RTD – SAS Executive Summary

Reporting Period: 01.01.2003 – 31.12.2003

1. *Achievements*

The *Research Theory and Development – Smart Adaptive Systems* Committee has been a highly active committee within EUNITE, contributing to a variety of important areas. There are 16 nodes within the committee, most of which have been active participants as key nodes. The particularly active key nodes have included:

• ASAP, UK	Jon Garibaldi
• DIBE, Italy	Davide Anguita
• OFAI, Austria	Georg Dorffner
• ENEA, Italy	Stefano Pizutti
• NTU, UK	Ahmad Lotfi
• UEP, Czech Republic	Petr Berka
• UPC, Spain	Cecilio Angulo-Bahon
• EIA, Spain	Josep Vehi
• DECSAI, Spain	Jose-Luis Verdegay
• UCT, Bulgaria	Nikolinka Christova
• IDSIA, Switzerland	Jurgen Schmidhuber

Also active in contributing research material for the committee have been Plamen Angelov (Lancaster University, UK; originally of Loughborough University, UK), Joaquim Armengol, Radim Jirousek (UEP, Czech Republic), Mincho Hadjiski (UCT, Bulgaria), Janusz Kacprzyk (Polish Academy of Sciences, Poland) and Robert John (DMU, UK), as well as numerous other researchers within EUNITE (but outside the RTD-SAS committee and numerous external contributors to material at research meetings).

Major achievements that members of the RTD-SAS committee have contributed to include:

- *The EUNITE Roadmap Document*. A significant body of material has been contributed to the EUNITE roadmap document by RTD-SAS members. The current state of the art in each of the original technique areas envisaged as being encompassed within EUNITE have been catalogued, including: artificial neural networks, evolutionary computing, fuzzy logic, and fuzzy control. In addition many other areas of artificial intelligence have also been covered including support vector machines, machine learning, probabilistic methods, hardware implementation, universal learning algorithms, interval methods and kernel methods. This roadmap represents a major achievement in collating such wide-ranging material on the state-of-the-art and potential future research directions in one document. This material should prove invaluable to both domain experts and newcomers to the field together.
- *The Bibliographical Database of Smart Adaptive Systems*. To accompany the RTD-SAS roadmap material, a bibliographical database of major research works in truly *smart adaptive systems* has been created in both hardcopy and online form. This bibliography brings together references from all the individual research fields mentioned above, categorises them according to their level of adaptivity (see Section 2 below), learning capability and EUNITE industrial application area. As well as categorising each reference, the bibliography provides short summaries of the contribution and importance of each reference. The purpose of this resource is not to compete with other online bibliographic databases (in which, almost universally, coverage as measured by total number of references is taken as the indicator of quality) but rather to act as a virtual expert to guide newcomers, novices and interested non-experts (such as industrialists) to the area. Thus the emphasis has been placed on the quality of the references rather than quantity, with each having been carefully reviewed and vetted by domain experts before being entered

onto the database.

- *SAS Research Theory & Development Meetings*. Numerous national and international research conferences, workshops, and special sessions have been organised by members of RTD-SAS, many of which have received modest financial support from the RTD-SAS committee. These have attracted contributions both from all members of the EUNITE Network of Excellence and the wider European and International research communities. The financial contributions of RTD-SAS have generally been used for purposes such as encouraging participation from junior (pre-doctoral) researchers or post-doctoral students at an early stage in their career, or for attracting leading researchers to give plenary talks at the meetings. In each case, smart adaptive systems have been the specific focus of the meeting (sometimes in the form of a special interest session at a wider-ranging or more general conference). These meetings have helped foster discussion on definitions of SAS, encouraged submission of research papers describing theoretical developments and real applications of SAS, and brought together researcher with an interest in SAS. More details on the specific meetings held during the EUNITE period are given below, in Section 5.
- *The EUNITE 2003 RTD-SAS Best Paper Award*. A generic *Research Theory and Development of Smart Adaptive Systems* best paper award was organised, judged and awared during the EUNITE 2003 annual conference, held in Oulu, Sweden. This award was seen as a mechanism to encourage and reward high quality research in the area, and was specifically aimed at identifying the best *smart adaptive system* in the eyes of the expert judges. The overall winner of the prize was a paper "Classification of Surface Form Deviations for Quality Analysis" by Döring, Eichhorn, Klose and Kruse. This paper described both theoretical aspects of a smart adaptive system as well as a practical implementation. The winning authors have been invited to give a plenary talk at the closing conference of EUNITE in order to further assist in the dissemination of the concepts and realisation of smart adaptive systems to the conference participants.
- *Contributions to the SAS Best Practice Guideline Publication*. Members of RTD-SAS have made specific chapter contributions to the SAS Best Practice Guidelines publication managed by the RTD-IM committee. This publication will, once again, collate the state-of-the-art of SAS as interpretted by domain experts and make this material available to the interested (though not necessarily expert) target audience. A typical person to whom this guide is targetted is, for example, an industrialist who wants to try to create a smart adaptive system in a specific application area, but who has previously had little or no experience in the area.
- *The ISAAC Server*. Industry is more interested in solving particular problems than in the algorithm used to solve them; academics, instead, focus mainly on the methodology and, sometimes, do not even provide comparison with existing methods. The *Internet Smart Adaptive Algorithm Computational (ISAAC) Server* was proposed to overcome these drawbacks and to help fill the gap between industry and research, providing a computational server for solving real-world problems, albeit on a smaller scale, through smart adaptive technologies. An Internet-based server has been constructed to allow an end-user to find application-solving methods and algorithms, submit their problem, and obtain the result in a batch fashion, through e-mail messages. Problems belong to several application domains such as classification, time series prediction, function estimation, data clustering, novelty detection, signal denoising, etc. Methods are based on techniques such as neural networks, fuzzy systems, genetic algorithms, etc. A user is allowed to submit a problem through a web page and obtain the result by email from the ISAAC Server. For further investigation, the user is addressed to the developers of the particular method that proved the most satisfactory to solve the problem.
- *Formation of New Research Collaborations*. A significant number of research proposals featuring collaborations between individual members of RTD-SAS and their institutions have been submitted as a result of general networking activities performed within EUNITE. Examples of such novel research collaborations include FP6 expressions of interest in integrated projects, FP6 proposals for Networks of Excellence, several FP6 STReP's and other smaller

scale alliances.

- *Other Activities.* RTD-SAS members have contributed to many other activities organised through other committees, including task forces such as *Taxonomy*, *Gene-Expression-Based Individualised Medicine* and *Safety Critical Aspects of Smart Adaptive Systems*, summer schools, supplying of benchmark problems, etc.

2. *Status of The Theory of Smart Adaptive Systems*

The committee has agreed a formal definition of 'adaptive' in the context of *smart adaptive systems* which in combination with a working definition of the word 'smart' defines the field. A further, more general and wider-ranging meaning of the word 'smart' has also been accepted which helps to define where the future of the area lies. The accepted fomal definition of 'adaptive' is a three-level meaning as follows:

1. adaptation to a changing environment;
2. adaptation to a similar setting without explicitly being ported to it;
3. adaptation to a new/unknown application.

In the first case the system must adapt itself to a drifting (over time, space, etc.) environment, applying its intelligence to recognize the changes and react accordingly. This is probably the easiest concept of adaptation for which examples abound, e.g.: customers preferences in electronic commerce system, control of non stationary systems (e.g. drifting temperature), or telecommunication systems with varying channel or user characteristics. In the second case, the accent is more on the change of the environment itself more than on a drift of some features of the environment. Examples are systems that must be ported from one plant to another without explicitly perform a porting of their main parameters, or financial applications that must be ported from a specific market to a similar one (e.g. a different geographical location). Other interesting problems are also the porting of compiled software from one processor to another (obviously without recompilation) or the porting of algorithms, usually developed in software, to a dedicated, and often resource-limited, hardware for embedded systems. The third level is the most futuristic one, but its open problems have been addressed already by a number of researchers, especially in the machine learning field where, starting from very little information on the problem it is possible to build a system through incremental learning.

The working definition that has been adopted by the RTD-SAS committee and the EUNITE network as a whole is that 'smart' implies that intelligent techniques must be involved in the adaptation of a system for it to be considered a 'smart adaptive system'. Hence a conventional classical control theory application, while adaptive within its own domain terminology and possible adaptive within level 1 defined above, is ***not*** a smart adaptive system because it contains no aspect of computational intelligence. That is to be a *smart adaptive system*, a system must be adaptive to a level defined above **and** must utilise one or more technique(s) from the list of: artificial neural networks, evolutionary computing, fuzzy logic, fuzzy control, support vector machines, machine learning, probabilistic methods, hardware implementation, universal learning algorithms, interval methods or kernel methods. This we denote *narrowly smart*. In this context, the following expanded descriptions are derived.

1. Adapting to a Changing Environment. The adaptation to a changing environment can be considered the basic characteristic of a smart system. A very revealing example of this kind of behaviour comes from the ANN field: most ANNs are built through the solution of an optimisation problem that finds the optimal weights and topology according to a certain criterion. After this learning phase, the network is frozen and put to work processing new data. Unfortunately, using this approach, the behaviour of the network is completely defined by its learning phase and, even though the network shows good generalization, it will always produce the same result when stimulated with the same data. What is needed to add 'smartness' is the adaptation of its parameters or topology to reflect the changing environment. Perhaps the most advanced field on this topic is the broad area of machine learning (ML) where the problem of adaptation through learning has been a core research issue for a long time. ML covers both symbolic methods (decision trees and rules, etc.), sub-symbolic methods (neural networks, Bayesian networks, etc.) and has several connections with traditional statistics (discriminant or regression analysis, cluster analysis, etc.).

Evolutionary Computation has started to address adaptation quite recently: examples are the adaptation of Genetic Algorithms to non-stationary environments, with adaptive and self-adaptive techniques, and the use of memory for storing good, partial solutions and reuse them later, when environment changes occur. Fuzzy Systems, as mentioned above, have been also used in dynamical contexts: an example is the adaptation or evolution of rule-based models, treated as a learning process. Some applications combine traditional techniques with Fuzzy Logic to build adaptive control systems. In the field of ANN, the problem of adaptation to a changing environment is well known, where the term *catastrophic interference* has been commonly used. The main problem is the need of both short-term memory, for incremental or on-line learning, and long-term memory for recognizing context drifting. A possible answer to this problem is the definition of an effective and reliable method for model selection, or in other words, the ability to predict with good confidence the generalization ability of a network: some results are starting to appear in the literature but questions remain about their applicability to practical problems.

2. Adapting to a Similar Setting. The adaptation to similar settings is a very interesting property that would avoid the explicit porting from one environment to another. Examples from ML that try to use a successful past solution and adapt it to a similar problem rely on the idea of reasoning by analogy. This is a principle that can be found in Case-Based Reasoning (CBR) systems applied to both simple and structured knowledge representations. A CBR system works along a cyclic process, retrieving the most similar cases, reusing the cases in the attempt to solve the problem, revising the proposed solution, if necessary, and retaining the new solution as a part of a new case. Some authors propose the use of fuzzy sets to build heuristics for solving classes of problems: examples include fuzzy sets being used to develop heuristic optimisation tools (including hard combinatorial optimisation problems). Examples based on neural networks are rarer; however the subsymbolic approach of the ANNs is a natural candidate for the embedding of neural computation on sensors and actuators. Indeed, while the idea of building a neural computer is slowly fading after many years of research, the proposal of neurointerfaces appears more promising. This approach introduces new problems for adapting successful neural algorithms, usually developed as software modules on conventional computers, to embedded analog or digital hardware.

3. Adapting to Solve a New Problem. The idea of a system that evolves by itself and finds autonomously the solution of a problem or suggests new ways to solve a problem could appear visionary, at least. In some way, we could assign to this area most of Data Mining research, as it tries to find new relations and information from a large variety of data. There is however, some interesting and specific research in this area coming from both the ML and ANN fields. If the constraint of starting from zero knowledge is relaxed, systems for knowledge revision and refinement can be cited: given an incomplete knowledge and a set of examples, they modify the knowledge to be consistent with the examples, thereby building new knowledge or, at least, improving existing. From the connectionist field comes another example that aims at creating new knowledge from scratch: the Creativity Machine (CM). According to its creator this system is able to generate new concepts and has been successfully applied to real-world applications.

On a more general level, the word 'smart' can be taken to have a meaning closer to the traditional dictionary definition of 'intelligent': i.e. *the ability to learn or understand or to deal with new or trying situations; or the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (as tests)* [The Merriam-Webster's Dictionary]. This is beyond even level 3 adaptation as defined above. By carefully picking words, level 3 adaptation could be interpreted as 'the ability to … deal with new … situations'. However, the intended meaning of 'smart' is surely above and beyond this strictly limited interpretation. Rather, it involves 'understanding' and 'the ability to think abstractly'. This we denote *generally smart*.

3. *Do Smart Adaptive Systems Exist?*

The answer is a qualified yes. There are, clearly, examples of theoretical and applied work in which intelligent techniques (as defined above) have been utilised to create systems that are adaptive to levels 1 and 2. All entries on the bibliographic database (currently numbering over 400 references) satisfy this minimum criteria –i.e. references are not allowed onto the database if they are not both smart and adaptive to the narrow definition. Finding examples of level 3 adaptive systems, even when 'smart' only to the more narrow definition has proven much more difficult. Having said this, there are at least 25 references on the database which have been carefully assessed by the experienced reviewers to exhibit level 3 adaptivity.

However, we are far from systems that are truly 'smart' according to the more general definition given above. It must be stated that whereas good working definitions of the term 'adaptive' have been generally accepted, the term 'smart' has eluded similar definition. We are left with the narrow definition which is little more than a wider set of specific techniques than the traditional 'soft computing' techniques, or the general definition which differs little from the old fashioned and no longer much researched concept of general artificial intelligence. Given then that 'generally smart' is currently little different to 'generally intelligent', it is hardly surprising perhaps that generally smart adaptive systems are not yet found.

In short, (narrow) smart (level 1 and 2, and fewer level 3) adaptive systems do exist; general smart adaptive (particularly level 3) systems certainly do not exist.

4. Predictions of Future Research Theory and Developments in SAS

As previously observed, the art of making predictions of the future is an extremely difficult task. It is easy to appear either unimaginative, with predictions which are little more than small incremental steps from the current state of thte art; or ridiculous, with predictions which are little short of nonsensical science fiction. With this thought in mind, therefore, it is hardly surprising that the RTD-SAS committee has, as a whole, been quite cautious in making firm predictions. Anguita has suggested that taking steps towards 'ubiquity' is the key next research topic, while Schmidhuber has suggested a 'radical change' might occur by 2015 when, by extrapolation of Moore's Law (that computing power doubles every 18 months), the fatest computers will match the human brain in terms of raw computing power.

What can be stated with certainty is that computing power will continue to rise at a significant pace over the next few years. Although general intelligence has remained an elusive 'Holy Grail' for computer scientists, it is instructive to look at advances in game playing that have come about, largely due to increases in computer power. Computers have now reached the highest human performance in othello, checkers (draughts), and chess. In conjunction with this power increase, is an increase in techniques available to exploit this power. Relatively recent techniques such as, for example, ant collony optimisation, artificial immune systems and swarm intelligence are gaining growing acceptance alongside the now long established techniques of fuzzy logic, evolutionary computation and artificial neural networks. The generally anticipated increase in computing power in conjunction with a far richer range of 'intelligent' techniques will surely allow progress towards 'smartness' in a wider range of application areas.

Increasing hybridisation of techniques and the gradual blurring of the distinctions between the various techniques would seem to be a trend that is likely to continue. Some of the biggest challenges in achieving truly level 3 adaptation in practical applications include questions on how to evaluate such systems, and the robustness and stability of them. How can a system be tested and certified to a high enough standard to be released commercially (e.g. in a consumer device) if it has the ability to modify itself to adapt to new environments? In some senses the consumer (and researcher) wants the best of both worlds. A system must be proven, measured, verified, demonstrated to achieve good performance in a domain, but must also be able to adapt to changes in the domain. It is not currently possible to resolve this conflict between predictability and adaptability. These areas must receive attention is smart adaptive systems are to be realised.

5. *Major Events in the RTD – SAS Area*

As stated in Section 1 above, a large number of conference, workshops and special sessions have been organised by members of the RTD-SAS committee. A selected list of major events includes:

- EUNITE 2001, Aachen, Germany.
- ANNAS 2002 (Artificial Neural Networks and Adaptive Systems) Session at the 16th European Meeting on Cybernetics and Systems Research, Vienna, Austria.
- Special Session at ECML/PKDD 2002, Helsinki, Finland.
- Special Session at 1st IEEE International Symposium on Intelligent Systems (IS 2002), Varna, Bulgaria.
- Special Session at ICANN 2002, Madrid, Spain.
- Two SAS Special Sessions at EUNITE 2002, Albufeira, Portugal.
- SAS Special Session at RASC 2002, Nottingham, UK.
- Workshop on Uncertainty Processing (WUPES) 2003, Hejnice, Czech Republic.
- Two SAS Special Sessions and an overall RTD-SAS Best Paper Award at EUNITE 2003, Oulu, Finland.
- Special Session on Adaptive Kernel Methods at 7th International Work Conference on Artificial and Natural Neural Networks (IWANN) 2003, Menorca, Spain.
- Special Session on Adaptivity at ECML/PKDD 2003, Cavtat (Dubrovnik), Croatia.
- To be held: Smart Adaptive Agent Applications (SA3) Joint ALAD and EUNITE Workshop at EUNITE 2004, Aachen, Germany.