AISB Quarterly

No.131

The Newsletter of the Society for the Study of Artificial Intelligence and Simulation of Behaviour

Integration of neural networks with production rules

An important insight into the simulation of behavior is that certain behaviors tend to be more dependent on lower levels of cognition than others. For example, to correctly simulate pop-out effects in visual search, it is important to understand how neurons are organized in the visual pathway.[1] To understand how people perform reasoning and decision-making tasks, it is important to provide formal models of the reasoning strategies that are being deployed.[2]

Obviously, most behavior that people exert comprises combinations of low-level psychophysical phenomena and high-level strategies. For example, performing a visual search task does not only require identification of the target stimulus, but also the perception of the to-be-searched information, storage in memory of that information, matching that information to the target stimulus, and motor actions that are necessary to provide a response – for instance, a button press.

A recent computational model of full-task behavior has been developed that takes the integration of multiple levels of behavior to a new level. Retrieval by Accumulating Evidence in an Architecture (RACE/A) [3,4] integrates memory models [5] that originate in the neural network literature with the cognitive architecture ACT-R.[6] ACT-R is essentially a production rule system that excels at explaining and simulating complete tasks at a rather high level of abstraction. By contrast, the theory on memory models that was integrated in the cognitive architecture is a detailed theory that is capable of predicting firing rates at the level of individual neurons.[7] The new integrated model will be able to simulate what certain neural behaviors mean on the level of task execution, and vice versa what taking a different task strategy means on the low level of neurons.

At the recent International Conference on Cognitive Modeling, held in Manchester, UK, we presented an example of a RACE/A model. The model described how Stroop interference effects (Figure 1) originate from the interaction of multiple sequential memory processes. While it was already known that the size of the interference depended on the particular order in which the trials were presented, this work also demonstrated that the amount of interference in each sequential memory process depends on the sequence of trials.

In the experiment that was presented (and the RACE/A simulation), participants were requested to perform two tasks at the same time. In one task, participants were to indicate from a tone whether it was high- or low-pitched. In the other task, participants were asked to name pictures (picture-word interference task, see Figure 1 for an example picture). An important instruction was that although the stimuli (tone and picture) appeared at almost the same time, the participants were required to first respond to the tone, and then to the picture. The instruction

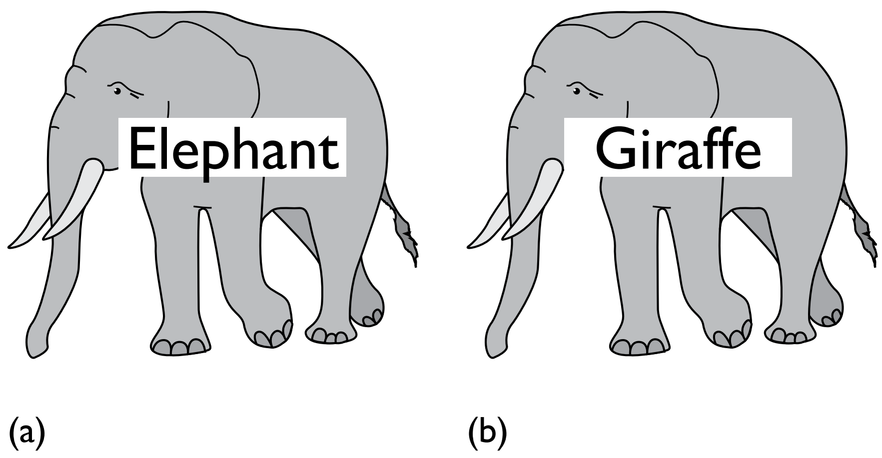

Figure 1. Example stimuli used in the picture-word interference experiments. The typical Stroop interference occurs when participants are asked to indicate in which color a word is printed, while the word also spells a color word (e.g., the word "red" in green ink). The interference is present as a slower response and more mistakes. The picture-word stimuli used in these experiments are thought to represent the same phenomenon.

In this issue:

Integration (cont.)

The AISB Quarterly is published by the Society for the Study of Artificial Intelligence and Simulation of Behaviour.

ISSN 0268-4179

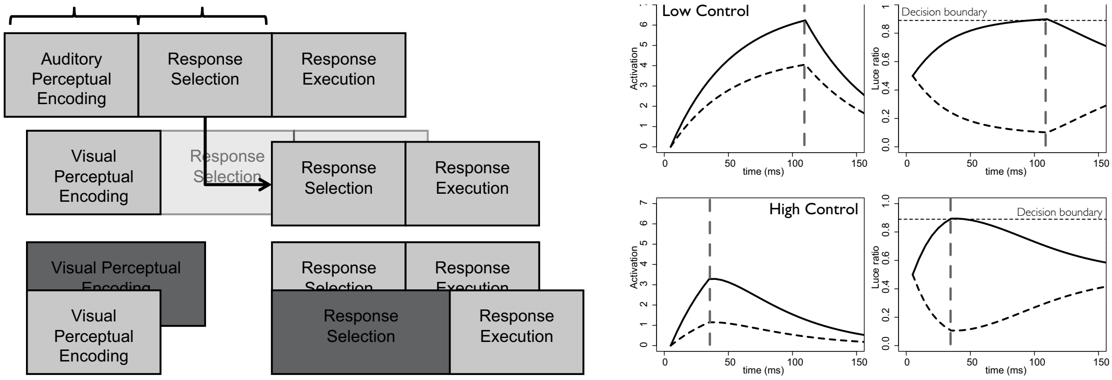

(a) (b)

Figure 2. a) Schematic representation of how inclusion of one task can create a delay in processing the other. Also indicated are the different stages in which interference effects may occur. (b) Schematic representation of how a fast initial retrieval may lead to a slow secondary retrieval.

on response order creates a delay in the processing of the picture-word interference task if the stimuli are presented really close together (Figure 2a). This delay is interesting, because it enables us to study in which stage of a task a particular interference effect is manifested.

As it turns out, interference effects are manifested in different stages, depending on the previous trial. The RACE/A model provides an explanation of why this is the case. The trials following a hard trial (with a stimulus similar to Figure 1b) are characterized by more focussed attention towards the picture (High Control).[8] The trials following an easy trial (Figure 1a) are characterized by less focussed attention (Low Control). Because less attention is directed at the picture (and more at the word), Low Control trials initially involve a tough memory retrieval. However, a second memory retrieval may be necessary to complete the task. For instance, participants (and the model) might retrieve information that relates to the motor program necessary to pronounce the name of the picture. The difficult initial retrieval causes the model to return to an equilibrium state relatively slowly, which gives the second memory retrieval a head start, relative to the High Control trials. In these trials, the initial fast retrieval results in a faster return to equilibrium (for more detail on the experiment or the model, see ref 3).

This example demonstrates that not only are the particular dynamics of memory of importance, but also the manner in which these are incorporated in the execution of a task, even in the execution of multiple (similar) tasks in succession. The RACE/A integrated theory of memory retrieval helps to understand these kinds of phenomena.

References

[1] Itti, L. & Koch, C. (2001) Computational modelling of visual attention. Nat. Rev. Neuroscience 2 (3), 194-203.

[2] Van Maanen, L. & Marewski, J.N. (2009) Recommender Systems for Literature Selection: A Competition of Decision Making and Memory Models. Proceedings of CogSci2009 (Amsterdam, NED).

[3] Van Maanen, L. & Van Rijn, H. (2009)The Locus of the Gratton Effect in Picture-Word Interference in, Howes, A., Peebles, D. & R.Cooper (eds) Proceedings of ICCM2009, Manchester, UK.

[4] Van Maanen, L.(2009) Context Effects on Memory Retrieval: Theory and Applications Doctoral thesis. University of Groningen: Groningen.

[5] Usher, M. & McClelland, J.L. (2001) The time course of perceptual choice: the leaky, competing accumulator model. Psychol. Rev. 108 (3), 550-592.

[6] Anderson, J.R. (2007) How Can the Human Mind Occur in the Physical Universe? New York: Oxford UP.

[7] Smith, P.L. & Ratcliff, R. (2004) Psychology and neurobiology of simple decisions. Trends Neurosci 27 (3), 161-168.

[8] Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S., & Cohen, J.D. (2001) Conflict monitoring and cognitive control. Psychol. Rev. 108 (3), 624-652.

Dr Leendert Van Maanen
Postdoctoral researcher
University of Amsterdam

Artificial Intelligence and the Frontiers of Genetics Research

The field of biology is in the midst of a well-documented data explosion. Advances in testing and measurement technologies, along with increases in computing power and storage now allow biologists to record not only static DNA profiles of organisms (the genome), but the massively complex sets of gene regulations that occur in cells known as the interactome.

The use of technology to sequence genomes from large populations of organisms allows the discovery of subtle differences in DNA that can provide valuable insight into the propensity for diseases such as type II diabetes [1] and common traits such as height [2] and the neuronal influence on obesity [3]. There is great potential for these genome-wide association studies to give us a better understanding of some of the most pressing human diseases and provide some genetic explanation for the infinite variation we see in human beings. As a result there is great interest in these genome-wide association studies (GWAS) and their ability to identify risk factors for disease and traits in the research literature and the media [4].

Gene expression data differs from sequence data in that it measures the level of expression of genes, effectively providing us with a snapshot of gene activity at a particular timepoint. Sequencing technology provides us with a static view of the genome and gene expression data with a dynamic one. Perhaps the best analogy is that of a complicated computer program. The set of instructions for the program (e.g. the code or DNA) can be known in advance, but it is only once the program is run and those instructions interact that the true nature of the program known and the debugging really begins. The levels of gene expression and their variation over time provides us with a host of information on the likely function of genes and their associations with cellular processes and, therefore potentially, disease. The discovery of networks of interactions between genes, known as gene regulatory networks, is a highly active research area and efforts to characterise the interactions between genes and other biological systems are currently being undertaken, but still have a long way to go.

Artificial intelligence techniques are making a real impact in this area where their abilities to mine large volumes of data in a reasonable timescale and provide hypotheses for further targeted biological research in the lab are invaluable. The techniques taken from machine learning and optimisation algorithms extend the capability of standard statistical testing and are able to mine the much larger spaces of the interactome. The following sections discuss the application of artificial intelligence techniques in two key areas: genome-wide association studies and gene expression data.

The first genome-wide association studies were published around four years ago, making them a very recent field of research, although the number of studies has exploded since then. Major studies which record data for thousands of individuals such as the Wellcome Trust Case Control Consortium and 1000 genomes projects now allow researchers easy access to large volumes of GWAS data. These studies investigate the effect of mutations in DNA across a population of organisms known as single nucleotide polymorphisms (SNPs). These small changes can occur within a gene, in non-coding regions of a gene or in intergenic regions. Typically, there are 400,000 to 1,000,000 of these for humans and so it is possible to exhaustively evaluate each single SNP and its association with a disease or trait using modern machines, providing a ranked list of SNPs and the statistical significance of their association with the disease or trait. However, to date, these single-SNP studies have only been able to account for approximately 5% of the heritability of such traits and diseases, where it is known that much more is inherited. One potential source of this missing heritability is in gene-gene interactions, where pairs, triplets or higher numbers of genes interact to create a phenotype. Inevitably, this creates a combinatorial problem that increases the search space of possible interactions to intractable proportions.

It is at this point that intelligent search and optimisation algorithms are required to search these much larger spaces. A number of approaches have been trialled including random forests [5] and genetic algorithms [6] (see [7] for a more comprehensive review). However, one of the most popular approaches is the use of ant colony optimisation. This search technique, based on the foraging behaviour of insects, has been shown to discover known single associations [8] and potential gene-gene associations [8, 9] from real-world datasets. The ants search by finding a path through the set of possible SNPs and locating those combinations that are most associated with the trait or disease (e.g. type II diabetes). Each successive population of ants uses the pheromone trail of those that have gone before it, to make decisions about path choices and the algorithm converges on a set of important SNPs. The search for SNP combinations is atypical of many search problems in that the algorithm must select a small number of SNPs (typically less than 5) from a database consisting of tens or hundreds of thousands. This problem difficulty, and the relative recent availability of GWAS data, means that this field is still in its infancy although it would appear to have a bright future.

The analysis of gene expression data using intelligent search and optimisation techniques is a more mature field, dating back to around the year 2000. This data can also be used to diagnose cancers and other genetic diseases through the use of optimisation techniques. The data created from microarrays will give expression levels of genes for an organism at a particular timepoint. When this is repeated for a population of diseased individuals and controls, a static dataset is created and the task is to interrogate that set and determine those genes that differentiate diseased individuals and controls. As with the SNP data, multiple genes in combination will be required to classify the individuals correctly (although higher order combinations are more commonplace here), rendering exhaustive search intractable. As a result of this complexity, many optimisation algorithms have been applied to this task and some of the most popular include support vector machines [10], neural networks [11], evolutionary algorithms [12,13] and hybrids of these to discover sets of genes related to cancers. More recently, multi-objective approaches [14, 15] have been used to simultaneously minimise the number of genes discovered and maximise accuracy of the discovered solutions. Attempts such as these to reduce the complexity of the discovered models whilst

AI and genetics (cont.)

maintaining accuracy are important to deliver results to expert biologists that can be easily investigated in the lab.

An additional use of gene expression data is to sample the same organism through time and therefore create a temporal record of gene expression. This data can then be used to determine genes of similar function, which will vary in similar patterns over time and can be clustered together using self organising maps [16] and other clustering techniques. If the data is sufficiently detailed, then putative gene regulatory networks can be discovered that describe the possible regulatory connections between genes. Models for these regulatory networks include Boolean and differential equation models and discovery can take place using the search and optimisation techniques previously described [17].

With the sequencing of the human genome and the advent of these impressive high throughput technologies, this is a very exciting time for bioinformatics that opens up many opportunities for artificial intelligence techniques. As shown in the examples here, often the complexity involved with searching these large databases is such that intelligent methods are required. The introduction of expert biological knowledge can help to reduce that complexity, but of course this can constrain innovation from the algorithms and reduce the likelihood of entirely novel discoveries. The union of AI and bioinformatics provides biologists with a greater number of tools to better understand disease and provides strong motivation to AI researchers to improve their techniques and make significant discoveries.

When the human genome was first sequenced, it was said that it was the first step to reading the book of life; all the letters in the book were known and the next step was to separate them into words. Some years after this, we are now some way along the road of separating the letters into words and understanding their meaning, but the challenge in the near future will be to determine how each word interacts with others in the book to create the whole story. I hope and believe that AI techniques will have a significant role to play in this incredible human endeavour.

References

[1] T. M. Frayling, T. M. et al (2007), "A Common Variant in the FTO Gene Is Associated with Body Mass Index and Predisposes to Childhood and Adult Obesity", Science, 316: 889-894.
[2] Weedon MN, et al., (2008) "Genome-wide association analysis identifies 20 loci that influence adult height", Nature Genetics, May; 40(5): 575-83. Epub 2008 Apr 6.
[3] Willer CJ et. al., (2009) "Six new loci associated with body mass index highlight a neuronal influence on body weight regulation", Nature Genetics, 41(1): 25-34. Epub 2008 Dec 14.
[4] Henderson M., (2007) "Fat Gene Found by Scientists" The Times, http://www.timesonline.co.uk/tol/news/uk/health/article1647517.ece [accessed 26/3/2010].
[5] Jiang, A., Wanwan, T., Wu, X., Fu, W. (2009) "A random forest approach to the detection of epistatic interactions in case-control studies", BMC Bioinformatics, 10: S65
[6] Hubley, R.M., Zitzler, E., Roach, J.C., (2003) "Evolutionary Algorithms for the selection of single nucleotide polymorphisms", BMC Bioinformatics, 4 (30).
[7] Kelemen, A., Liang,Y., Vasilakos, A. (2008) "Review of Computational Intelligence for Gene-Gene and Gene-Environment Interactions in Disease Mapping" in Studies in Computational Intelligence (SCI) 85: 1-16.
[8] Christmas J., Keedwell E., Frayling T. Perry J., (2011) "Ant Colony Optimisation to identify Genetic Disease Association for Type 2 Diabetes", accepted to Information Sciences doi:10.1016/j.ins.2010.12.005
[9] Wang, Y., Liu,X.,Robbins, K. Rekaya, R. (2010) "AntEpiSeeker: detecting epistatic interactions for case-control studies using a two-stage ant colony optimization algorithm", BMC Research Notes, 3: 11.
[10] Yeang CH, Ramaswamy S, Tamayo P, Mukherjee S, Rifkin RM, Angelo M, Reich M, Lander E, Mesirov J, Golub T. (2001) "Molecular Classification of Multiple Tumour Types", Bioinformatics, 17 (1): S316-S322.
[11] Khan et. al. (2001) "Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks", Nature Medicine, 7 (6): 673-679.
[12] Ooi, C.H., Tan, P., (2003) "Genetic algorithms applied to multi-class prediction for the analysis of gene expression data", Bioinformatics, 19 (1): 37-44.
[13] Keedwell, E.C. and Narayanan, A. (2003) "Genetic algorithms for gene expression analysis, in Gunther Raidl et al (Eds.) Applications of Evolutionary Computing" LNCS 2611, Proceedings of EvoBIO2003 1st European Workshop on Evolutionary Bioinformatics, pp 76-86.
[14] Keedwell E, Narayanan A,(2010) "Gene Expression Rule Discovery with a Multi-objective Neural-Genetic Hybrid", IEEE International Conference on Bioinformatics & Biomedicine, Hong Kong, 18th - 21st Dec 2010.
[15] Garcia-Nieto, J., Alba, E., Jourdan, L., Talbi, E., (2009) "Sensitivity and specificity based multiobjective approach for feature selection: Application to cancer diagnosis", Information Processing Letters 109: 887-896.
[16] Tamayo et. al. (1999) "Interpreting patterns of gene expression with self-organizing maps: Methods and application to hematopoietic differentiation", PNAS March 16, 96 (6): 2907-2912.
[17] Keedwell, E., Narayanan, A., (2005) "Discovering Gene Regulatory Networks with a Neural-Genetic Hybrid", IEEE/ACM Transactions on Computational Biology and Bioinformatics, July-September, Vol 2 (3): 231-243, IEEE Computer Society

Dr Ed Keedwell
Lecturer in Computer Science
University of Exeter

Simulating Human Multitasking with a Cognitive Architecture

Multitasking -- doing several tasks at the same time -- is one of the most astounding of all human abilities. In some situations, people multitask effortlessly: chatting with colleagues while eating lunch, or helping with a child's homework while washing dishes. In other situations, people have an extremely difficult time performing two tasks: sending a text message while driving, or reading a book while listening to someone else's conversation. What makes multitasking so easy, or so difficult, in these different situations? And can we simulate human multitasking with a computational system, predicting whether multitasking will be easy or difficult for some particular combination of tasks?

Over the past five years, we have been making strides toward this goal by studying and simulating multitasking using a cognitive architecture (see Byrne [1]). A cognitive architecture is simultaneously a computational framework and a working theory of human cognition. The architecture allows for specification of running models that simulate both cognition and the perceptual/motor actions that occur in service of cognition; the models often interact with either simulated or real task environments, and thus their behaviors can be visualized and analyzed in much the same way as human behavior on the same tasks. At the same time, the architecture strives to be a faithful representation of cognition, and the models developed within the architecture are generally compared against empirical data collected from human behavior.

Cognitive architectures provide an ideal context in which to simulate human multitasking. We have developed a theory of multitasking, called threaded cognition, which is incorporated into the ACT-R cognitive architecture.[2] Threaded cognition (Salvucci & Taatgen, 2011) states that all people have a core ability to interleave ("thread") the execution of multiple tasks.[3] At the same time, core limitations in cognitive processing result in interference in certain situations, and for certain combinations of tasks. Its instantiation in the ACT-R architecture means that a user can take models of two different tasks and immediately simulate behavior for the two tasks being performed concurrently.

As a simple example, consider a laboratory dual-task scenario: for one task, a person presses a key in response to a visual stimulus; for another, the person speaks in response to an auditory stimulus. It has been shown for variations on such dual tasks that, even when the modalities of the tasks do not overlap (visual/manual in the first, aural/vocal in the second), people are slower to respond to one of the tasks. Threaded cognition accounts for this interference because of ACT-R's cognitive processor: in essence, a person can only perform a small "thinking step" for one task at a time, producing a cognitive bottleneck that slows down the second task. Interestingly, the dual-task interference can sometimes disappear with enough training and when the two tasks are not prioritized. This result arises in the threaded cognition account from an ACT-R learning mechanism that reduces each task's memory load, which eventually interleaves the two (unprioritized) tasks in a way that neither task interferes with the other.

The theory also extends to applied, complex, task domains. For instance, we have studied the effects of a driver's interaction with in-vehicle devices while driving (see Salvucci & Taatgen, 2011). The interleaving of other tasks, especially visual tasks, not surprisingly can produce large effects in a driver's ability to steer and/or respond to external events. Perhaps more surprisingly, even purely cognitive tasks such as mentally rehearsing a list of numbers can produce observable effects on driver behavior. Again, the combination of ACT-R and threaded cognition helps to account for these phenomena.

Threaded cognition is not restricted to concurrent multitasking, doing multiple tasks at the same time; it can also be used to model sequential multitasking, in which tasks are alternated. More specifically, threaded cognition predicts that alternation between two tasks is much more efficient if the resource demands (especially maintenance of context) of these tasks do not overlap, and that people will try to switch between tasks when the risk of losing context is minimal.

The simulation of multitasking behavior by means of threaded cognition and ACT-R thus has two important implications. First, these theories and their associated computational mechanisms help us to better understand the intricate workings of the mind with respect to people's fascinating multitasking abilities. Second, the computational theories allow us to make predictions through simulation -- e.g., a prediction of the distraction potential of a new in-vehicle device -- thus grounding the theory with the promise of guiding the design and development of applied systems.

References

[1] Byrne, M.D. (2008). Cognitive architecture. In The Human-Computer Interaction Handbook. CRC Press: 93-113.
[2] Anderson, J. R. (2007). How can the human mind occur in the physical universe? New York: Oxford University Press.
[3] Salvucci, D.D., & Taatgen, N.A. (2011). The Multitasking Mind. New York: Oxford University Press.

Dr Dario D. Salvucci, Drexel University
Dr Niels A. Taatgen, University of Groningen

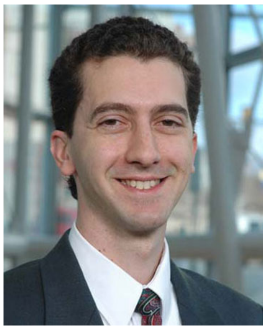

Modeling Dynamics of Multimodal Cues for Spontaneous Agreement and Disagreement Recognition

Social Signal Processing (SSP), an emerging research domain, aims at bringing social intelligence to computers. A social signal is defined as a communicative or informative signal that, either directly or indirectly, provides information about "social facts", namely social interactions (e.g. turn taking), social emotions (e.g. empathy), social attitudes (e.g. agreement), or social relations (e.g. friendship). Although the comprehension of all social signals is important for social intelligence, social attitudes, which include cognitive states like interest, boredom, agreement and disagreement, are particularly common in any interaction whether that is among humans or between humans and machines. A machine that is able to detect such social signals could improve the experience offered to the user and even attempt to elicit a desired attitude. The focus of my Ph.D. is on three important social attitudes: agreement, disagreement, and interest. The primary goal of the project is to propose computational models that are able to capture the audiovisual dynamics that characterise the expressions of these three social attitudes and to detect them in a continuous audiovisual stream of spontaneous behaviour.

Agreements and disagreements occur daily in human-human interactions, and are inevitable in a variety of everyday situations. These could be as simple as finding a location to dine and as complex as discussing about notoriously controversial topics, like politics or religion. They occur in dyads, group meetings, televised debates, and even when interacting with multimedia material. In fact, one could argue that either agreement or disagreement occur almost always, as a verbal or nonverbal manifestation of them is bound to occur every time an opinion is expressed even in the simplest of conversations . Similarly, one's interest may be aroused for various reasons, such as a webpage with information one might be looking for, a painting one might find beautiful, a conversation on a topic one feels passionate about. Agreement and disagreement are frequently expressed verbally, but the nonverbal behavioural cues that occur during these expressions play a crucial role in their interpretation. That is naturally the case also for interest and other social attitudes.

The main aim of my Ph.D. is investigating whether it is possible to automatically infer social attitudes, and more specifically interest, agreement and disagreement, from complex constellations of nonverbal cues (facial expressions, prosody, gestures, postures, etc.) detected through widely available sensors, specifically microphones and monocular cameras. A secondary aim of my Ph.D. is to also explicitly describe what are the cues and dynamics that make each of the above attitudes unique.

The first step towards this aim was a survey of all nonverbal cues that could be relevant to agreement and disagreement [1] and the use of these cues towards creating models that can recognize spontaneous episodes of agreement and disagreement [2]. In [2], it was shown that (i) it is indeed possibly to recognize agreement and disagreement by using only non-linguistic cues, (ii) the temporal dynamics of these cues are vital to the task, and (iii) Hidden Conditional Random Fields, a relatively new computational model, is able to model these dynamics better and outperforms traditional machine learning techniques. Finally, a method to automatically analyze the Hidden Conditional Random Fields was presented which made it possible to confirm the findings of social psychology regarding what cues are most prevalent during the expression of these two attitudes.

References

[1] K.Bousmalis, M. Mehu, and M. Pantic, "Spotting Agreement and Disagreement: A survey of nonverbal audiovisual cues and tools," in Proc. IEEE Int'l Conf. Affective Computing and Intelligent Interaction, 2009, pp. 1-9

[2] K.Bousmalis, L.-P. Morency, and M. Pantic, "Modeling Hidden Dynamics of Multimodal Cues for Spontaneous Agreement and Disagreement Recognition", in Proc. IEEE Int'l Conf. Automatic Face and Gesture Recognition, 2011

Konstantinos Bousmalis
Department of Computing
Imperial College London

Conference Report: IEEE World Congress on Computational Intelligence 2010

The 2010 IEEE World Congress on Computational Intelligence, bringing together the IEEE International Joint Congress on Neural Networks (IJCNN), IEEE Congress on Evolutionary Computation (CEC) and the IEEE International Conference on Fuzzy Systems (FUZZ), was held for the first time in Europe, in Barcelona, from 18th to 23rd July.

On the day before the main conference began a number of tutorials and workshops took place, including an interesting talk from Arthur Kordon, a member of the Computation Intelligence (CI) group at Dow Chemical. He highlighted the differences between academic and business research, listed a number of reasons why businesses are, or should be, interested in CI, particularly in the area of competitive advantage, and described a couple of successful CI projects at Dow Chemical. He included a description of the application of genetic programming to a real industrial problem; a mathematician took three months to determine the mathematical function to describe experimental results for a real industrial product, while the genetic program took ten (unsupervised) hours to perform the same task.

The theme of industrial applications was carried through the conference, where an aerospace stream described the application of evolutionary techniques to aircraft design and the design of spacecraft trajectories.

In "Efficient aerodynamic optimization of a very light jet aircraft using evolutionary algorithms and RANS flow models", Emiliano Iuliano and Domenico Quagliarella described their work in designing the configuration of a wing and engine nacelle for a small business jet, concentrating on the aerodynamic performance in two specific operating conditions. The multi-objective evolutionary process was subject to a large number of geometric constraints, such as wing area and fuel tank volume, and aerodynamic constraints, such as lift and stall configurations. To reduce the cost of large numbers of long simulation runs, optimal wing section shapes were incorporated from past research. The properties of the evolved wing and engine designs were tested using a C F D system.

The process designed a natural laminar flow wing which was successful over a range of working conditions. The authors suggested that although the problem can be solved using standard evolutionary techniques, significant computational power is required, though it is well within the bounds acceptable for an industrial development process.

The problem of wing design is potentially within the huge number of design parameters and the many different operating points under which the wing must perform. A related but simpler single objective problem was described in "Double shock bump design optimization using hybridised evolutionary algorithms", by DongSeop Lee, Jacques Periaux, Jordi Pons-Prats, Gabriel Bugeda and Eugenio Oñate. A double shock control bump is a lump on the trailing edge of the top and bottom surface of a wing whose aim is to control the effect of the transonic shock wave in order to reduce drag, especially at the critical flight conditions where a double shock wave occurs. Each shock bump was defined by just three parameters: the length of the bump, the maximum height of the bump and the position of the maximum height within the bump.

The evolutionary algorithms used were based on the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), a popular benchmark choice at the conference. The first method was based on HAPMOEA, using hierarchical multi-population topology; the second replaced the hierarchical multi-population topology with Nash-Equilibrium. The research showed that the hybrid HAPMOEA-Nash- Equilibrium method was beneficial, both in terms of computational cost and design quality. The designed bump on an existing aerofoil resulted in significantly reduced transonic drag.

"Constrained global optimization of low-thrust interplanetary trajectories", by Chit Hong Yam, David Di Lorenzo and Dario Izzo, described the application of evolutionary computation techniques to the design of spacecraft trajectories from Earth orbit to the orbits of other planets in the solar system. The simulated spacecraft utilised a nuclear electric propulsion system and the aim of the research was to explore the trade-off between the starting mass of the spacecraft and the inter-orbit flight times. This represents an enormously complex non-linear programming problem, with a high number of non-linear constraints. Each trajectory was described as a series of timed engine burns at defined start times and in specific directions.

The authors applied the Sims-Flanagan transcription method to produce the non-linear programming problem, and then two different methods (basin hopping and simulated annealing with adaptive neighbourhood, both hybridised with a local search) to explore the solution space for global optimisation. The method made no use of expert knowledge and was initialised at randomly-generated trajectories.

One of the example trajectories described in Yam's presentation seemed highly unlikely. The trajectory took the spacecraft into an imaginative foray into the region of space beyond the target planet, from which it wended its way back to its destination. Although seemingly implausible, manual evaluation of the result demonstrated that the solution was indeed possible.

The authors are or have been members of the Advanced Concepts Team at the European Space Agency. As part of World Space Week the team ran an open competition for trajectory optimisation (see http://www.thespacegame.org/) which attracted 4,200 entrants, of whom 1,900 provided solutions. There is a link through the website for a new, permanent experiment in trajectory optimisation to which solutions may be submitted.

Dr Jacqueline Christmas
Postgraduate research student
University of Exeter

Report on AISB 09's "New frontiers in Human-Robot Interaction" Symposium

The AISB 2009 Convention took place on 6th – 9th April at the Edinburgh Conference Centre, located in the suggestive venue of the Riccarton Campus of Heriot-Watt University in Edinburgh. The theme of the convention was "Adaptive and Emergent Behaviour and Complex Systems" (chair: Dr. Nick Taylor). During the congress a number of parallel sessions took place, addressing very interesting and current topics in Artificial Intelligence (AI) and Robotics, directed by eminent scholars in these fields. The multidisciplinary approach of the addressed topics made it possible to acquire a multifaceted knowledge of the "state of the art" developments in these areas, giving wider perspectives on current and future studies in both basic and applied research.

A Symposium (http://homepages.feis.herts.ac.uk/~comqkd/HRI-AISB2009-Symposium.html/) was held in the same location on 8th and 9th April during the Science Festival (6-18 April 2009; http://www.sciencefestival.co.uk/). It focused on "New Frontiers in Human-Robot Interaction" (chair: Kerstin Dautenhahn, University of Hertfordshire). During the two days young and experienced researchers in AI and HRI field presented results and ongoing work from very different approaches and perspectives.

Improvements and developments in computers and complex machineries quality raised questions about how people can actually make good use of these technologies, as for example studies in usability, Human-Computer Interaction and Human-Machines Interaction point out. Development of robots has led to another field of application of this kind of knowledge, and now Human-Robot Interaction (HRI) is a growing research area with many contributions from different fields which have a big impact not only on the economic domain, but also on our well-being and lifestyle. Benefitting from contributions to studies within different areas of research and science, this symposium was aimed at providing a collaborative arena to discuss recent findings and challenges in HRI. In this highly interdisciplinary research field, in fact, an open dialogue among different approaches is desirable in order to contribute to the synthesis of a body of knowledge collecting HCI, robotics, psychology, social sciences, and other fields' results and observations.

The topics varied from the employment of robots in different fields (for example, personal and health care, assisted therapy and assistive technology, companions and helpers at home, schools and other educational environments or training, and more dangerous and complicated situations such as search and rescue) to more technical issues, such as sensor and interfaces, architectures for socially intelligent robots or user-centered robot design for creating the possibility of good interaction between humans and robots. This, in particular, was the focus of the symposium, as defined by the large number of contributions on Human-Robot Interaction (HRI), either from the ergonomic perspective or the psychological point of view. Some methods and methodologies to perform and analyze HRI were also proposed and discussed in the presentations.

Outstanding scholars in this field were invited to give keynote speeches, to fulfill the relevance of the argument and to make state of the art demonstrations in some of the covered fields. Holly Yanco (University of Massachusetts Lowell, USA) presented a talk entitled: "How to Partner with a Robot: The Design Space of Human-Robot Interaction", and addressed the delicate question of the design of human-robot interaction though many examples of robot systems in different application domains, such as assistive robots or the ones used for urban search and rescue, or educational ones.

Takanori Shibata (National Institute of Advanced Industrial Science and Technology, Japan) presented instead one of the most well-known and developed robots, used both as a simple companion and as therapy in hospitals: the famous Paro, a baby seal robot. In his presentation, "Integration of Therapeutic Robot, Paro, into Elderly Care in Denmark ", Takanori Shibata stated that a large number of institutions for the elderly, especially in Denmark, have started to use this incredibly tender and warm robot for caring for the elderly with dementia, and take advantage of his peculiar features for therapy. In some cases, in fact, Paro seemed to trigger an improvement in the mood and basic activities of the residents, showing how an interaction between a human and a robot can be not only psychologically, but also emotionally and sentimentally relevant for rehabilitation and care giving in such institutions.

Ruth Aylett (Heriot-Watt University, Scotland) delivered a presentation on "Affect, empathy and graphical characters". She pointed out some relevant questions about how to stimulate the desire of caring about a character from a user, and showed that it is necessary to create an "affective loop" between the user and avatar, since it seems that it is a basic requirement for good companionship and human-robot relationships. She considered robots with different shapes and appearance, and showed how certain graphic and aesthetic features of robots are more suitable than others for a good relationship – in particular, a pleasant and emotionally relevant graphic in human or animal shape- and how these features are linked to the purpose the robot is designed for.

During the Symposium the contributors addressed largely differentiated issues in AI. For example, Muhammad Ali et al. (CNRS-LAAS, France) proposed an "Architecture Supporting Proactive Robot Companion Behavior". Proactive behavior is a certain kind of anticipatory behavior that a robot should have in order to successfully perform collaborative tasks with humans, which allows it to determine by itself if, how and when it can intervene and help a human in his work or activities. The architecture they proposed is basically derived from human frontal lobes organization and functioning (especially in planning processes), and allows the robot to select certain high level goals from a correct recognition of the situation and scenario (thanks to a Chronicle Recognition System; CRS) that can be refined and changed by means of constant updating and refreshing. Similarly, and based on neuro-cognitive mechanisms about joint and shared actions in humans, Estela Bicho et al. (University of Minho, Portugal) presented a design for robots that allows the joint coordination of actions and goals as a dynamic process that integrates contextual cues, shared task knowledge and the predicted outcome of the user's motor behavior, which includes a basic form of error monitoring and compensation. Aimed at the same capacity of an anticipatory behavior of machines was also the study from Aris Alissandrakis and Yoshihiro

New frontiers in Human-Robot Interaction (cont.)

Miyake (Tokyo Institute of Technology, Japan) on routine home tasks, showing how the human feedback provided to the robot is necessary for learning tasks correctly and fulfilling human expectation, in a system perspective.

More technical issues were also touched upon, like in the "Photogeometric Sensing for Mobile Robot Control and Visualisation Tasks" presentation by Alonzo Kelly et al. (Carnegie Mellon University, USA; DCS Corporation, USA), in which new types of sensors were presented which can produce high fidelity imagery from a camera installed on a moving robot, with important effects in remote visualization of recorded environments. "Haptic Control for the Interactive Behavior Operated Shopping Trolley InBOT" was another very specialised presentation in which Michael Göller et al. (FZI Research Center for Information Technology, Karlsruhe, Germany) presented a shopping trolley which can be guided by haptic control, thanks to a sensible handle-bar that is able to translate different pressure and gestures in commands to move itself and even follow the person from afar, or "park" itself aside in order not to obstruct the way. Again, remote control of robots difficulties were considered by Alan Atherton and Michael A. Goodrich (Brigham Young University, USA), who proposed "Supporting Remote Manipulation with an Ecological Augmented Virtuality Interface" in order to improve operators' situational awareness (SA) in robot-guidance. The work by Alberto Valero et al. (SAPIENZA - Università di Roma, Italy) also tried to extend some cognitive-based knowledge into the main issue of Human-Robot Interaction in robot guidance to improve SA of operators, showing that integration of survey and route knowledge in a desktop interface can lead to a better SA, with respect to a claimed advantage of operator mobility using, for example, a Portable Device Application (PDA).

Cognitive science's findings and knowledge can also be usefully employed to shape HRI designs, as Tom Carlson and Yiannis Demiris (Imperial College London, UK) showed in their work "Using Visual Attention to Evaluate Collaborative Control Architectures for Human Robot Interaction", in which they illustrated that, contrary to their expectations, assistive and collaborative technologies require more users' attention and cognitive resources, concluding that further work still has to be done in this field. In some cases robots are also thought to be useful in therapy, like Mark B. Colton et al. (Brigham Young University, USA; Honda Research Institute, USA) propose in their work "Toward Therapist-in-the-Loop Assistive Robotics for Children with Autism and Specific Language Impairment". Empirical findings within the field of language ("Heuristic Rules for Human-Robot Interaction Based on Principles from Linguistics-Asking for Directions", by Andrea Bauer et al. Technische Universität München, Germany; "A WOz Framework for Exploring Miscommunication in HRI" by Theodora Koulouri and Stasha Lauria, University of Brunel, UK) and proxemics ("An Empirical Framework for Human-Robot Proxemics"; Michael L Walters, et al., University of Hertfordshire, UK) were also discussed in order to build basic knowledge for interaction and collaboration between humans and robots, especially in communication and physical interaction between them.

Furthermore, evidence from neuroscience studies, namely, the existence of mirror-neurons systems in human brain, which allows a person to recognize other people's actions and movement schemes, can give some interesting insights into human perception of robots' movements and behavior, like a study from Aleksandra Kupferberg et al. (Ludwig-Maximilians-Univ. Munich, Germany; Technical Univ. München, Germany), "Video Observation of Humanoid Robot Movements elicits Motor Interference" , showed. The different perception of robots' appearance was also investigated by Michael L. Walters et al.'s (University of Hertfordshire, UK) study of "Preferences and Perceptions of Robot Appearance and Embodiment in Human-Robot Interaction Trials", showing that some robots' physical features are perceived by humans as expression of temperamental traits, like the presence or absence of intelligence or intelligent behavior and the reliability of the robot as a helper.

Moving into the psychology of personality, sociology and social psychology domain, "The role of expectations in HRI" from Manja Lohse (University of Bielefeld, Germany) uncovered a series of problems connected to what humans expect from a robot and their willingness to accept them as helpers or companions. The link between expectations and appearance of robots was investigated in children in the presentation "The boy-robot should bark! – Children's Impressions of Agent Migration into Diverse Embodiments" by Dag Sverre Syrdal et al. (University of Hertfordshire, UK). The creation and evaluation of some scales which can give a measure of human expectations and acceptance of robots was discussed in "The Negative Attitudes towards Robots Scale and Reactions to Robot Behaviour in a Live Human-Robot Interaction Study" presentation by Dag Sverre Syrdal et al., (University of Hertfordshire, UK), as well as in "The USUS Evaluation Framework for Human-Robot Interaction Astrid Weiss et al. (University of Salzburg, Austria; University Paul Sabtier, France); a huge framework considering Human-Robot Collaboration with humanoid robots addressing usability, social acceptance, user experience, and societal impact (abb. USUS) was proposed. Robots' impact on society was also investigated in another talk by Astrid Weiss et al. (University of Salzburg, Austria; University Paul Sabtier, France; National Institute of Advanced Industrial Science and Technology, Japan), "Addressing User Experience and Societal Impact in a User Study with a Humanoid Robot"

The aim of all these talks was to establish a good relationship between humans and robots and promote collaborative behaviors within them. Connected to this, proposals for improved design of robots were presented in a number of talks, including the "Iterative design process for robots with personality" presentation by Bernt Meerbeek et al. (Philips Research, NL; Univ. of Tech. Eindhoven, NL), "Affective-Centered Design for Interactive Robots" by Laurel D. Riek and Peter Robinson (University of Cambridge, UK), "Baby steps: A design proposal for more believable motion in an infant-sized android", presented by Silpa Wairatpanij et al. (Indiana University, USA), and "Creating Trustworthy Robots: Lessons and Inspirations from Automated Systems" in the presentation by Munjal Desai et al. (University of Massachusetts Lowell, USA; Carnegie Mellon University, USA).

After the Symposium, a brief discussion was held between the participants

New frontiers in Human-Robot Interaction (cont.)

analyzing the usefulness and feedback of this initiative, pointing out and stressing the strengths of these interventions and the use that the received comments and suggestions from people working in the same areas but with different perspectives can have for everyone's work, especially when a project is in progress and feedbacks and feedback or observations are needed to figure out if it is moving in the right direction or not. The chance to discuss these arguments and issues from different points of view with more experienced researchers and scholars totally enriched participants' minds, while the presentatons covered wide areas and topics, allowing the attainment of broader knowledge for both speakers and listeners.

Chiara Saracini,
"Sapienza",
University of Rome
Chiara.saracini@uniroma1.it

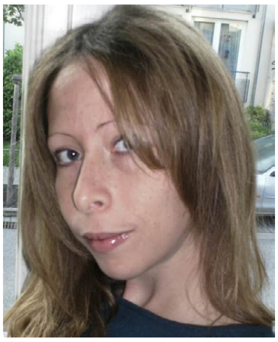

2011 Loebner Prize

The annual instantiation of the Turing Test, the Loebner Prize, will take place on Wednesday 19th October 2011 in the South-West of England in the College of Engineering, Mathematics and Physical Sciences at the University of Exeter.

The prize, named after its founder and philanthropist Hugh Loebner, is an annual world-wide contest to test the state-of-the-art in artificial intelligence (AI). The top four AI programs will compete to win a share of a prize fund worth $5750, with additional prizes made available this year for a new Junior Loebner Prize, a version of the contest where children will act as judges.

For more information or if you would like to get involved in this year's contest, please visit the contest website at http://loebner.exeter.

Society News

AISB/IACAP World Congress 2012

In line with the celebrations of Alan Turing's life and work in 2012 and given the significance of Turing's work both to AI and to the philosophical ramifications of computing, the 2012 AISB Convention will be replaced by the AISB/IACAP World Congress 2012 (initial website: http://events.cs.bham.ac.uk/turing12/). IACAP is the International Association for Computing and Philosophy (http://www.ia-cap.org/).

The intent of the conference is to stimulate a particularly rich interchange between AI and Philosophy on any areas of mutual interest, whether or not directly inspired by Turing's own work. Two plenary keynote speakers have so far been secured, namely Aaron Sloman (a Fellow of AISB) and Luciano Floridi (former President of IACAP and also AISB Fellow).

The Congress will have a similar structure to AISB conventions, and a call for symposia etc. will be announced in due course. It will be held at the University of Birmingham, 2nd to 6th July 2012. It will be co-chaired by Tony Beavers, President of IACAP, and John Barnden, Vice-Chair of AISB, with Manfred Kerber (Birmingham) as Local Chair. The Organizing Committee is largely constituted of AISB and IACAP representatives.

Some research themes we have identified as being especially appropriate for the Congress are as follows, but are in no way exclusive:

- the fundamental nature and limits of computation
- computational theory of mind
- the nature and possibility of AI
- testing for intelligence
- consciousness (natural or artificial)
- creativity (artistic and otherwise), aesthetics, etc.
- people's attitudes towards and relationships with intelligent machines
- ethics of AI and computing in general, and how AI may ultimately affect ethics
- the philosophical nature & ramifications (e.g. for notions of self) of intelligent software agents in cyberspace.

We are also interested in holding, in conjunction with the Congress, a series of artistic events (in any area of the arts) that are relevant to such themes.

Please contact Professor John Barnden (J.A.Barnden@cs.bham.ac.uk) if you have ideas for any aspect of the Congress (e.g., a symposium proposal) or for possible cultural events.

David Everett Rumelhart

David Everett Rumelhart, one of the pioneers of the renewed interest in connectionism, passed away on 13th March 2011.

Rumelhart, who has been an eminent figure in mathematical psychology, artificial intelligence and analytic approaches to cognition, made a significant contribution to the field of cognitive neuroscience by contributing to the development of the notions of parallel distributed processing and connectionism.

Connectionism is an area of machine learning eschewing the golden but old-fashioned 'Artificial Intelligence and symbolic approaches to cognition', in favour of distributed models mimicking the operation of the nervous system. His Parallel distributed processing: Explorations in the microstructure of cognition, co-authored with James McClelland in 1986, was at the time a landmark volume in cognitive science and sold in excess of 100,000 copies.

This work was in part responsible for the renaissance of research in artificial neural networks, which had been stymied by earlier difficulties and relative success of the classical Artificial Intelligence based on symbolic approaches.

Rumelhart's many accolades included a membership of the National Academy of Sciences of the USA, a MacArthur 'Genius' Award, and the SEP Warren Medal received jointly with McClelland in 1993. The Glushko-Samuelson Foundation 'David E. Rumelhart Prize for Contributions to the Theoretical Foundations of Human Cognition' was established in his honour in 2000.

His passing creates a significant gap in connectionist and cognitive science communities.

Dr Slawomir J Nasuto
Reader in Cybernetics
School of Systems Engineering
University of Reading

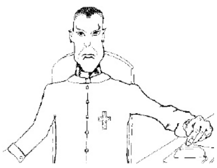

Dear Aloysius...

Fr. Aloysius Hacker answers your questions

Cognitive Divinity Programme
Institute of Applied Epistemology

About the Society
The Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB) is the UK's largest and foremost Artificial Intelligence society. It is also one of the oldest-established such organisations in the world.

The Society has an international membership of hundreds drawn from academia and industry. We invite anyone with interests in artificial intelligence or cognitive science to become a member

AISB membership includes the following benefits:

- Quarterly newsletter
- Student travel grants to attend conferences
- Discounted rates at AISB events and conventions
- Discounted rates on various publications
- A weekly e-mail bulletin and web search engine for AI-related events and opportunities

You can join the AISB online via:

http://www.aisb.org.uk

Dear Aloysius,

I have been accepted to study for a degree in AI at the University of Poppleton, but am unable to afford the £9,000/annum fee. I'm desperate to study AI. Please help!

Yours, Impoverished Student

Dear Impoverished Student,

You've come to the right person. In collaboration with an employee of the Student Loans Company, my Institute has developed a service that will solve your problem. The WIPE SLATE™ (Wangle Investment in your Personal Education: Student Loans Awarded and Then Erased) will enable you to complete your degree at zero cost to yourself. Just apply for a loan to cover all your maintenance and tuition fees costs, plus my own very reasonable remuneration. WIPE SLATE™ will then communicate with the Student Loans Company's computer to ensure that all records of your loan are wiped from its memory. To balance its books, your charges will be deducted from the bank account of Mr. Nick Clegg MP. Your education will be completely free.

Yours, Aloysius

Dear Aloysius,

I am the glorious and much loved leader of a middle-eastern state. A conspiracy between small and unrepresentative groups of gangsters, imperialist fifth columnists and religious fundamentalists has been stirring up opposition against me. I fear for my life, but cannot abandon my loyal subjects when they need me to defend them against the plotters. Can your Institute's world-famous AI technology come to my aid?

Yours, Glorious Leader

Dear Glorious Leader,

I think we can help you. A combination of holographic and speech technology can be used to enable you to defend your homeland against the conspirators whilst fully protecting your life. Using our SHYSTER™ (Speech and Holography Yield a Simulation That Emulates Reality) system, your image and speech can be transmitted to a press conference or palace balcony from a distant place of safety. The technology is now so convincing that everyone will assume that you are still at your post. We can even offer a post-life service. With automatic speech recognition and generation, together with the last holographic avatar technologies, you can appear to live forever. We have successfully tested this technology in even the most exacting conditions: a 'live' interview with Jeremy Paxman;, so we are therefore confident that it will fully meet your requirements.

Yours, Aloysius

Dear Aloysius,

Working for an international bank in the City of London, I developed the AI software that automates trading in the money markets. My software, and that of my colleagues, is largely responsible for the dramatic changes in the banking scene over the last few decades. My bosses, none of whom has any understanding of computing, have awarded themselves large bonuses for presiding over these changes, while doing little to bring them about. Meanwhile, we AI researchers, economists and mathematicians, who did create this revolution, have seen scant reward. Is this fair?

Yours, Hard Done By

Dear Hard Done By,

Indeed, it is not fair at all. You deserve as much credit for the current state of the economy as the bankers - and you should receive similar rewards. Fortunately, my Institute has developed just the product to right this historic wrong. BOONDOGGLE™ (Bonus Organiser Offers Numerous Dollars for Overly Gung-ho Gambling that Leads the Economy). As a computing expert you can install BOONDOGGLE™ in your bank's payroll system and you should see your efforts appropriately rewarded.

Yours, Aloysius

Agony Uncle Aloysius, will answer your most intimate AI questions or hear your most embarrassing confessions. Please address your questions to fr.hacker@yahoo.co.uk. Note that we are unable to engage in email correspondence and reserve the right to select those questions to which we will respond. All correspondence will be anonymised before publication.