they also added greatly to the complexity of the task.

If one considers the activities involved in the operating of an aircraft, one can identify a number of distinct actions. Trajectory control involves control of the aircraft attitude and position around the centre of gravity. The aircraft also operates within one of a number of **flight modes** for example climb, altitude hold etc. Finally flight management covers issues of planning based around a map representation, navigation and route modification, take off and landing approaches with significant input from air traffic control.

Each activity operates close to a particular **level of control** with a particular time constant. Trajectory control operates over the short term sub-second timescale. Tactical decisions over the control of operating-mode functions over a somewhat longer (sub-minute) timescale, while strategic route management takes places over a timescale of minutes[2].

As more subsystems have been introduced into the cockpit, the task of the pilot has increasingly moved to one of the **management** of these subsystems. This management task consists of arbitrating and resolving conflicts between subsystems, decision-making, and recovering from problematic and failure situations. The ability to do this is based on the pilot possessing the necessary training and experience. **Additional subsystems** have been incorporated with somewhat mixed results, for example ground proximity warning systems [Boy2006][3]. This has become stabilised as the appropriate ways of presenting the information and integration into the management task have become established.

One of the main external variables is the **weather**. Information on the current situation is captured from a number of sources such as sensors and communications from the ground. This is presented to the pilot in a number of ways involving both continuous measures

[2] This division into the three levels of reactive, tactical and strategic will be familiar to developers of automation technology and appears in some form in a great many cognitive architectures descriptions

[3] It might be valuable to analyze such 'mixed results' and what 'lessons learned' that can be drawn from these experiences. Such analysis may already exist in the literature.

(eg wind speed) and symbolic data. The pilot is then required to anticipate the forthcoming weather conditions based on experience and thus decide which action to take, including, if necessary modification of the trajectory and route.

Statistics on accidents are rigorously managed thanks to the existing legal frameworks, as well as the procedures (near-miss reporting[4]) and technology put in place (eg black box) Deficiencies in maintenance have now become a significant (root) cause of failures and cognitive systems may have a role in a decision support function.

Trends: The current economics of air transport is such that the management of the saturation of airspace which, under existing safety constraints on separation time (90 seconds) and space has become the major driver. The current growth in traffic of approx 5% pa is making this a considerable challenge. One approach has been to try to decentralise the decision-making and implementation by giving more autonomy to aircraft using the principle of **contracts**. These specify the terms under which an aircraft intends to operate for a particular leg of the journey and is negotiated with air traffic control. This allows delegation of authority whilst managing deviances and penalties agreed in advance.

There is also a move towards **integrating air-based and ground-based systems** for greater efficiency, although traditionally the suppliers and operators have been different players with distinct histories (priorities, objectives and regulatory contexts).

2.2 Automobile as a cognitive system

In comparison to aerospace system, road transport exhibits a number of significant differences[5]. Superficially, basic assumptions about the two types of system appear to differ greatly. For example, an aircraft is a dynamic system that cannot be expected to simply stop in midair whereas a car can do this. However, such an objection is not such a problem when one considered that a car cannot safely stop on a motorway, and a taxiing aircraft can come to a halt. This means that the specific situation or state determines the modes of operation that are available to the vehicle.

It is interesting to note that since the development of automobiles took place over roughly the same timescale as the development of aircraft, one could argue that the available technologies were very similar. However the political and economic environment ensured considerable divergence in approaches. Perhaps more significant is the stage of development in the two sectors and differences in perception and vocabulary rather than

[4] It may be interesting to examine what can be learned, and to transfer some of the near miss analysis and gained know-how into other areas such as automotive. However it is understood that the recording system for near misses is not applicable to the automotive world, but there may be ways to achieve sufficient incident recording in other ways.
[5] It may be interesting to see what can be learnt from the experiences of the aerospace systems, and where the differences mean that one should refrain from drawing parallels and why and in what way this insight could help implement and measure implementation success in automotive applications.

any deeper divergence in the objectives to enhance safety comfort and performance. However the differences are worth examining to inform the discussion.

One notable difference is that the **diversity of vehicle types** involved in the road environment is far greater, from large trucks to bicycles via occasional users such as tractors and earth moving equipment. The age and condition of the vehicle also varies widely but the impact is the greatly moderated by the use of vehicle certification in most countries. Mandating of newer equipment to vehicles is less prevalent than for aircraft.

The diversity of **drivers** is also far wider than for aircraft pilots, in terms of training and experience, as well as cognitive ability. The initial training of the human driver is at a much lower level, although this has been improving. Moreover the accumulation of experience over time can be considerable but this is not certified beyond simple metrics such as number of year without insurance claim. Whilst many rules of behaviour and interaction are coded in law, road drivers as a group appear to have greater autonomy to establish informal rules than in aerospace. However both the safety and the performance of a journey remains primarily a function of driver skill, much like in the early days of aviation.

The **local road environment** is arguably more complex with many nearby vehicles and complex signage absent in airborne system. **Control** has both decentralised and centralised aspects, with a dynamic aspect provided by traffic lights. The radio delivery of traffic news provides a form of dynamic communication but is centralised in the sense that it is not targeted to specific vehicles. Signage is typically static although dynamic (adaptive) signage is appearing in some areas. Some types of environment are much more controlled than others, for example motorways compared to the city driving environment. This makes for less demand of the anticipation to a smaller number of possible states.

As with aerospace, **weather** plays a role. Indeed a wider range of variability is accepted in the different geographical regions since more discretion rests with the driver.

The three **levels of control** – reactive, tactical and strategic – apply just as well to automobile as for aircraft, albeit with some slight modification due to the nature of the vehicle. The reactive vehicle-attitude-level includes longitudinal control (velocity and acceleration and a degree of lateral guidance) whereas the tactical-level includes lateral control with modes such as lane changing.

The last few years has seen the introduction of a wide range of **new technological subsystems:** electronic power steering EPS, electronic stability program ESP, anti-lock braking ABS, parking assistance, lane keeping and so on. These subsystems may provide one of a number of possible functions in support of the driver: **information** (description); **warning** (prescription) or **substitution** (partial or global). Also expected of the subsystem is ease of management of the various support systems (compatibility of function) and adaptability to the context and selectivity of the assistance provided. This brings us back to the task of the pilot as arbitrator between the subsystems when the information they provide is in conflict.

A number of these subsystems clearly operate at the lowest **reactive-**level of operation. As such they often operate under a very tight time constraint. The performance of the driver under these conditions is often quite modest which means that automated systems offer considerable benefit operating as a substitute. Other subsystems such a GPS clearly operate at the longer-term route-planning **strategic**-level providing descriptive information. This is somewhat less problematic since the time constraint is considerably more relaxed, although problems may arise if the information provided starts to affect decisions made at the tactical level.

One particularly powerful subsystem is of a driver assistance system (**ADAS**) which uses an additional sensor such as an infrared camera which provides information complementary to that already available to the driver. The infrared signal is able to detect pedestrians or other potential obstacles that the driver may not be able to see in time. A typical system might be configured to warn the driver of pedestrians by the side of the road to avoid an accidental collision. The possibility also exists that the system autonomously take action by e.g., braking the vehicle if the driver fail to take action and a near-accident situation is developing[6].

Considerable work has been done to better understand the issues involved in developing and evaluating driver support subsystems and a number of issues appear to be involved as shown below:

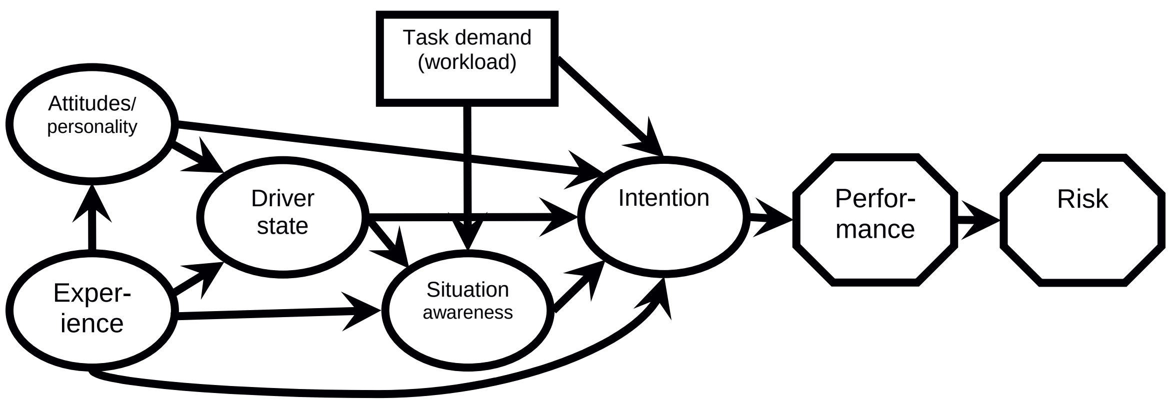

Several studies have shown that the driver may reject the alert, even if they have previously agreed that the alert would be useful. Both the **principle** of alerting the driver and the **means of delivering the alert** must be carefully adapted. Experiments with active speed limiters has shown that even if the driver has agreed to the alert, they may reject an intervention that they feel is intrusive such as limiting movement of the accelerator pedal. On the other hand, intervention by an automatic braking ABS to the breaking servo is widely accepted. In the case of GPS, personalisation options are now widespread. Experience with lane detection systems developed to enhance safety has

[6] The definition of a near-accident situation is not universally agreed and attempts that have been developed such as time to collision, (TTC) have several limitations and draw-backs.

shown that they have been used in ways not originally imagined [reference missing], for example to allow drivers to continue driving after they have become drowsy[7].

Given that the technology generating an alert will necessarily generate some **false alarms** and it is well known that multiple false alarms reduces the confidence of the driver, this may well lead the driver to switch off the subsystem, or worse, for it to become distracting. The underlying probability of the event to be detected is typically already very low [Parasuraman 1997] while the probability of detection should be reasonably high (well above 50% say) which places very high demands on the acceptable false alarm rate. Current systems are not able to integrate sufficient **contextual information** to filter out false alarms, and cannot tell what the user has already noticed.

Trends: The high social and economic cost of accidents in the road traffic system is a strong motivation to try to reduce their incidence. Saturation of the underlying road infrastructure has become an economic driver.

Whilst the overall autonomy of the automobile is very low today, as the number of additional support subsystems has multiplied in recent years, one can imagine that further subsystems will be proposed in the future. **Integration** of these multiple subsystems and multiple levels of control will take on an ever greater importance. As with the aircraft, as autonomy rises, the task of the driver will change into one of subsystem management and cooperation with the intelligent co-driver where issues of predictability and dependability take on greater importance, anticipating situations as they develop and problems before they occur.

[7] This is a continuation of longstanding debate on passive safety features encouraging dangerous driving. Certain systems do appear to exhibit a negative effect though one should be careful and search for examples of positive effects. ESP has recently been shown to have significant benefits, only surpassed by seat-belts in cars. In the end it may be an issue of how to balance conflicting demands on a system. Such balancing should consider cognitive factors and may be a justification to further research in the area.

3. Potential research themes for cognitive systems

In this section an attempt is made to try to define themes where the interests of cognitive systems researchers and the application domains coincide.

3.1 Context-aware detection

Perception technologies based on infrared and visible light sensors offer the potential to perform useful tasks, especially at the middle tactical level. Conversion of the image signal into a usable decision signal remains problematic however. This in part due to the utilisation of the decision or alert within the whole driver support system, as it may involve more than a simple servo or guidance and may involve mode switching (turn left/right, lane change, begin overtaking manoeuvre etc). These take place under moderate time constraint, with the problem of mode confusion, that is, which modes to switch to in a critical situation which in turn depends on the situation. Thus to allow the driver to reliability integrate the decision signal requires consideration of the context in which it will operate. A framework to carry this out would require consideration of the following factors:

The **context of operation** affects the ability to perform a reliable detection. For example the detection task at night-time will be quite different to daytime operation. This will affect the acceptable true rejection/false acceptance detection performance since the baseline performance of the operator will also vary. Similarly operation in an urban environment will be different from that in a country road. Likewise weather will have a strong effect. Some aspects of the context will be detectable from other sensors such as location and time. In other instances a particular external event may lead to particular data patterns. For example after a concert a lot of people may spill onto the street which would lead to too many alarms in a pedestrian detection system. If such as event could be identified or inferred then this could be used to adjust the detection system. Useful within this is the definition of **scenarios** and use cases covering both normal and abnormal operation.

Capturing the current **operating rules** is also important. This may originate from a rule book but also from observed practice, especially in the case where variations in local driving style and preference are observed. For instance rules at a pedestrian crossing differ across Europe from the pedestrian having priority in Denmark, to the vehicle having priority in Portugal. Certain cues may be available and could be detected, including traffic density, vehicle separation, use of turn-indicators in other vehicles, etc. Acquisition of this knowledge can be thought of as part of the cognitive problem.

The coding of **operator skill** levels and the tracking of their evolution thanks to learning. This should be representative of the drivers own competence. In the case of pilots, regular re-training schedules are established and documented. For automobiles, this would have to be inferred, though the system could be in a position to generate evidence that a particular condition has been experienced[8]. Simple tests to evaluate the extent to which skills have become procedural are available and could play a role.

Intention estimation is important to help ensure that the alert is usable. Current systems cannot tell if the driver has already noticed a situation. Alternative sensors (brain activity, eye gaze, response time…) may be required to assess **driver state,** making it possible to model how the human adapts to the situation.

If the operator is able to provide a **feedback** to the detection subsystem, this may be used to set driver preferences (suppress or override) and threshold levels for various classes of detectable event and possibly even make use of near-miss information.

As learning by the system will play an important role in improving the performance, issues of **generalisation** and over-generalisation start to play a role. Learning may take place over a particular mission or trip, but generalizing to a new mission may cause a problem if the class of mission is not appropriately recognised, for example a business as against a leisure trip.

Interaction instabilities. When two adaptive systems operate, for example the human operator and the support system, there is the potential to create interaction instabilities which must be managed. The short term control stability a known problem in control theory, but presents a novel aspect when involved longer term instability due to learning. In this case modelling of the operator becomes an important task.

Once a large number of systems are deployed, a further possibility becomes possible, that of distributed learning and **information sharing** between units.

Increasing number of on-board systems calls for greater **integration**. This involves managing the cognitive load on the driver, as well as appropriately modelling the level at which the alarm and decision intervene. As with aerospace, the pilot/driver is expected to retain final responsibility.

The theme of context-aware detection can be expected to be general interest as it could be applicable to a wide range of detection tasks, whether of pedestrians, ground proximity, or of faults on manufactured parts.

3.2 Coordination of multiple cognitive agent using contracts

The presence of multiple autonomous systems within a dynamic environment presents obvious risks in terms of conflict with a strong potential for a loss of safety and task performance. As with the aerospace and automotive domains, this will be experienced by autonomous cognitive systems in other domains.

---

[8] Similar issues regarding practice of fine motor skills can be heard from surgeons & professional musicians

The approach of formalising the terms of the interaction in the form of a **contract** seems attractive for a number of reasons. They allow one to specify the terms under which an autonomous systems intends to implement a particular element of its strategy when this is likely to make use of a limited resource which must be negotiated with other users of this resource. In the case of aerospace this would be negotiated with air traffic control as a centralised authority. This allows delegation of authority whilst managing deviances and penalties agreed in advance.

Given these parameters it should then be possible to reason over sets of contracts especially if contracts are defined in a symbolic form. If this form is also human readable then this would lend them open to human overview, though this does not seem to be obligatory. One challenge will be how to define and enforce penalties and generally police the contracts. Perhaps it will fall to an independent agent to avoid conflicts of interest. If this were a human agent, then the use of human readable contracts is obviously desirable.

Related problems in a number of other fields may be very relevant here. These include game theory from **economics** which allows the taking into account of the value and costs associated with negotiating parameters, including mechanisms such as markets, the use of auctions and other adaptive behaviours.

Work in the area of **communications** may also have a bearing as this also concerns the development of protocols to optimise the use of a scarce resource (bandwidth) whilst ensuring the avoidance of instabilities and lock up is a prime concern.

The issue of establishing frames of reference in dialogues may find a parallel with similar issues in **linguistics**. Further discussions with participants from these domains may be beneficial in developing this theme further.

3.3 Further potential applications

3.3.1 Further potential applications in rehabilitation research

In the field of rehabilitation, assistive technologies could enable those who depend on them to live a more independent and easier life. However, as currently available data demonstrate, the technology on the market does not meet this goal. For example in the field of electrically powered wheelchairs, studies show that it takes months and sometimes even years to learn how to control the device so that it actually makes achieving everyday tasks easier. Some people in need even lack the required visual acuity or motor skills and never manage to learn how to control electrically powered wheelchairs. The statistics are even worse when only considering people who cannot operate a traditional joystick and require specialty controls such sip-puff devices to steer their wheelchair. One reason for this long skill acquisition process is the high number of input commands, which are necessary, to achieve the, from the assistive technology, desired behavior. Hence, a cognitive rehabilitation system could offer sophisticated behaviors to its user and apply methods which were e.g., developed in the field of robotics to ease the everyday life of the people in need:

On one hand, the behaviors also developed and applied in the aerospace and automotive sector could be generalized to the rehabilitation sector. A typical example is a collision warning or avoidance system. This is for wheelchair users especially important, as they can hardly turn around, check the environment behind them and might not notice stairs going down right behind them. A system with appropriate sensors and data processing methods warning them or stopping the wheelchair if the users want to drive to the back might make life for these people in need safer. Another example is an appropriate **intention estimation** behavior which, on the basis of an embedded sensory system and the past behavior of the user, predicts the user's future behavioral goal and drives him/her autonomously to the desired goal position. This has the potential to significantly reduce the number of input commands required from the user and, thus, make the skill acquisition process shorter.

On the other hand, the specific requirements of the rehabilitation sector make the development of specific methods and algorithms necessary. There is a great inter- and intra-individual variability within and between different types and degrees of disabilities. This is why a highly cognitive assistive system should have appropriate means to judge on the **current status** of its user, on his/her currently available cognitive and motor skills and give the support the user needs in this moment to make his/her life easier. This is especially crucial in order not to take over too many tasks, so that the person in need does not unlearn skills, which he/she still has. The support given should be exactly the one that the user needs in the current situation with his/her current level of abilities.

Besides this development of new and sophisticated behaviors, it is another important step to **integrate** the variety of available assistive technologies on the market in a dependable manner without the danger of interaction instabilities. If, for example, the wheelchair is enabled to drive its user from the kitchen to the living room, home automation should be connected to the wheelchair system to open/close the doors automatically at the same time, switch on/off the lights when necessary, etc. Such a comprehensive and highly cognitive assistive system would not only have the potential to revolutionize the field of rehabilitation technology, but also give important insights into the impact of cognitive systems and open new research topics such as the adaptation of technologies to the inter- and intra-individual differences of users.

3.3.2 Further potential applications

Other potential application areas in robotics and business systems were briefly discussed at the workshop and may merit further discussion at another occasion.

4. Conclusions: outcomes and recommendations

A number of potential research subjects have been identified and justified. The advanced state of development of the aerospace field indicates several opportunities within automotive and rehabilitation research, and possibly other areas, whilst bearing in mind the different conditions such as operators skill level etc. A number of topics worthy of further investigation are suggested. Some of these may usefully involve wider disciplines including economics, communications and linguistics. Further discussions with participants from these domains may be beneficial in developing this theme further.

The requirements and desired capabilities are closer to being understood. We expect to see a degree of technological convergence and potentially a parallel evolution of the role of the driver as manager of subsystems. However more discussion is required. The workshop report is only one step and should be seen as part of an ongoing discussion.

5. References

[Boy2006] Boy, G., Barbé J. and Giuliano S., 2006, Automation and assistance in aeronautics and automotive: Diversity versus homogenization? AAAF Conference Proceedings - Towards common Engineering & Technologies for land, maritime, air and space Transportation, Paris, November 20-23, 2006.

[Parasuraman1997] Parasuraman, R., Hancock, P.A., & Olufinboba, O. 1997, Alarm effectiveness in driver-centred collision-warning systems. *Ergonomics 40*, 3, 390 - 399.

6. Acknowledgements

The authors extend their thanks to the workshop participants (Guy Boy, Meike Jipp, Jan-Erik Källhammer, Christof Eberst, Martin Spott, Oliver Carsten**,** Farida Saad), the Ecole des Mines, Paris for the use of facilities, and David Vernon and eucognition for supporting the workshop.