Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Artificial Intelligence in Games

A Survey of the State of the Art

M.A.J. Bourassa and L. Massey

Defence R&D Canada – Ottawa

Technical Memorandum
DRDC Ottawa TM 2012-084
August 2012

Canada

Artificial Intelligence in Games

A Survey of the State of the Art

M.A.J. Bourassa
Defence R&D Canada – Ottawa

L. Massey
Royal Military College of Canada

Defence R&D Canada – Ottawa

Technical Memorandum
DRDC Ottawa TM 2012-084
August 2012

Principal Author

Original signed by M.A.J. Bourassa

Approved by

Original signed by J. Tremblay-Lutter
Head/CARDS Section

Approved for release by

Original signed by C. Macmillan
Head/Document Review Panel

Abstract

The use of computer games as a means of training has seen considerable growth in industry and the military. In particular, the military has found computer games to provide a cost-effective, realistic, and safe environment to augment actual field exercises. Computer games are also easily made available to individuals or small teams. A challenge to the use of computer games for training is that computer games must frequently be populated with computer generated allies and enemies known commercially as *non-player characters* (NPCs). The quality of the training or game experience will therefore be dependent on the quality of the non-player characters who should demonstrate realistic human-like behaviour. This paper surveys the current state-of-the-art of the application of artificial intelligence techniques to create realistic behaviour in computer generated characters in games. The survey considers both commercial and military computer games in order to establish the state of the technology. Conclusions are drawn on the maturity and applicability of the surveyed approaches and highlights directions for future research.

Résumé

Dans les domaines industriel et militaire, on assiste à une forte croissance de l'utilisation des jeux informatiques comme outils d'instruction. Dans le cas plus particulier de l'arme, on constate que les jeux informatiques fournissent un environnement rentable, réaliste et sûr qui vient compléter les véritables exercices en campagne. De plus, les jeux informatiques sont faciles d'accès pour les personnes et les petites équipes qui ont accès à un ordinateur. Cependant, il faut y introduire fréquemment des alliés et des ennemis virtuels, appelés dans le commerce " personnages non-joueurs ". Il va de soi que la qualité de l'instruction ou du jeu dépend en grande partie de la qualité des personnages non joueurs. En effet, les personnages générés par l'ordinateur doivent avoir un comportement reproduisant avec réalisme celui d'un être humain. Le présent document recense les applications de pointe de techniques d'intelligence artificielle à la création de personnages virtuels réalistes dans les jeux. Nous avons examiné les jeux tant commerciaux que militaires afin de dresser un portrait courant de la technologie de pointe dans le domaine. Nous présentons ensuite des conclusions quant à la maturité et à l'applicabilité des approches recensées, après quoi nous présentons quelques voies que pourraient emprunter les travaux de recherche à venir.

This page intentionally left blank.

Executive summary

Artificial Intelligence in Games

M.A.J. Bourassa, L. Massey; DRDC Ottawa TM 2012-084; Defence R&D Canada – Ottawa; August 2012.

Background: Computer games occupy a large and profitable market niche as much for their entertainment value as for their use in "serious games". The latter are game applications applied to the world of modeling and simulation. Serious games are increasingly viewed as a cost effective, and often safer, means of training personnel in demanding and dangerous occupations. Serious games are becoming a larger presence within the Department of National Defence where they are used for training and operational development.

Serious games used for operational training must necessarily be as realistic as possible. Commercial and serious games have both seen great leaps in the quality of their acoustic and visual renderings. It is fair to say that the current graphics quality is good enough to convincingly immerse players in a synthetic world. But the quality of the game play is also greatly affected by the quality of the behaviour of the computer generated opponents. Therefore, future improvements in commercial and serious games is likely to come from improvements to the credibility of computer generated opponents. In short, the latter must behave believably, that is, in a human-like manner.

Principal results: This paper surveys the use of artificial intelligence (AI) techniques that have been, or are currently being, used in the gaming world to enhance the credibility of computer generated opponents and allies. The techniques outlined draw from modern computer and neuroscience, as well as from decades of research in sociology and psychology. The survey looks at the use of AI in both commercial entertainment games and in serious games, in particular, computer generated forces in use by the military such as: OneSAF, JSAF, and VBS2.

The principal conclusion is that there is no clear leader in game AI technology either in industry, government, or academia. Commercial games, which might be expected to be at the vanguard of AI development, are not. The reality is that the video game industry, naturally, holds the gamer experience as the highest priority, and all game design decisions are made in that context. Consequently, commercial game development has focused on the environment of the game: graphics and character *models*. The impact is that AI research by the video game industry is meagre and only a small portion of the computational resources in a game are allotted to AI. The relatively small market for serious games destined for governments and the military, means that many of these games are derivatives of commercial games with some adjustments to suit clients or in-house developments. Research for seri-

ous games is necessarily constrained by costs that, unlike commercial games, are unlikely to be offset by sales profits.

Academia remains the domain that is most actively pursuing innovative AI research; however, while considerable research is taking place it is not necessarily focused on applicability to games. Increased efforts have been made in the last decade to increase the interaction between academic AI researchers and industry AI developers. This rapprochement promises to bring accelerated improvements in the integration of advanced AI in games.

Generally, there is no clearly superior AI technique for enhancing computer generated character behaviour, though incremental progress has been made over the years. For the most part, commercial games provide an illusion of AI by: marginally enhancing very basic algorithms for decision making, cleverly timing character animation, and the use by the computer of game information unavailable to the player. This is a reflection that human intelligence is still not completely understood, and that existing approaches only model aspects of intelligence relevant to narrow problems such as path-finding. A new branch of AI, known as Artificial General Intelligence, has recently begun to focus research on a more generally applicable AI architecture that may provide the foundation for a believable game character AI.

Significance of results: This survey served to highlight that game AI is still an open problem. The situation presents both a challenge and an opportunity. On the one hand, there is no single software package that could be acquired to implement a clearly superior AI in a CGF. On the other hand there is tremendous opportunity to participate in progressing AI research. The recent increase in computational power afforded by multi-processor CPUs, means that more complex AI approaches can be explored in a game context. The latter, when combined with keen client interest and first-hand knowledge of military operations, means that DRDC is well-positioned to progress and lead research in the improvement of serious game AI.

Future work: Industry and government requirements for progress in the field mean that considerable intellectual interest and resources are available to tackle the problem. Future work will be the identification of the most immediately promising AI techniques and their application to an existing CGF in use within DND. The most promising research directions are: the inclusion of emotions and needs in behaviour modulation and more naturalistic implementations of basic behaviours such as path-finding and collision avoidance.

Sommaire

Artificial Intelligence in Games

M.A.J. Bourassa, L. Massey ; DRDC Ottawa TM 2012-084 ; R & D pour la défense Canada – Ottawa ; août 2012.

Contexte : Les jeux sur ordinateur remplissent un créneau de marché important et profitable. Ils sont utilisés à des fins de divertissement, ainsi que pour des applications sérieuses dans le domaine de la modélisation et de la simulation. De plus en plus, on considère les jeux sérieux comme une méthode rentable, et souvent plus sécuritaire, de former le personnel dans le cadre d'occupations exigeantes et dangereuses. Les jeux sérieux remplissent un rle croissant au sein du Ministère de la Défense nationale o ils sont utilisés pour la formation et le perfectionnement opérationnel.

Les jeux sérieux utilisés pour la formation opérationnelle doivent bien sûr être aussi réalistes que possible. La qualité du rendu sonore et visuel des jeux commerciaux et des jeux sérieux s'est considérablement améliorée. On peut affirmer que la qualité du graphisme de ces jeux permet de plonger les joueurs de manière convaincante dans un environnement synthétique. Par contre, la qualité du jeu dépend aussi grandement de la qualité du comportement des adversaires générés par ordinateur. En conséquence, les améliorations futures des jeux commerciaux et sérieux seront probablement fondées sur une crédibilité accrue des adversaires virtuels. Ces adversaires doivent agir de manière crédible, c'est-à-dire humaine.

Principaux résultats : Le présent document donne un aperçu de l'utilisation de techniques d'intelligence artificielle (IA) actuelles ou précédentes dans l'industrie du jeu pour améliorer la crédibilité des adversaires et des alliés virtuels. Les techniques décrites proviennent de l'informatique et des neurosciences modernes, ainsi que de dizaines d'années de recherches en sociobiologie et en psychologie. L'enquête examine l'utilisation de l'IA dans les jeux de divertissement commerciaux et les jeux sérieux, en particulier par rapport aux forces générées par ordinateur dans les jeux militaires (p. ex. : OneSAF, JSAF, VBS2).

La principale conclusion tirée est qu'il n'y a pas de chef de fil évident en matière d'application de technologies d'IA aux jeux, que ce soit dans l'industrie, le secteur public ou le milieu universitaire. On pourrait s'attendre à ce que les jeux commerciaux soient à l'avant-garde du développement de techniques d'IA, mais ce n'est pas le cas. L'industrie du jeu par ordinateur met naturellement l'accent sur l'expérience offerte au joueur et prend des décisions en conséquence en matière de conception des jeux. La conception des jeux commerciaux est donc axée sur l'environnement de jeu : graphisme et *modèles* de personnages. Par conséquent, l'industrie du jeu par ordinateur effectue peu de recherche en IA et seulement une petite partie des ressources informatiques de jeu servent à l'IA. En raison du marché relativement restreint pour les jeux sérieux à l'intention du secteur public et des

applications militaires, plusieurs de ces jeux sont dérivés des jeux commerciaux et modifiés en fonction des besoins des clients ou des développements à l'interne. Les recherches relatives aux jeux sérieux sont limitées par les coûts qui, à l'encontre des jeux commerciaux, ne seront probablement pas couverts par les profits.

C'est dans le milieu universitaire que sont effectués des travaux de recherche novateurs importants en matière d'IA. Par contre, ces travaux ne portent pas nécessaires sur l'application de ces techniques aux jeux. Au cours de la dernière décennie, on s'est efforcé d'accrotre les interactions entre les chercheurs universitaires en IA et les développeurs d'IA de l'industrie. Ce rapprochement devrait accélérer l'intégration des techniques de pointe en IA aux jeux.

En général, aucune technique d'IA ne se démarque par rapport à l'amélioration du comportement des personnages virtuels, bien qu'on ait réalisé des progrès graduels au fil des ans. Dans l'ensemble, les jeux commerciaux offrent l'apparence d'une intelligence artificielle au moyen d'une amélioration minime de simples algorithmes de prise de décision, d'un chronométrage astucieux de l'animation des personnages et de l'utilisation par le logiciel de renseignements de jeu non disponibles au joueur. Cela reflète une compréhension incomplète de l'intelligence humaine et indique que les démarches existantes ne modélisent que des aspects d'intelligence liés à des problèmes bien particuliers, par exemple, la découverte de parcours. Un nouveau domaine d'IA, l'intelligence artificielle générale, axe ses recherches sur une architecture d'IA d'application plus générale qui peut servir de fondement au développement d'une IA de personnage de jeu crédible.

Importance des résultats : Les résultats de l'enquête démontrent que l'application de l'IA aux jeux demeure un problème. Cette situation présente à la fois un défi et une possibilité. D'une part, aucun logiciel disponible ne permet l'élaboration au moyen d'une IA nettement supérieure de forces générées par ordinateur (CGF). D'autre part, cela présente une excellente occasion de participer aux recherches en cours en matière d'IA. L'accroissement récent de la capacité de traitement grâce aux multiprocesseurs permet d'explorer des démarches d'IA plus complexes dans un contexte de jeu. De telles démarches, ainsi qu'un réel intérêt de la part des clients et une connaissance directe des opérations militaires, mettent RDDC en bonne position pour mener et faire progresser les recherches sur l'amélioration de l'IA pour les jeux sérieux.

Travaux futurs : L'industrie et le secteur public ont besoin de progrès dans ce domaine - il y a donc un intérêt réel et des ressources sont disponibles pour attaquer le problème. Il faudra cerner les techniques d'IA les plus prometteuses et déterminer leur application à un système de CFG utilisé au MDN. Les avenues de recherche prometteuses comprennent l'inclusion d'émotions et de besoins dans la modulation du comportement et une application plus naturelle de comportements de base comme la découverte de parcours et l'évitement des collisions.

Table of contents

1 Introduction

There is a requirement for realistic computer simulations for use by the military, police and rescue teams of countries around the world. The use of computer simulations, or "serious games", to explore or prove concepts, or as part of a training system, is well established. The simulations have the advantage of replacing large quantities of equipment and large numbers of personnel by synthetically generated equivalents, thus potentially reducing costs and infrastructure requirements. An added benefit is that simulations can allow the creation of environments tailored to the needs at hand, for example, one not subject to inclement weather [1]. Finally, in the case of military training, it is recognized that experience and reflection on lessons learned are the best source of education. Simulations can provide such learning without the attendant risks that some situations may incur in real life [2]. Needless to say, the more closely a simulation mimics reality, the more worthwhile it may be. Consequently, every effort must be made to ensure the highest quality of a computer simulation from its depiction of the environment to the behaviour of the computer created entities that populate that environment.

This paper outlines the current state-of-the-art of the behavioural aspects of computer generated entities in computer simulations and games. To that end, the survey reviews the existing and most promising Artificial Intelligence (AI) technologies that could lead to improved Computer Generated Actors (CGA) in military simulations.

CGA are computer-controlled characters present in a simulated environment. They respond to both the virtual environment and human operators/players in a human-like manner. CGA are alternatively known as computer generated characters, Non-Player Characters (NPC), Computer-Generated Forces (CGF) and as Semi-Autonomous Forces (SAF). This paper will use the acronym CGA.

Although the paper's focus is on CGA at the individual level, CGA behaviour at higher echelons such as units, brigades, etc. will be considered. Composition issues, such as grouping or movement in formation [3], will also be considered when they illustrate interesting applications for the creation of realistic, intelligent behaviours.

CGAs have been used extensively in non-interactive computer modeling and simulation such as for Monte Carlo simulations in operational studies used to collect large, statistically significant outcomes. Of interest in this paper are CGAs that are used in 'Human-in-the-Loop' (HiL), interactive simulations. In HiL, a human is involved as a participant in the simulation while the computer generates CGA allies and opponents.

HiL interactive simulations that are used for training, aim to present a human participant (player) with a synthetic environment populated by CGAs that behave in a plausible, human-like manner. In order to evaluate the effectiveness of CGAs, it is necessary to define what exactly constitutes human-like behaviour. In this paper, the focus is necessarily on HiL since it is the area where unrealistic CGA behaviour is most likely to impact the user experience.

Important aspects of human-like behaviour are situational awareness and the ability to adapt to changing circumstances [4]. Behaviours are loosely gauged by the following criteria:

1. *Fidelity* – the accuracy of the representation of simulated characters during the simulation. A CGA should be as indistinguishable as possible from human-controlled actors [5]; and
2. *Believability* – the impression of reality a CGA will give a human participant.

The requirements for believability originated from the *Oz Project* at CMU, lead by Bates [6]. Believability has been recently studied in detail by Mateas [7], who lists the following aspects of a CGA that will result in believability:

1. consistent, lifelike behaviour;
2. illusion of intelligent life, including personality and emotions;
3. social interaction;
4. adaptation to circumstances, ability to change and evolve; and
5. self-motivation.

A wide range of AI techniques are, or could potentially be, used to give CGAs the aforementioned human-like abilities. In this paper, AI techniques that have the potential to provide such realistic, believable behaviours for CGAs are reviewed. The review is restricted to AI applications to CGA. For other simulation and CGA-related issues, such as software or hardware configuration, architectures, standards, inter-operability, integration, evaluation, etc., the reader is directed to the work of Bruzzone et al. [8]. Additionally, AI is not considered exclusively from the point of view of optimality, but instead from the "intelligence" of the resulting behaviours. This means that the behaviour "looks right" (believable) in the eyes of a human observer [9]. This view of AI forms the basis of more plausible simulation and will have a direct and positive impact on the quality of the training of human participants.

Finally, this paper only provides an overview of techniques and concepts to keep the text light and concise. The reader is encouraged to refer to the cited literature for deeper technical details, and also for references to previous work.

The following sections will first review the current state-of-the art of CGAs in modeling and simulation (M&S). It will be followed by sections on the following topics: autonomous agents, games, artificial life, complexity theory, and embodied AI. Many of the techniques used in different fields overlap and this overlap will reinforce the pre-dominance or potential of some ideas compared to others.

2 Modeling and Simulation (M&S)

Simulations often have analytical purposes such as planning and monitoring. In the military, one may be interested in simulating force structure, strategies or logistics with the motivation of making battleground predictions. These simulations work in accelerated time for cost and time efficiency reason, with scenarios repeated multiple times to accumulate data. The focus here is on human in the loop, interactive, real time simulations in which CGA fidelity and believability is of utmost importance. In this context, a CGA behaviour or action is activated in realtime based on some stimuli generated by the system or from external sources (such as a human user) and is used to control resources, weapons, or movements of troops or equipment. The goal of AI is to select the most appropriate action or adopt a behaviour in a way that is both realistic and human-like.

2.1 AI and Computer Generated Actors

Stytz and Banks [10] provided a series of papers giving a good overview of AI techniques used in CGA in the context of military simulations. The techniques outlined include the following:

1. *Blackboard systems* allow for the sharing of data among intelligent agents. Each agent acts as a knowledge source and places information on appropriate panels of the blackboard. Other agents can then access that knowledge to include in their decision making process. The blackboard structure contains the CGA state information and represents its situational awareness.
2. *Expert Systems* (ES) are knowledge bases (KB) (containing both rules and facts) elicited from human experts coupled with an inference engine to use that knowledge and to make decisions or recommendations.
3. *Case-based reasoning* (CB) works by matching a problem with an existing case in a case base (a database of previous situations). The case base includes solutions to past problems, and the CB system adapts those solutions to the current situation.
4. *Evolutionary computing* (EC) is an optimization technique that allows for searching a solution space using genetically inspired concepts. It is a form of learning that adapts to the environment based on a pre-defined fitness function.
5. *Bayesian* and *Bayesian Belief Networks* are based on Bayes' probabilistic theory and excel at representing and manipulating uncertain information.
6. *Fuzzy logic* systems are also useful to represent and reason about uncertain data by attributing a truth value that is not restricted to true or false but can take a continuous value between 0 and 1.
7. *Neural networks* (NN) are inspired by biological neurones and excel at recognizing patterns and learning with noisy or incomplete data.

The authors also describe Modular Semi-Autonomous Forces (ModSAF), a military software architecture for simulation and CGA creation. ModSAF is organized around tasks with each task implemented by a finite state machine. Related tasks can be grouped into *task frames* to allow for the composition of tasks of any complexity. A mission assigned to a CGA is a set of such task frames.

The same authors have also presented a three part paper surveying the state of AI applied to CGA [5, 11]. Although this series of papers was published in 2003, the most recent references in them date back to 1999 and most references are actually older than that. The information in the papers is therefore 10 to 15 years out of date. The main points of this survey are highlighted below to convey the essence of older work, without delving into the specific description of individual work. Note that the general conclusion and state of the field of M&S as it relates to intelligent CGA, have not changed very much since then.

Stytz's survey of AI applied to CGA drew the following conclusions:

1. CGA research has focused on reasoning system architecture, software architecture, and human behaviour modeling.
2. A multitude of reasoning systems and behaviour models has been proposed. The rest of these points will present an overview of them.
3. Reasoning systems have mostly been implemented with rules/production systems (such as expert systems) and task frames (as described for ModSAF previously).In some cases, decision trees (a decision or classification mechanism where each branch of a tree structure represents a possible outcome based on the value of an observation evaluated at the parent node) and case-based reasoning have been used, often in conjunction with task frames.
4. There is, in general, no learning taking place in CGA architectures, but the authors state it would be very desirable since it is not possible to code every eventuality in a knowledge base (or case base). This leads to current CGA lacking realism and being brittle when faced with unforeseen circumstances and incomplete knowledge or information.
5. No singular reasoning system or behaviour models has been identified as being superior to any other.

2.2 Impact of Architecture

As mentioned in the introduction, although architecture and implementation are very important when developing CGA software, the objective and focus of this paper is on the actual AI techniques and aspect of generating believable CGA. Nevertheless, it is worth briefly reviewing architectures and implementations to identify some relevant points expressed by Stytz [11]. While most of the citations in the latter document date to the late nineteen nineties, they remain relevant today. The following observations are the most pertinent to designing believable CGA:

1. Architecture and implementation issues dealing with CGA assembly from software component libraries or the structure of the KB, are not the most important issues. The most relevant ones are the conception of the expert systems and its interaction with task frames as well as the use of inheritance to specialize knowledge from frames containing general knowledge.
2. Most research has not addressed architectural issues. Software agents are noted as an architectural solution that is gaining popularity (see Section 3). Software layer architectures used with object orientated methodologies or components, was also becoming popular to manage the complexity of software. (Note: Both cases are even more evident today.)
3. In general, systems still have an explicit strong connection between the architecture and supported reasoning system. This makes supporting more than one reasoning system type difficult. Simulation standards such as HLA have helped in this respect by providing a service oriented architecture.
4. The authors conclude that there has been no standard or metric to compare performance and capabilities of various architectures. There is therefore little information available to guide the development of systems that might exploit the advantages of some techniques while avoiding pitfalls.

2.3 Knowledge Representation

As pointed out above, most reasoning and human behaviours are modeled with large knowledge bases implemented with rules. The major difficulty encountered in developing and extending such knowledge bases is the acquisition and hand-coding of all that knowledge in a usable format (representation). This is in fact a well known problem in AI since all intelligent behaviour require some form of domain or general knowledge to be hand-crafted [12, 13]. Gonzalez et al. [14] have proposed an interactive acquisition mechanism that elicits tactical knowledge for CGA from dialogues with a subject-matter expert. The partial model thus acquired reduces the hand-coding effort by 60% but the author admits that no rigorous comparison has been done on the effectiveness of the knowledge compared with totally hand crafted systems.

Another school of thought is that instead of coding large knowledge bases, complex models of human behaviours and deep reasoning systems, one might want to provide a CGA with a shallow intelligence that gives the illusion of intelligent behaviour. This can be achieved by equipping CGA with a rich inner-life that includes daily routine activities and purely reactive behaviour. This is sufficient to obtain believable behaviour from simpler CGAs:

crowds, supporting actors, etc. Riedl [2] and Mateas [7] used such an approach for leadership training simulations where the aim was to generate interactive and adaptable scenarios using purely reactive planning. Another solution to address the difficulty of acquiring knowledge is the idea of composable behaviours for CGA, where users compose custom behaviours for CGA from a repository of more basic behaviours [15, 16, 17]. This idea is currently in use in current serious games such as OneSAF.

Furthermore, conventional crisp logic, that is the bimodal logical statements used in the rule base of most CGA rule-based systems, is often inadequate to deal with the uncertainty arising from incomplete information in real world situations. To capture the approximate reasoning of human beings dealing with imperfect information, Sui et al. [18] implemented a CGA in which conventional logic rules are complemented by fuzzy logic rules. They have built applications in air-to-air combat and ground robot swarm simulation with their hybrid conventional-fuzzy logic rule base. The authors fail to discuss and compare the performance of this approach with a traditional rule-based system. Fuzzy rules, like crisp rules, do not scale well.

CGA often work in a purely reactive manner: do something until the situation changes, then do something else. In such cases, a behaviour model for CGA only aims at reproducing observable human behaviours. Deciding on an action is based merely on past actions and current environment state. This is best implemented as a finite state machine, which causes the CGA to stay in a given state until a state transition event occurs. In implementation, state machines implement decision-making with rules. Transitions can be implemented with simple conditional statements or with decision trees. Finite state machines are very predictable and do not scale well to large state-spaces.

Reactive behaviour that simply reproduces expected human behaviours when a state change is triggered may be appropriate when implementing a CGA based strictly on military doctrine. However, this approach will inevitably lack realism since warfare, even in a simulated environment, is rarely as clear-cut as what doctrine or operating procedures describe. Higher-level cognitive processes must be taken into account to make CGA more human-like and thus believable. One possible solution is to use *cognitive models*. Cognitive models attempt to model human goals and intentions while determining a plausible course of action that mimics human behaviour. They offer a much richer ”mental” environment to make decisions based on human-like cognitive abilities and processes. Several cognitive models are discussed in the next paragraphs.

The action/cognition behaviour model (ACBM) is one cognitive-type model described by Landweer [19] in which complex behaviours are composed from simple functions. For increased realism, CGA do not have direct access to the actual environment state but instead build their representation of the world based on their perceptual capabilities and memory. An example of a cognitive model developed for military decision-making is presented by Weaver [20]. In the latter work, decision-making is modeled with a neural network with three outputs that give the scale of the decision.

Busetta, Ronnquist, Hodgson, and Lucas [21] implemented a CGA based on the *BDI* (Belief-Desires-Intentions) cognitive model in which human behaviour is deemed to have three underlying factors: beliefs about the world, desires that need to be satisfied, and intentions of acting to satisfy these desires. In this context, a CGA performs some actions based on its intentions until it is forced to reevaluate its intentions by events in the simulation environment. The main difference between *BDI* systems and other models is the ability to represent and act on intentions. Beliefs corresponds to what a CGA can possibly do (its knowledge) while desires are a subset of beliefs describing what the CGA desires to do based on environmental stimulus. However, the choice of action is not determined solely by this subset: it can be further expanded or constrained by intentions. Indeed, in the *BDI* model, CGA choices may not be taken from the subset of desires and may even be beyond its beliefs. Therefore, an intention (the chosen set of actions) could fail to achieve a goal. In fact, intentions can be changed after a series of actions have been taken towards a goal. Although not reaching a goal may seem like an undesired property when considering optimality and effectiveness issues, in reality it represents a typical trait of human-like behaviour that can be useful to implement plausible CGA. There are many more such models, some addressing specialized aspects of behaviour, such as *INCIDER* [22] which looks at modeling human factors in friendly fire casualties. Such factors include stress, risk taking attitude, experience and confirmatory bias.

Two well-known general-purpose approaches in modeling cognitive aspects of complex intelligent behaviour are the *Soar* [23] and the *ACT-R* architectures [24]. Many CGA have been developed based on these architectures [5]. For instance, *ACT-R* was applied to the modeling of a pilot landing with synthetic vision system [25]. As well, Best et al. [26] investigated simulations for Military Operations on Urban Terrain (MOUT), in which *ACT-R* is used to improve the cognitive realism of CGA. The simulation takes place as part of a video game, where the attackers (human players) face a team of CGA defenders. Realistic team behaviour is frequently lacking in CGA, but in the case of this work, the authors report good responsiveness between both human and CGA teams. The objective with *ACT-R* was to obtain more realistic motor and perceptual behaviour, including space representation and navigation. Realistic behaviour would include for instance memory decay and shortcomings, as well as fatigue and stress. The authors have implemented their CGA as *ACT-R* based autonomous agents. The concept of agents will be reviewed in detail in Section 3. Of note is that their ideas have been implemented in software and in robot, in both cases using *ACT-R*. Other work on teaming behaviour using an extended version of *ACT-R* is being done by Kennedy [27], in which the idea is for the CGA (a robot in this case) to be aware of its teammates' capabilities so that they can better cooperate towards a task. The application tested is alarm response by groups of robots and humans.

With cognitive and behaviour models, one once again faces the challenge of acquiring large amount of specialized domain and general-purpose knowledge, and then translating it into a usable representation. While such models can lead to more evolved behaviours, they suffer from similar weaknesses as rule-based systems: incompleteness of knowledge and imperfect reasoning systems. Reece [28] summarizes the situation as follows:

> …complex but brittle AI mechanisms for inference, planning, reasoning and decision-making are often used that provide superhuman capabilities in some instances and catastrophic failures in others, without any of the limitations and resourcefulness of human cognition.

For these reasons, and because constructing models can be very time consuming and costly, one may be interested in machine learning techniques to acquire the rules governing behaviours. Machine learning will be addressed in Section 4 of the paper but an example of learning in the field of M&S is presented here. Although not applied to the specific area of generating believable CGA for interactive simulations but rather for predictive battlespace simulation, the work of Wedgwood et al [29] looks at using evolutionary algorithms to determine which rules each agent should use in the context of UAV coverage. UAV behaviour is evolved using a global fitness function that is based on the performance of all UAVs working together to cover a given region. The solution is also guided by inputs from the user specifying, for example, that two UAV sensor coverages should not overlap. The authors report good solutions with relatively little human guidance. Of further interest, the multiple UAVs interact among themselves and with the environment leading to the emergence of novel behaviour. Furthermore, this type of behaviour by UAVs is similar to what human pilots might do, hence displaying realistic and believable actions.

Agent and multi-agent AI technology, which is referenced several times in this paper, is becoming increasingly popular for intelligent CGA implementations. A chapter is dedicated to this topic later. Agents are useful in the CGA context because they provide a natural metaphor for decentralized decision-making and control and also because they offer an open and flexible architecture for expansions and adaptation to unpredictable battlefields [30].

In summary, Stytz [5] made the following observation about the state of the art in CGA modeling:

> In general, researchers advocate that a model of human behaviour should account for attention, intent, situation assessment, decision-making, skills, training, fatigue, environmental stressors, education, experience, and motivations, and do so in some combination. To date, very few systems have attempted to model more than a few of these factors, and most models are generally assembled using an ad hoc approach instead of employing a strong methodology. Early efforts at human behaviour modeling used finite state machines; recently, these efforts have been coupled with agents to try to improve the fidelity of the

> modeled behaviours. Concurrent control and cognition in conjunction with arbitration as well as case-based reasoning coupled with machine learning have also been investigated. Intelligent agents have also been used to model behaviour by having one agent responsible for each type of behaviour and then arbitrating or scaling their individual outputs.

The situation, as represented by Stytz, has changed little since 2003. There is an increased focus on agent-technology and on machine learning, but many agents are still implemented as simple state machines. Where such agents implement no adaptive AI, they serve as nothing more than a vocabulary shift in the field of CGA, bringing no real advance in developing agents that learn from the environment and from the users. The conclusion is that much work remains to be done to achieve believable CGA.

3 Autonomous Agents

Agent technology is about producing intelligent, autonomous entities. The entities can be software, or embedded in a hardware platform such as a robot. An agent-based system consists of a bottom-up design, where the overall system functions as the sum, sometimes synergistic, of individual agents behaviours. By comparison, non-agent AI systems take a top-down approach where the whole system functionality is pre-specified and individual behaviours are generated and controlled centrally.

Autonomous agents (AA) are agents that rely on their own sensors to determine a course of action or a behaviour in an unpredictable environment [31, 32]. Franklin and Graesser [33] identify four properties that would qualify a CGA as an agent:

1. Reactive - reacts to game-environment situations and translates these percepts into appropriate action.
2. autonomous - support its own life, makes its own decisions i.e. not hard-coded behaviours.
3. goal-oriented - can achieve a given mission considers actions to take, consequences and ability to reach goal state.
4. continuous time - continuously sensing the environment and acting (vs discrete time, turn by turn actions or plays).

Autonomous agents can be framed more generally within the context of Autonomic Computing [34]. Autonomic systems are self-regulating, self-managed, self-fixing systems and as such, form a super-class under which autonomous agents can be deemed to fall. This offers a general framework for agent comparison.

Just as with the field of M&S CGAs, there is a desire in agent research for believable behaviour [35, 36]. Although agents are closely related to M&S CGAs, AA generally have a specific purpose while in S&M, CGAs must behave in a much wider variety of situations related to the whole range of real-world conditions present in the simulation environment [5]. Each AA application can relate to a specific desired behaviour. For instance, navigating robot agents can be related to the movement of troops in a simulation. Hence, any autonomous software or hardware agent displaying human-like intelligence has the potential to find fruitful application in military CGA. Autonomous agents can also be used to populate virtual simulation worlds [37].

Multiagent systems (MAS) are systems in which many interacting autonomous intelligent agents pursue goals or perform tasks [38]. MAS can be useful in the modeling of groups of individuals where each individual possesses its own variable behaviour. MAS can also be seen as a distributed computing architecture. Such systems can see the emergence of complex behaviours from relatively simple ones. This latter property makes MAS a good choice for CGA implementation to, for example, model the decision process of a commander [39]. However, problems that are hard with single entities such as collision-free path finding, become harder with MAS and grow combinatorially with the number of agents to become intractable [40].

3.1 Rules

Autonomous agents often use hard-coded rules rather than using techniques from cognitive theories and autonomous cognitive science [41]. As such, agents are frequently implemented as production rule systems and finite state machines with periodic sensing of the environment. Such agents frequently suffer from a lack of realism due to the hard-coded nature of their behaviours and from a combinatorial explosion as the number of states increases. They are inherently brittle since rules are prone to be incomplete representations of what can happen and of the circumstances under which events can occur [42]. Thus, rule-based agents' performance can vary broadly depending on the agent designer's domain knowledge and their ability to anticipate potential situations the agents may encounter.

Rule-based autonomous agents lack 'situational awareness', which leads to unrealistic or even totally unacceptable behaviours. Situational awareness is central to intelligent CGA since it drives the decision making process and consequently what actions a CGA will take. Realistic situational awareness by human-like agents is affected by a multitude of factors like weather, culture, training and fatigue. Human-like agents must necessarily account for such factors and, like humans, must also be able to select between doctrinal and non-doctrinal actions depending on the situation [30].

3.2 Cognitive Models

Agents have been created that exploit cognitive and behaviour models identical, or similar to, those covered in the previous section [42]. These agents react to the environment with a high degree of human-like realism, including moderating factors such as stress and fatigue. Fletcher [30] describes such a general approach to agent-based CGA that display cognitively plausible, variable behaviour at the individual levels and at larger military groups levels such as coalitions. Wang [41] describes a layered architecture and theoretical framework for autonomous agents that draw from cognitive informatics and computational intelligence. The basic idea behind Wang's work is that intelligent behaviours in autonomous agents develop hierarchically at three different layers: imperative (similar to instinct, rule-based behaviours), autonomic (self-managed, unconscious) and autonomous (higher intelligence, for instance contributed by mammals neo-cortex).

An alternative to cognitive models is the use of ontologies [1] to model conceptual relation-

1. An *ontology* is an explicit formal specification of the terms in a domain and the relations among them. [43]

ships. As such, an ontology embeds a priori knowledge implicitly in a vocabulary. Ontologies are an increasingly popular concept in Knowledge Management [44] and their use has been investigated for making agent-agent and agent-world social interactions more realistic [45].

Modeling a sound, yet realistic, decision making process is an important goal for human-like CGA behaviour. Naturalistic decisions are the type of decisions people make based on their experience, and such decisions may be neither optimal, totally rational, or unbiased [46]. Sokolowski [39] developed a multi-agent model that reproduced the naturalistic decision making process of a senior military commander in the field. The model was validated by comparing model decisions to those made by actual military officers. The model was shown to generate non-optimal decisions similar to those of the humans. The underlying agent's experience and actions, elicited from case studies and human experts, were encoded in *frames* and evaluated using fuzzy logic. The frames used differed from the group of related tasks mentioned in Section 2. In Sokolwoski's work, a frame is an AI data structure used for representing a stereotyped situation [47]. The authors noted similarities between their work and the *BDI* cognitive model presented in the Chapter 2. Interestingly, the fuzzy model used can adjust its goals in order to resolve conflicts. The multi-agent architecture used was chosen over other AI implementations such as finite state machines, rule-based systems, case-based reasoning and neural networks because it scaled well to large knowledge domains, could explain and represent the decision process, and because its goal orientation and interaction allowed the emergence of complex behaviours from simple ones.

3.3 Emotions

Realistic human behaviour modeling should include emotions. For example, fear and excitement have an impact on memory, judgment and decision making. Aguirre et al [48] proposed a model based on emotional intelligence to represent mood and emotion and to control the behaviour of avatars (virtual creatures) in 3D virtual environment. The idea of emotional intelligence comes from the work of Damasio [49] and of Goleman [50] and involves self-awareness of one's own emotions and of others'. Aguirre's work focussed on building an ontology of emotions and of their facial representations, and how avatars express their emotions based on emotional states that are influenced by the environment. They report realistic avatar behaviour and motion but stress that the work is very preliminary and much more investigation is required. Most models of emotions used in autonomous agents have focussed on either the physiological or the cognitive aspect of emotion. There is a need to balance both aspects to obtain believable emotional behaviours [51].

Cattinelli et al [52] report on a simple model, based on probabilistic finite state automata, that displays rich *emotional* interactions between a robot and a human user or other robot. Emotional states were predefined and the paper modeled the *transitions* from one emotional state to another. The model introduced probabilistic and time-varying state machine transition functions that resulted in "rich and dynamic interactions" between agents as they adapted to each other's emotional state. The model also used both manual tuning and *reinforcement learning* (RL), to successfully guide the CGAs towards desired interaction goals. The latter demonstrated the potential of hybrid approaches to achieve interesting results.

Bosse et al. [53] also looked at emotion regulation in agents, that is, how agents deal with and manage emotions. The premise was that agents must be able to reason about the mental state of other agents and adjust their actions accordingly, for instance to anticipate other agents actions or to manipulate their behaviour. In a military simulation this might involve a CGA increasing its level of aggressiveness when it determines that a human trainee shows signs of lower morale or lack of desire to fight. The model proposed was an extension of the *BDI* model that included emotions modeled with rules that took into account the duration of a triggering event required to generate an emotional response of a certain level and duration. Emotional levels in turn triggered rules governing actions. Emotions were seen as shortcuts to actions, bypassing rationality. It was observed that the model's rules had to be handcrafted. As such, the emotional modeling suffered from the same knowledge acquisition bottleneck discussed earlier. However, the approach proposed by Bosse et al. has the advantage of being adaptive in that an agent observes the environment and other agents and adjusts its behaviour accordingly. Bosse et al. tested their approach in an adventure game in which a guide had to adjust its behaviour based on the level of fear of the traveler under his responsibility. The simulation proved successful because the guide, over time, became better at predicting and managing the level of fear of the traveler. Although the test displayed some realistic affective behaviour, the implementation is limited in two ways. First, it does not interact with human users and so is not able to react to the user's emotions. Second, it is not able to elicit emotions automatically from environmental cue, that is, the model encodes a fixed number of emotions cued only by certain events.

3.4 Learning

As mentioned earlier, cognitive model based agents and similar architectures require large amount of knowledge handcrafting and still do not fare well in complex, non-linear or uncertain situations [54]. A more adaptable autonomous agent would learn from the environment rather than expect everything to be predetermined by a rule set. The learning process involves: observation, situational awareness, coordination, reasoning and decision-making. For instance, Camponogara [55] used an autonomous agent that attempted to learn to control traffic in a city. In the latter instance, the agent was intended to display some of the characteristics of policemen or of traffic control centre staff. One recognizes in this example that the subjects have traits in common with those of a military leader. Therefore, the AI techniques and lessons learned from traffic agents would be applicable to the creation of military CGA. In this work, reinforced learning was used to train collaborative agents to take the correct actions in directing traffic. This idea of RL is indeed very attractive in applications where a reward or penalty from the environment is readily available.

Learning can be very difficult and may require much domain knowledge to guide or restrict an otherwise intractable search space. For instance, Klugl et al [56] report on an experiment in which the agent's rule-base was replaced by a classifier system to learn the required behaviour. Their results show brittle action and perception modeling by the learning system. Shoham et al. [57] claim that the problem of reinforcement learning for multi-agent environments is ill-defined. Reinforcement is really about how best to select the next action based on observations of the current situation. The issue the authors raise is, how to define "best". They conclude that there are four possible approaches:

1. Descriptive - this approach sets criteria based on how humans learn;
2. Distributed - agents act independently and learning converges to optimal individual solutions;
3. Equilibrium - drawing from game theory, some form of equilibrium, such as Nash, defines successful learning; and
4. AI Agenda - each agent seeks the best learning strategy given the other agents present in the environment.

The authors reject the focus on equilibrium because of its lack of prescriptive power; equilibrium describes when learning should stop, but that condition can be weakened in the presence of multiple equilibria or when the strategy space is so large that convergence is unlikely in the time of the game. The authors also reject a descriptive view due to the lack of objective evaluation criteria; there still exists no definitive theory of how humans learn. The distributed approach is condemned because optimality can result in CGA that act too well to be believable. [2] Finally, the authors endorse the "AI agenda" with its focus on measuring learning success relative to the type of agents. This is substantiated by the observation that the actions (past, present, and future) and learning styles of each agent in a multi-agent environment, necessarily inform the learning and actions of any other agent it interacts with.

In another approach to AA learning, Deutsch et al. [58] investigated agent learning through memorization of past decisions and their outcomes as a means of providing feedback to new, similar situations. This is similar to case-based reasoning except that the case base is built incrementally as the agent accumulates experience. As well, Heuvelink [59] proposed a simulation training system for the Netherland Armed Forces with a traditional rule-based expert system augmented with a beliefs base accumulated over time. The former allows for rational reasoning while the latter offers biased reasoning making the agent behaviour more human-like. The approach is based on an episodic memory model with a decaying

2. The latter is a significant difficulty faced in the commercial game industry. A computer opponent that always acts perfectly creates a poor game experience for a human player.

component. The decay of memory episodes ensures beliefs become less available over time unless new evidence (episodes) continues to support them. The reasoning on beliefs is biased by this availability (evidence and decay level) just like a tired person may make a decision based on the first, non-optimal solution that comes to mind. This has the potential to make more believable CGA. Several issues are still to be resolved such as speed of reasoning with complex beliefs and consistency of beliefs.

An interesting idea proposed in Kechagias et al. [60] is the application of *data mining* [61] techniques to historical data (for example, from previous human-human or human-computer training sessions in a simulated environment) to build knowledge models for intelligent agents. Indeed, every interaction on a computer is a learning opportunity. The agent can thus adapt its behaviour to human players and, over time, gain in experience and intelligence. Typically, with the current technology, the opposite is true; the human trainee learns about inflexible agent strategies and soon learns to defeat the agents. This negatively affects the effectiveness of the simulation as a training tool, and for the human trainee, the simulation environment rapidly loses its credibility.

The learning approach suggested by Kechagias [60] is *off-line learning*, but many *on-line* learning algorithms exist that could permit a CGA to learn during simulation. In both on- and off-line learning, an important issue is to find the right time to learn since learning can be disruptive computationally [62]. The Centre for Research and Technology Hellas (CERTH) research group involved in the work of Kechagias et al. makes available the *Agent Academy*, an open source multi-agent development environment that includes training and re-training of agents using the *Weka Machine Learning Toolkit* [63] and other algorithms. The main idea is to couple agents with data mining techniques so that they can discover new rules as data is accumulated. This could, for example, be used to learn new engagement rules when enough observations of enemy behaviour have been obtained. The software package also provides an evaluation scheme for a given agent that observes *other* agents and evaluates them. The learning proposed in the work by Kechagias was tested in a legacy supply chain management system where agents coordinated various activities and customer requests. The authors report that their system successfully adapted existing business rules autonomously. Note that since the extracted rules were not always satisfactory, they were presented as “suggestions” to a user who then made the proper recommendations to ensure proper service to customers.

A potentially useful characteristic for artificial agents would be an ability to plan and this has been the subject of intense research, mostly under the form of rule-based expert systems and cognitive modeling. The need for planning in serious games is not a given. For example, one could program an agent to use the standard operating procedures (SOPs) a human might use in a given scenario. This is often the case in the military where SOPs and doctrines are defined for a wide range of tasks and operations. However, the world, and military operations particularly, rarely evolve as expected. Imagination and adaptation become important, as well as an ability to plan under uncertain and changing situations. An example of research in the area of reasoning and planning under incomplete information is Giannikis and Daskalopulu [64]. In their paper, agents in the domain of commerce are developed with an ability to use logic to make a best guess when faced with knowledge gaps. The authors investigate how default rules can be obtained incrementally without resorting to proving what information cannot be determined from the agent knowledge base. This work is only at the prototype phase but might be useful in the context of generating default, yet coherent, actions given incomplete information. The proposed approach adds some flexibility to the otherwise rigid logical rule paradigm of agent implementations. However, the representation is based on bimodal logic, a school of thought common in AI for many years but being increasingly questioned as unreasonably formal and unsuited to modeling realistic, human-like behaviour [65].

Autonomous agents have provided an alternative paradigm for rethinking CGAs but progress has been incremental. In large measure, the work with agents seems simply to have been a reapplication of the earlier algorithms of modeling and simulation, applied to a multitude of interacting agents. As a consequence, agent approaches have not yielded much improvement in CGA human-like behaviour. In contrast to modeling and simulation CGAs, where a single computer interacted with the world, agent methodologies have been successfully used to explore interactions between multiple CGAs. Overall, autonomous agent concepts have reinvigorated research in the field of AI, particularly since the approach lends itself admirably to the needs of a highly profitable industry, that of commercial video gaming.

4 Commercial Video Games

4.1 Overview

This section reviews the use of AI in commercial video games. The advent of so-called serious games [59, 26] demonstrated the applicability of games to the military. In fact, Fong [66] claims that both the game industry and the military simulation community, share common interests and goals, for instance, high-fidelity and plausible CGA. Fong believes the two will eventually merge. One can now find API and development environments geared specifically for training simulations using a game approach [67].

There are several video game genres including: Adventure, Role-Playing (RPG), Simulation, Action, and Strategy. Adventure games are typically oriented towards quests and puzzles with *Myst* [3] being the foremost example. Role-playing games originated with the pre-video era game *Dungeons and Dragons* [4], and are characterized by the players assuming roles in an extended story-line. There is generally little AI involved in these game genres.

Simulation games attempt to simulate, to some degree, one or more aspects of real life. *The Sims* (human, social simulation), *SimLife* (animal/biological simulation) *SimCity* (a city-building game), *Creatures*[69], and *Tamagotchi* (small, hand-held, simulated electronic "lifeforms") typify Simulation games. This game genre does include AI since the player assumes the role of a "god" overseeing the behaviour of electronic lifeforms. The latter have an internal AI which the player attempts to understand and influence towards some goal.

Neither of the preceding genres may seem immediately applicable for military uses as their emphasis is general not on confrontation. However, they do have an ability to simulate non-confrontational processes. The United States Air Force Research Laboratory (AFRL) recently commissioned the development of a game in support of its Intelligence Advanced Research Projects Activity. Together with the University of Arizona, the AFRL will develop an educational video game or serious game to train intelligence analysts and measure their proficiency in recognizing and mitigating the cognitive biases that affect intelligence analysis. [5]

Action games are typified by games found in arcades and the focus is generally on some form of confrontation. A sub-genre of these games are Shooter games. First-Person Shooter (FPS) games typically involve the player assuming the role of a combatant with projectile weapons, executing a mission either solo or as a member of a team. These games involve

3. *Myst* took the player on a first-person journey through an interactive world. The player's goal was to unravel the game's back-story by progressing through the world guided by clues. The graphics quality of the game led many to suggest it was as much a work of art, as it was a game.[68]

4. http://www.wizards.com/dnd/

5. http://uanews.org/node/42719

situational awareness, rapid decision making, quick reflexes, tactical planning, and, in team settings, good communication skills. Not surprisingly, these games have been used in military training and recruiting. Games such as *Halo* (science fiction, combat scenario) [6], *Unreal Tournament* (armed players fight in an arena) [7], *Call of Duty* (WWII and modern military combat) [8], and *Ghost Recon* (future World War) [9] are in the first-person shooter genre. While these games often pit player against player, the games are also populated by CGAs that are referred to as Non-Player Characters (NPC) or "bots". Bots usually have the role of advancing game play along a particular story line and are where the video game industry generally applies AI techniques.

Strategy games are an extension of the board game paradigm. Turn-based games include *Chess*, *Go*, and *Risk* [10]. Real-time Strategy (RTS) games typically have the player and an opponent moving simultaneously in real time. Typically, each player directs the overall strategy of bots under their control by assigning goals and, for some games, rudimentary tactics. The bots are typically semi-autonomous and following a command, can perform rudimentary pathfinding and combat. The genre has exemplars for most time periods and settings. Representative RTS games are *Starcraft* (science fiction), *Age of Empires* (several historical periods) [11], Naval War Arctic Circle (near future) [12], and *Command and Conquer*(historical and near future) [13]. Some games such as *Shogun 2* (Feudal Japan) [14] combine real-time and turn-based gaming.

While each of the preceding game genres have some measure of AI, considerable industry development is focused primarily on bot AI for RTS and FPS games. The remainder of this section will focus on these two genres which are the most relevant to CGF AI improvement. In 2004, Michael Buro of the University of Alberta wrote a call for AI research in RTS games.[70] He described the state of AI performance in RTS games at the time as poor and the section will conclude with an assessment of the current state of the art in FPS and RTS games from an industry perspective.

4.2 General Concepts

Millington [71] presents a model of game AI that divides it into three tasks in two domains:

1. **Group AI** - strategy, that is, the coordination of NPCs towards a common goal;

---

6. http://www.halo.xbox.com/en-us/
7. http://www.unrealtournament.com/us/main.html
8. http://www.callofduty.com/
9. http://ghostrecon.us.ubi.com/
10. http://www.hasbro.com/risk/
11. http://www.ageofempires3.com/
12. http://www.paradoxplaza.com/games/naval-war-arctic-circle
13. http://www.commandandconquer.com/
14. http://www.totalwar.com/shogun2

2. **Character AI** - the AI of individual NPCs which includes:
   (a) CGA decision-making; and
   (b) CGA movement.

Millington states that most action games use only Character AI. Decision-making is deciding what to do next. The movement, and associated animation modules, then execute the chosen behaviour. However, some games like *Half-Life* [15] and *Ghost Recon* also use Group strategy AI to coordinate teams of individuals. Millington describes various game AI techniques in detail, including pseudocode, advantages and disadvantages, and implementation issues. The reader is referred to the book for details and what follows is an overview of AI techniques used in games.

4.3 Decision Making Approaches

Heuristics (rules of thumb) can be used in decision-making to make a decision that is not optimal, but that represents a compromise between effectiveness and efficiency. Heuristics may not work at all in some situations but may work quite well in others. Additionally, heuristics can result in more realistic behaviour since they often mimic the way humans work out solutions.

Decision trees, or binary decision trees, are used for making decisions based on a tree structure, where each node of the tree represents a decision point (a condition that decides which branch will be taken next), starting from a root node and going down to the decision. Random decisions trees include random branch selection for added realism.

CGA in games often work in this manner: do something until something changes. This is best implemented as a finite state machine (FSM), which causes the CGA to stay in a given state until a state transition event occurs. Finite state machines also implement decision-making, in fact transitions can be implemented with decision trees. Finite state machines are very predictable but do not scale well to large state-spaces. A means of mitigating the scaling problem is to use hierarchical finite state machines (HFSM).[72, 73]

Rules, decision trees, and FSMs can lead to rigid and predictable behaviour. This is because conditions in finite state machines and decision trees are binary, that is, based on true or false values. Fuzzy logic allows for grey-scale values between true and false (as numeric values in [0,1]) and can be used to model uncertain situations or yield variances in decisions. Fuzzy state machines are an example of such an implementation.[74]

The above techniques deal with reactive-situations: IF stimulus/state, THEN select an action, behaviour, or decision. Sometimes, CGA must display goal-oriented behaviours, that

15. http://www.valvesoftware.com/games/hl2.html

is, the appearance of desires to achieve some objective. For this, a CGA must plan. Planning has to do with action selection to reach a goal, but the actions are situation dependent. This means that some actions may be disabled by a previous action or may be otherwise unavailable at some time. Planning by looking ahead and anticipating the outcomes of actions is then required to choose the best course of action. Games use rules-based systems and scripting for goal-directed behaviours. Scripting is when a pre-determined action is executed in response to a stimulus. The scripts are small sub-programs executed by the main program in response to events in a game or to a particular game state. Scripting does not scale well and considerable design forethought is required when using this approach.

Yildrim and Stene [75] examine the various types of game genres and identified those for which CGA are required. They conclude that in Role Playing Games (RPG), there is generally little intelligent CGA required. FPS games require the generation of both friendly, neutral (crowds), and enemy CGA that behave intelligently. RTS, in which both humans and synthetic players compete for resources to achieve goals in a real-time, also require the presence of CGAs. In Simulation games, the human player controls and oversees groups of computer generated characters and whole civilizations. Some limited use of intelligent CGA has taken place in this context to provide basic intelligence to the characters being controlled. Video interpretations of classic games, such as board games and puzzles, do not require CGA.

Having identified what games require CGA's, Yildrim and Stene [75] investigate whether or not CGA, in the context of games, are agents. To frame their answer, they provide a rapid and superficial survey of the AI techniques used in current and past games. The list is a familiar one: decision trees, deterministic and non-deterministic finite state machines, flocking and emergence, neural networks, genetic algorithms and fuzzy logic. The authors mention that most behaviour in game AI is hard coded and thus pre-determined, albeit at times with some randomness to display variable behaviour [76]. The authors also touch on a common industry approach to simulating AI, that of 'cheating'. These are techniques which give the illusion of intelligence. The resemblance is actually achieved by giving the CGA an advantage over the human player such as: faster reflexes, access to the player state, or a view of the whole game world. Although such an approach maximizes the NPC's capability it can also occasionally result in implausible actions. Autonomy and goal oriented behaviour are important features of being a true intelligent agent. Intelligence is then defined as the ability to acquire and effectively apply knowledge. However, the authors observe that learning rarely occurs in commercial games even though it is an important aspect in achieving adaptability and autonomy. A possible reason identified, is that game characters are required to make real-time decisions and therefore it does not seem practical to implement time-consuming learning mechanisms. The overall conclusion of the paper is that some game AI characters are often far from qualifying as intelligent agents.

The preceding gives a good idea of the state-of-the-art in commercial games and this will be revisited later in the section. It is also interesting to note that the main limiting factor identified by the authors was processing power. If this were true, then in the long run, true intelligence and learning can be expected to be more and more present in commercial games. In the context of military simulations, processing power can often be procured given a sufficient budget, so this may not be an issue. This still raises the important issue of the time complexity and computability of AI algorithms, topics that are studied actively in the academic community.

Schaeffer et al [76] present an accurate picture of the state of the art in game AI using the game *F.E.A.R.* [16] and work conducted at the University of Alberta to illustrate the current situation. They identify entertainment as a measure of success rather than performance. They stress the importance of making CGA act realistically, plan their actions and navigation autonomously and learn from both their mistakes and from the user. The authors state that learning "on the fly" is very rare in games but would go a long way in improving the realism and enjoyment of the game. These important issues will be revisited later in this section. An interesting game mentioned is the experimental game *PaSSAGE* [17]. *PaSSAGE* dynamically adapts its CGA as the game progresses based on what it observes of the human player's game-play. The result is a dynamic environment unique to each player's playing style and experience. The authors explain that the game environment will study a player and maintain a history of the game states (for global consistency reasons) so that, for instance, a player who has shown an interest in fighting will be presented more opportunities to fight than a player who has not.

Games also use model-based reasoning, as is often the norm in interactive simulations. Furthermore, models can be augmented with a learning capability. Ulam et al [77] have investigated a hybrid system that uses a model to reason about blame attribution when it fails in the context of reinforced learning. They claim that the agent developed in this manner learns to perform a task as well as an agent designed by an expert.

Real-time strategy (RTS) games such as *Starcraft* provide realistic insights for military training simulation. Contrary to classic games, they involve a large number of entities and resources in a complex environment. They also work in continuous time, and each player makes moves simultaneously. Such environments require rapid planning and decisions by players. CGA, whether opponents or friendly autonomous synthetic agents, in the game must therefore also display these abilities.

Traditional planning approaches, such as those in classic games like chess, are based on exhaustive game tree searches. Those approaches do not scale for use in complex RTS environments. RTS games also differ from classic games in that player decisions are not global in their impact. Sailer et al [78] propose an approach to planning that is based on high-level abstractions dealing with strategies rather than low-level, tactical actions. The strategy se-

16. *F.E.A.R.* stands for **F**irst **E**ncounter **A**ssault **R**econ. The game's official site is http://www.whatisfear.com/
17. http://hcsoftware.sourceforge.net/passage/

lection algorithm is based on Nash-equilibrium approximation. In a Nash-equilibrium setting, no player has an incentive to deviate from a given strategy. For optimal results, strategy selection has to be randomized. The authors have demonstrated their approach with eight strategies and have shown that a Nash-optimal strategy selection produces agents that defeat individual strategies. The approach has not been tested against human players or complex rule-based game AI agents but is an example of academic approaches to gaming.

The ideas of agents and planning as discussed previously in Section 3 resurface here, but with video games it is usually in the specific context of adversarial planning and opponent agents. In such environments, agents aim at defeating a human player. This is certainly very applicable to the area of CGA in military simulations. However, for video games, due to the interactive nature of games and the computationally intensive nature of planning, it must generally be completed off-line [79]. More generally, one may be interested in modeling not only opponents, but also allied or neutral players, and in all these cases, players can be described as embedded agents with various levels of autonomy. Two main approaches exist to designing a game agent. The first is traditional AI where game agent behaviours are governed by knowledge-engineered rules obtained from subject-matter experts. The alternative to this lengthy hand-coding process is to derive game policies, tactics and strategies by learning them. For instance, Miikkulainen et al [80] have investigated neuroevolution, the idea that neural network controllers for agent behaviours can be evolved using genetically inspired concepts aimed at maximizing a fitness function. They have shown that game agents learning this way can display flexible behaviour, adapting to new environments. Although neuro-evolution by itself is interesting to evolve a single behaviour, what is more important in a realistic scenario is the learning of multiple behaviours using multi-objective fitness functions that reward different behaviours so that a variety of behaviours applicable to various situation can arise [81].

4.4 Learning

Learning can derive a model that is purely reactive (a direct mapping of current state to action) or case-based (matching the tuple $< state, context >$ to an action). In some cases, continuous mapping of state to action would be required for situations similar to war or adventure games, but discrete mapping can also be useful as they present military strategy as movements of troops, similarly to games like chess.

The work of Lucas [82, 83] focuses on Temporal Difference Learning (TDL) [84]. TDL is a form of reinforcement learning [85] that involves optimizing a reward. Reinforcement learning is a popular form of learning in game AI. In artificial intelligence, a reward is a numerical signal received from the environment. Lucas investigated architectural issues such as action selection vs value function. With action selection, the learning agent observes the current state of the game and selects the next action with no look-ahead. The main advantage is the reduction in processing demands since there is no need for an internal game model and training is done off-line. However, off-line training itself can be long and unpredictable; some form of filtering is often needed to eliminate illegal actions. The value function approach, on the other hand, looks at the hypothesis space to select an action that generates the best possible future outcome. To achieve this, a model of the game needs to be hand coded.

Lucas also compared TDL and evolutionary algorithms for learning to adapt the game-play, based on interactions with the user. TDL has the advantage that it can learn incrementally at each time step, without a model of the environment. Evolution, by comparison, requires a population of strategies that are evaluated on a game, and only at the end of each game can the fittest one be selected. The approach is very wasteful since learning is slower and computation time is greater. Evolutionary algorithms do have the advantage of generally being more robust.

Finally, Lucas looked at various function approximation schemes for the mapping of state to action. This third aspect considers learning internal to the game rather than externalized to adapt to a user. Lucas concluded that for this aspect of learning, the best approach is evolved Multi-Layered Perceptrons (MLP) or neural networks. He also points out that each choice of architecture, learning algorithm and function approximation needs careful consideration. For example, in some cases, table-based function approximation often learns more reliable solutions than MLP, particularly when coupled with TDL. Lucas takes an information theoretic view to determine the best learning approach. Within that framework, he concludes that TDL should be used for learning things that take lots of bits to encode (complex behaviours) because evolution is too wasteful from an information point of view. However, the question of how this information theoretic view translates into reality is still open. Indeed, TDL may learn faster and in a less wasteful way, but does it learn the right thing? Lucas' work seems to demonstrate that evolution is better in this respect.

A work that fits within the general framework of learning to mimic humans and behavioural cloning, was presented by Bryant and Miikkulainen [86]. In their paper, an embedded game agent is trained using human-computer interactions during a strategy game where the human follows pre-determined policies. The network takes input patterns obtained from game sensors, and outputs one of seven discrete actions. The evolution of the neural network uses the Enforced Sub-Populations (ESP) algorithm, which makes use of separate chromosomes for each neuron and maintains a distinct population for each neuron. Breeding is done separately within each population, and neurons co-evolve into individuals that work well together to solve the specific problem for which the network was trained. What may be the most powerful idea in this paper is the addition of Lamarkian learning to ESP. Lamarkian learning is based on a 19th century idea that biological phenotypical information (physically observable manifestations such as structure and behaviour) can be passed to off-spring along with genotypical information (the inherited code). Although the idea has not been proven in biology, it has potentially powerful applications in neuroevolution. Specifically, the weight changes learned during a "tuning" phase by using an evolved neural network's back-propagation learning abilities (phenotypical information), can be inserted directly in to the neurons' chromosomes (genotypical information). The authors report encouraging results with Lamarkian neuroevolution in that specific game policies and strategies were effectively learned. The authors also state that human behaviour mimicking is a difficult problem to solve as there are issues involved in defining what constitutes "good behaviour". For instance, does the game score effectively measure the visible intelligence learned by agents? The authors propose several possible measures including the use of qualitative psychological evaluation methods.

An interesting idea related to neuroevolution is the work of Reeder et al. [87] who investigated the interactive evolution of neural networks to control the shooting and movement of synthetic actors in a war game. The evolution is said to be interactive because the user can set the parameters of the fitness functions in real-time to obtain quicker convergence to a realistic evolution.

Yannakakis [88, 89] has studied the important issue of maintaining player satisfaction in mixed interaction, real-time games where humans and CGAs interact. In this work, satisfaction is equated to fun. Although it is not clear what causes fun and what fun actually is, one thing is known: fun games increase user satisfaction, augment the desire to play and provide a motivating environment for learning [90]. The work of Yannakakis investigates how to optimize player satisfaction and proposes ways to measure fun quantitatively by observing the physiological responses of humans to games. The physiological responses measured were: heart rate, blood volume pulse, and skin conductance.

One of Yannakakis' conclusions is that one can better model and enhance player satisfaction by learning the mapping between game features and user entertainment preferences. Hence, not only can machine learning be used to improve the behaviour of CGA, but it can also be used to optimize the quality of experience between humans and the computer environment. For this purpose, Yannakakis uses neuroevolution, a recurring theme in the application of AI to games. The latter approach results in a better modeling of human satisfaction than psychological-based, custom-designed modeling. The model thus acquired allows one to adjust opponents (red forces in a military game simulation) in real-time according to human user preferences and style. The adaptive, dynamic game generated in this manner, although simplistic and played by children, is preferred in 76% of cases over the static game. In short, Yannakakis has shown that neuro-evolved games are potentially successful predictors of entertainment value of games and can help in automating the design of games for optimized player satisfaction.

"Serious games" and are used, for example, to train people. Avery [91] conducted research into a training game in the context of dynamic Cold War scenarios on behalf of the Australian Department of Defence. The game consists of budget allocations into different types of defensive and offensive weapons taking into account intelligence and counter-intelligence information. The user must make the right choices that will preserve international peace. Avery's goal was to allow the game to adapt over time to the human trainee decisions and style. In its static form, the game is easily beaten by humans because it is predictable and so negatively affects the quality of training. The opponent player, an autonomous agent, is encoded as a fuzzy rule base in the form of

$$\langle rule, fuzzymembershipvalue \rangle \Rightarrow \langle action \rangle$$

The work of Avery uses previous human game interactions to evolve a fuzzy rule base which is a very approximate version of the human strategies. This fuzzy rule base is then used to create various competing artificial players who play previously recorded games; the evolutionary algorithm iterates, replacing the worst agent player with the best agent player at each iteration. The author conducted a user study to evaluate the approach and found that the users considered the evolved game much more interesting than the original static game. This work reinforces both the importance of the entertainment even in serious games, and the power of evolutionary computing.

FPS games are currently applied to fighting force training in a simulated environment. [92] In these games two teams engage each other in a simulated combat environment. Hladky and Bulitko [93] investigated the use of hidden semi-Markov models and particle filters to learn a position predictor, so that a synthetic player could behave in a more believable manner than current computer generated players. A CGA's ability to predict another player's position is learned from previous human expert players' game logs. The position predictor so trained demonstrated both skillfulness (achieving average human performance in tracking the adversary) and believability (making realistic mistakes, similar to the ones a human adversary would make). The approach was evaluated with user testing with the *Counter-Strike: Source* [18] anti-terrorist game.

In addition to modeling individual enemy entities in a military training simulation setting, one may also be interested in the modeling a group of entities, namely units at various levels from sections to whole armies. Strategy computer games such as *Age of the Empire III* [19] propose different approaches to dealing with such formations. The first is static formations, which lack realism, as they cannot adapt to unforeseen situations. Another is to let formations emerge using complexity theory-based approaches such as swarm computing or artificial life. This can also lack realism due to unpredictable behaviour and a difficulty in controlling what behaviour emerges. Van der Heijden [94] proposed a third approach, based on a stochastic optimization algorithm to select the best formation. Data collected during the game is used to classify opponents and apply an opponent model. Direction and speed are determined by following a leader entity. Losses during combat are accounted for and the formation dynamically re-adjusts itself. The approach is tested with the freely available *Open Real Time Strategy* (ORTS [95]) game environment, designed specifically for RTS game AI research. The results show that the dynamically-generated formations

18. http://www.valvesoftware.com/games/css.html
19. http://www.ageofempires3.com/

outperform five out of six opponents; the only opponent not defeated is one using what the author calls very high quality low-level AI, that is, very good hand crafted rules.

Case-based reasoning is another AI technique that has been used in the context of games.[96] It has been applied in an innovative way by Sander et al. [97] to provide rapidly adaptive games that do not rely on complex and lengthy learning algorithms. In this approach, game knowledge is acquired automatically in the form of a case base, from which an evaluation function is obtained and cases are used directly and immediately by the game. The results show an ability to win but also to maintain a tie, which from a user satisfaction point of view can be interesting as it provides a motivation to keep playing to beat the opponent. This idea is similar to using the history of the game to adapt the game to the player.

4.5 Large Scale Games

Finally, to realistically simulate a battlefield, hundreds or even thousands of entities may need to be generated and simulated, and interact with the humans beings being trained on the system. *Massively Multi-player Online Games* (MMOG) and *Massively Multiplayer Online Role Playing Games* (MMORPG) are promising environments to achieve this. This game architecture has the further advantage of allowing the training of people who are not co-located (a current problem with existing simulations). Mayo et al [98] investigated training of US Army troops in asymmetric warfare in a MMOG game environment. The goal was to train on cross-cultural communications in which CGA are generated with OneSAF Objective System to augment the human players involved in the scenarios. The functional evaluation of the system was made by soldiers involved in the training simulation and was of a non-technical nature. Most recognized the value of this training tool as an intermediate between real-life basic and advanced training, to increase decision making and situational awareness. Soldiers particularly appreciated the capability to measure and replay movements and rules of engagement violations as well as cause damage and visualize its effects (including injuries). On that latter point, participants appreciated the possibility to modify various types of improvised explosive devices and to blow them up on demand to support training in an asymmetric environment. They also identified important requirements such as the ability to select individual equipment and to control characters configurations and scenarios.

4.6 Industry Perspectives on AI

Whenever the topic of artificial intelligence in games is discussed, there is invariably the observation that much could be learned from the video game industry. The latter is actually not the case and, in fact, the opposite is true; the video game industry learns from others such as academia and M&S. The reasons for this stem from the goals and pressures of the video game industry.

Video games are “big business” and profit is the sole goal. Profit generation mandates that a gaming company must get their product out before competitors to seize market share, and that their game provides a positive user experience. These factors have a negative impact on the application of AI in games.

A positive player experience is the highest priority in video game development. All the money and time spent on development and production is done with this in mind. In an interview with Matthew Rice, the game AI developer for the FPS game *F.E.A.R.*, he reveals that a player notices the AI in a game under two conditions, “... the AI does something great, or he ran against the wall for the last three seconds.” [99]. Conversely, game reviews constantly comment on the quality of game graphics and animation. The result is that investment in AI is done only when necessary and only to the minimal extent required.

A positive player experience means that players must win, therefore the AI must be challenging but not dominant. A telling statement from an AI developer is that the “substantial” improvements in AI from *F.E.A.R.* to *F.E.A.R. 2*did not “ ... mean that the AI (bot) is any more difficult to kill, it just means that the environment is richer, and the player is more engaged in the combat.” and “ [the bot is] not more difficult, he’s just more realistic.” [99] This statement underlines that AI is not used to improve upon bot performance but instead is used to enhance realism, for example through more realistic mannerisms. The ramification is that complex AI techniques are not needed to satisfy industry needs and no investment is put towards developing AI further than rudimentary levels. Finally, time is of the essence in the video game industry and complex AI development can raise more problems that incur delays, than it can enhance player experience.

It is important to reiterate the magnitude of the task facing the video game industry to develop AI. A bot in a game does not have any awareness of the artificial world around it in the sense that a human does. As a result, any improvement in AI must be accompanied by a significant investment in providing cues in the game for the bot to react to. A game’s layout is usually overlaid by a “navmesh”, invisible to the player, that explicitly tells a bot that, for example, a particular bush is “good cover”. Similarly, every piece of information that a bot might require for decision making must be coded in to the fabric of the game’s playing field.

A bot also lacks any autonomous, corporeal existence. In short, a bot cannot perform a physical act that has not been previously animated. Therefore, the richer the set of possible actions a bot may take, the greater the number of animations must be produced to render those actions on a screen. The relationship is not a linear one, the ability to “jump” mandates animation for all the possible types of jumping required: forward, backward, sideways, near, far, etc.

Despite the challenges, the industry is including more AI in games. FPS games have crowds, that is, numerous simple NPCs moving with seeming purpose but not directly involved in the player’s story line. NPCs also display more realistic mannerisms and speech recognition but the nature of the AI used by the NPCs has changed little. [99, 100, 73, 101]

To convey a sense of the rate of progress in video game AI, consider that in 2004, Buro published his call for AI research in RTS games. [70] Four years later at the AIIDE 2008 conference [20], while there were many promising approaches *on the horizon*, there was still a perceived need to bring together AI researchers and game designers before the techniques presented could move in to development. [100] [21] By 2010, the situation had progressed somewhat in that some reviews of game AI's were commenting on the *style* of the bots' play rather than any evident flaws. Nonetheless, the technology behind the AI remained behaviour trees, finite state machines, and scripted behaviours. [102]. In 2011 industry developers still opined the need for game designers to work with AI practitioners. [101] Tellingly, the developers also observed:

1. game AI is still not "intelligent" as it is still reactive and does not "look ahead" to foresee what might happen;
2. additional computational power is unlikely to improve game AI unless the techniques applied are also changed, that is, existing hardware accommodates the techniques currently used; and
3. the need for game AI has actually decreased with the growth of multiplayer online games.

Generally, the video game industry operates under constraints that make the use of AI attractive, but that make it unlikely to lead AI development. The experience of the video game industry is relevant in highlighting the difficulty of increasing the realism of a game. The approach that has evolved is for the video game industry to focus efforts on improving the rendering of the game for maximum visual and aural experience, while academia has focused on approaches to AI. Collaboration between the industry and academia continues to be encouraged, and the AIIDE conference and the IEEE Computational Intelligence Task Force on Real Time Strategy Games [22] are in the vanguard of these efforts.

20. The AI and Interactive Digital Entertainment Conference (AIIDE) has hosted annually for seven years by Stanford, Palo Alto, USA. The conference brings together academia and industry to address AI topics.

21. Interestingly, even with the advent of multiple processor computers and graphical processor general computing, thought to be enablers of truly complex AI, the industry had not "... found a strong AI use for it." [100]

22. http://gameai.itu.dk/rtsg

5 Artificial Life, Complexity Theory and Embodied AI

Initially simulation used a single global AI which was applied to all the entities in a game. The approach was therefore monolithic. With the introduction of agents, the AI became distributed amongst the entities in a game. This distribution came closer to replicating reality as agents interacted with each other. An early observation was that the outcomes of agent interaction could not always be predicted. This circumstance occurred naturally because as agent behaviour became more complex, it became impossible for designers to foresee all the possible outcomes of design decisions. While this was certainly a challenge for game designers, it was also seen as a possible way ahead for AI development. The question was asked if it was possible that intelligent behaviour might be generated through agent interaction instead of being explicitly being coded into an agent. This idea is explored in this section.

5.1 Complexity, Emergence, and Artificial Life

Nobel Prize winner Herbert A. Simon was involved in the founding of AI [103], studied decision-making with uncertainty [104, 105] and is considered the founder of Complexity theory [106]. Complex systems are composed of large numbers of elements, such as agents, that interact among themselves and with their environment in a non-linear fashion. As a result, the collective of agents can exhibit properties and behaviour are not explicitly coded by the programmer. The phenomenon is known as emergence and can, in principle, allow one to avoid the difficulty of handcrafting all the rules of expert systems, production rule agents or behavioural models. One can instead let behaviours emerge naturally. Emergence has the further advantage of providing interesting, adaptable, and novel behaviours without having to consider all possible situations. This can make a CGA much more interesting but it is a double-edged sword. Since the behaviours that emerge are not foreseen, the issue arises of controlling these behaviours so that they correspond to what is desired, they are coherent (locally and globally), and they are realistic.

Artificial Life is related to complexity theory first by its approach based on biological system inspirations and second, by its philosophy of decentralized, bottom up behaviour emergence. Kim and Cho [107] survey the use of artificial life including many examples of applications to CGA or CGA related topics. One such example is the *MICROB* project (Making Intelligent Collective Robotics) that looks at self-organization of robots to perform a collective task, in this case a soccer game [108]. Similarly, *bio-mimicking* aims at reproducing the shape and/or the behaviour of biological beings, including humans. For instance, the Jet Propulsion Laboratory and Caltech have looked at the application of this technology to aircraft inspection by robots [109]. The applicability to CGA design is that mimicking humans ultimately leads to perfectly plausible, if not indistinguishable artificial beings, either as robots or as virtual characters.

An example of artificial life and emergence is the work of Hawick and Scogings [110]) who have demonstrated that very simple predator and prey artificial life entities will spontaneously self-organize in social groups given a spatial environment. The result is noteworthy as the entities possessed only two attributes, health and age, and a restricted set of micro-rule: breed, eat, kill, and pursue. The authors' *Animat* model achieves a dynamic equilibrium, making decisions autonomously based on the environment. The equilibrium is dynamic in that segregation of prey and predators into multiple social groups changes over time. This "emergent behaviour from spatial effects" is typical of complexity theory where emergent behaviour or structure is the result of the interaction of simple elements. The application to CGA could be in the self-organization and partitioning of troops based on geographical features and presence of enemy troops. The advantage of emergence is that one does not need to embark on a time consuming and costly hand-coding of behaviour: these rules are implicit, they arise naturally from interactions among CGA and with the environment. One can hypothesize that such emergent behaviour might be more believable or at the very least less predictable and more interesting to a human user.

5.2 Applied Emergence

Emergence has been studied in detail by Sweetser [111] in the context of strategic games. Sweetser's work focused on comparing scripted games, where agent behaviours are predetermined in hand-coded rules, with games in which behaviours emerge from interactions among agents and between agents and the environment. She used cellular automata [23] in a limited domain of a strategy game (the physical modeling portion of the environment) and influence maps [113]. Influence maps divide the world into a layered grid. Each cell of the grid contains local world information such as combat unit strength, available resources and traversability. The information is propagated with a decay, to neighboring cells. Agents can then decide which cell to move to based on a weighted sum of the various values stored in the cells. Sweetser investigated two movement approaches: comfort (move from danger to safety) and desire, a goal-directed behaviour (a mission to accomplish, such as a target to reach). The experiments conducted concluded that equal weighting of goal-directed and reactive behaviour in agents results in the most realistic behaviours. The behaviour obtained under emergence was found to be rich, interesting, and plausible. For instance, the author reports an example where a military unit en-route for an objective stopped when it started raining rather than risking being stuck.

Schaeffer et al. [76] give another interesting example of emergence. They report on *F.E.A.R* (Section 4.3), a commercial FPS that uses a simple strategy to emerge complex behaviours.

23. A cellular automata is a discrete state machine equipped with a set of simple rules that dictate the current state, based on the past state of itself and, usually, surrounding automata. Wolfram [112] provides an introduction and thorough study of a simple form of cellular automata.

Given a goal, each CGA has a set of the possible actions that can be called upon to reach that goal. Each CGA may take a different approach to reach a common goal, causing unexpected cooperative behaviours. For instance, the authors reported that as CGA aimed at attacking the player, some took cover at positions near the human player, while others went around, thus flanking the player and offering cross fire. This tactic was not coded explicitly in the program but emerged naturally from the existing set of simple actions that could be taken by each CGA.

Swarm intelligence is a stochastic optimization technique based on populations of entities [114]. Popular demonstrated populations include: birds (flocking), fish (schooling), particles, etc. Swarm intelligence has been used for UAV path planning [115] and other movement of troops in a military simulation. Swarming can also be useful to model grouping behaviour of troops. The behaviour introduced is often rigid and fixed by a set of predetermined rules. In an attempt to make swarming techniques adaptable, Morihiro et al. [116] introduced reinforcement learning as part of a self-organized grouping of agents. The authors report that the approach resulted in agent grouping and anti-predator behaviours that are diverse and robust.

An example application of anti-predator behaviour is escape and evasion tactics in a military scenario. Danielsiek et al. [117] look at the realistic movement of troops (5-7 individuals) in an influence map, given the presence of enemies and obstacles, as an emergent property of the interactions of three simple flocking rules: separation, cohesion and alignment. Their idea is to integrate an influence map, with attractive cells for friendly forces and repulsive cells for enemy forces, into the cost function of the A* path finding algorithm [71]. A parameter can be set by the user to control the aggressiveness of the troops, in which case they may run through enemy forces rather than avoid them if they "feel" they can defeat the enemy. The experiments have show that flocking makes all the difference; the troops move and fight together. The authors do note that the emergence of single file movement or the accommodation of the slowest individual in a group, caused more casualties. [24]

Embodied, or situated AI, was proposed by Brooks [65, 118, 119]. The approach uses interaction with the environment, as perceived through an agent's sensory organs, to derive behaviour. Embodiment research is biologically inspired and implements its learning mechanisms at different levels of abstraction, from biological plausible to abstractions of biological systems such as evolutionary and neural computing.

An application of the embodied AI concept was made by Thompson and Levine [120] who evolved an agent implemented with neural networks. The experiments used a subsumption architecture [118] which is a modular and hierarchically-layered approach to agent

24. This observation is only germane to video game play. In a military simulation such emerging behaviour would be very attractive and realistic. Lower performance can be an outcome when pursuing realism and of believability.

control. The lower levels are reactive to the sensory information while the higher levels address more abstract behaviours that can override lower level decisions. The agents used by Thompson and Levine comprised two layers, each made up of a small unsupervised neural network of $2-4$ inputs only. The setting was an adversary game where agents were evolved to navigate between to points while avoiding obstacles and enemy fire. The authors reported that the subsumption architecture allowed for incremental learning of complex behaviour better than a standard neural network. When faced with a change in the environment, such as an obstacle placed on the previously evolved path, the agents learned to go around the obstacle rapidly.

The latter work was only a first step, and much more work is needed to develop more complex behaviours. Embodiment research has not yet entirely achieved Brook's vision of completely autonomous agents. Ziemke [121] notes that the issues of embodiment are generally treated superficially in research and furthermore, that embodiment fails to recognize the important effect of emotions.

6 Artificial General Intelligence

Artificial intelligence is a field that originated approximately 50 years ago. Its inception gave great hope that it would lead to an understanding of the mind and perhaps show mankind a path to achieving greater intelligence, superseding that of humans. AI has yet to live up to that goal. Instead, the original idea of studying intelligence has been largely taken over by methodologies. Where researchers once considered "theories of the mind", with the Turing test as the benchmark, one now finds information processing techniques used on particular problems dominating AI research. This essentially follows the trend of technology toward computational approaches to virtually all aspects of science, that is, simulations. Recently, there has been a resurgence of interest in the roots of AI that has led to a branching of AI research into two distinct categories: narrow and artificial general intelligence.

Narrow AI is the current term used to refer to the numerous machine learning techniques described in literature, such as: neural networks, fuzzy logic, genetic algorithms, Bayesian networks, etc. Generally these methodologies have been applied to singular problems such as classification or clustering. These techniques have also been applied, to some extent, within the gaming industry as described in the preceding sections. The application in each case addressed specific aspects of the problem space such as path finding or decision making. The hallmark of narrow AI methodologies is that a solution applicable to a specific problem translates poorly, if at all, to other problems.

Artificial General Intelligence (AGI), by contrast, seeks to mimic human intelligence in that it can adapt to any problem. The first book [122] addressing the subject opens with the following statement:

> The AI projects discussed in this book, however, are quite different: they are explicitly aimed at artificial general intelligence, at the construction of a software program that can solve a variety of complex problems in a variety of different domains, and that controls itself autonomously, with its own thoughts, worries, feelings, strengths, weaknesses and predispositions.

AGI therefore seems a logical, possible approach to CGA's. The difficulty is that AGI is still in its infancy and as yet only has promising avenues of exploration. The roots of AGI are what inspired early attempts at general intelligence resulting in efforts such as the following:

1. GPS (General Problem Solver)[123]
2. CYC [25]

25. http://www.cyc.com/cyc/cycrandd

3. NARS (Non-Axiomatic Reasoning System) [26]
4. Soar [23]
5. ACT-R [25]
6. Bayesian networks
7. Neural Networks
8. Evolutionary Programming
9. Artificial Life
10. Program Search

Each of these approaches has had some limited success though many were not purely AGI-oriented. Goertzel [122] concludes that an integrative approach may be the most promising, that is, combining some or all of the approaches mentioned above. To this end, Goertzel et al. propose an architecture, the *Novamente Cognition Engine*[124] and it's open source development model, *OpenCog* [125, 126]. The architecture integrates many concepts from cognitive and complex systems science disciplines into a framework designed to operate on large scale distributed systems. Similarly, Bach [127] has proposed the *MicroPsi* architecture that is somewhat less ambitious but nonetheless falls into the category of AGI. The hallmark of both these approaches is that they reach beyond problem solving. Each of these projects seeks to imbue an agent with nothing less than a sense of existence, that is, autonomy and interaction with an environment based on needs beyond simple goals external to the agent. To accomplish this each project touches on aspects common to living creatures: emotion, memory, body parameters, etc. Both architectures have had some success as AI for virtual agents though none has yet been used in a game environment such as an RTS or FPS. [27]

Like the cognitive architectures that preceded it, AGI is clearly applicable to CGA development. However, the indications by researchers in the field are that considerable work remains to be done and it is not clear whether any current AGI applications can live up to its promise. Nevertheless, the field is one that should be closely watched as it is the only one truly probing the questions: *what* is general intelligence and *how* can such intelligence be implemented artificially.

---

26. http://www.cogsci.indiana.edu/farg/peiwang/papers.html
27. The software for MicroPsi can be found at http://www.cognitive-ai.com, while that for OpenCog is at http://opencog.org

7 Discussion

This survey observed four very broad themes in the application of AI to synthetic environments. The first is that modeling, simulation, and games are the principal forums for the application of AI but progress has been slow. The second is that the agent paradigm has become the prevalent paradigm, replacing monolithic architectures. The third is that learning is felt to be the next area of advancement. Finally, there is a return to a view of AI as a general, as opposed to the narrowly focused, problem-solver. This section summarizes these themes.

7.1 Modeling, Simulation, and Games

This survey of the application of AI to the generation of human-like, believable CGA, began with a look at the current state-of-the-art in M&S and video game CGA design. The field relies mostly on hand-crafted, rule-based systems (video games) or on cognitively-based models of human behaviours (M&S).

Cognitive models like *ACT-R*, *Soar* and *BDI* have been useful as they potentially equip CGA with rich, human-like behaviours. An aspect of realistic human behaviour modeling discussed is the inclusion of emotions and modulating factors such as fear and excitement. However, these cognitive AI approaches require intensive and costly knowledge acquisition efforts, and still often fail to capture the high variability of the environment and of human behaviours. Additionally, they do not scale well as the complexity of the problem increases. Decision and behaviour trees, and finite state machines dominate the video game NPC AI, that is, where AI techniques are applied. Most video games still rely on giving the illusion of AI by allowing the NPC the benefit of omniscient game knowledge or scripted behaviours.

The reality is that, after years of study, neither M&S nor video games are closer to achieving human-like, believable CGA AI. This point of view reflects the fact that the techniques used to implement AI have remained basically the same for two decades. Improved graphics and processor speed have certainly increased player experience and the implementation of realistic synthetic environments, but neither of the latter relates to AI. If a further confirmation of this state is needed, one need only look at the burgeoning growth of memberships for online gaming; players would rather play another human than the game AI's. As such, one might consider M&S and video games as forums to experiment with CGA AI, but academia is where AI development is being carried out.

7.2 Agents

Autonomous agents have become a popular AI technology to implement CGA. Agents are used in two ways. First, CGAs *are* the agents. This approach means that the agents are often hard-coded using either rule-bases, state machines, or some cognitive model. As such, these agents can suffer from some of the same limitations as non-agent CGA, namely, brittleness and lack of adaptability in realistic, dynamic and uncertain environments. Nevertheless, the modularity of this approach is appealing for development and the possibility of emergent behaviour amongst CGAs.

The second way the agent paradigm is used is to have a collection of agents constitute the AI of each individual CGA. In this approach, each component agent performs some particular aspect of the AI. Such an agent architecture often has advantages over other homogeneous AI implementations such as finite state machines, rule-based systems, case-based reasoning and neural networks. The component agents can combine a variety of approaches permitting better scaling in large knowledge domains. Furthermore, agent interactions in such multi-agent systems may allow the emergence of complex behaviours from simpler ones.

The agent paradigm will continue to dominate AI development. Its appeal lies in its modularity and the facility with which "systems of systems" can be created. The latter allows the possibility of emergent behaviour which holds the possibility of providing richer behaviour sets for CGAs.

7.3 Learning

The suggested solution to the problem of knowledge acquisition is to let the CGA adapt and learn, rather than hand-crafting the rules and models governing their behaviours. While learning has been experimented with in academia, few examples of learning are present in industry. Some games like *Creatures*[28] and *Black and White*[29] include learning as a fundamental part of the game but neither of these games is relevant to military concerns.[30] Therefore, it remains unclear how learning should be used for improving CGA AI.

The root of the problem is three-fold:

1. Humans learn in a sensory rich environment;
2. Learning is unpredictable; and
3. Learning is hard.

When learning is considered for CGA, most consider it in a human sense. One might expect a CGA to, like a human, learn the synthetic environment and then behave intelligently in

28. http://creatures.wikia.com
29. http://lionhead.com/Games/BW2
30. The learning that takes place in those games is not intended to make the NPC's more intelligent. Instead, it is a plot device used to avoid repetitive behaviour. It is notable that unexpected behaviours do arise though they could not be considered believable.

that environment. Unfortunately, the illusion created by high level graphics hides the fact that a CGA has a paucity of sensory information. A CGA has no muscle feedback, no sense of balance, smell, hearing, or even vision. In the absence of rich sensory information, learning can only be very limited.

Learning is unpredictable. In as much as a CGA can potentially learn good behaviours, it can learn equally poor behaviours. This is, in part, because of the lack of any sensory information to provide negative feedback. But it is also because the space of possible solutions to a problem usually contains a greater number of suboptimal solutions than optimal ones. This poses a significant liability for any attempt at implementing learning during game play (*online* learning).

Learning is hard. As the preceding paragraphs point out, a CGA must be provided with all the necessary parameters for learning by the developer. This means that careful game design must populate the synthetic environment with appropriate cues for the CGA to use. Furthermore, all learning behaviour must be carefully designed, monitored, constrained, and tested to ensure that suitable behaviour is the result. For these reasons, few games have online learning and offline learning is far from trivial.

Despite the challenges, learning is still touted as being a desired feature of a game. If that is the case then two questions remain: *why* is learning required and *what* is to be learned?

The reasons *why* learning is desired *seem* obvious, that is, to have CGA's that can adapt, innovate, and challenge a human player or each other. Millington dismisses the idea that CGAs could learn, from scratch, how to defeat a human opponent as "almost pure fantasy" [71]. The challenges cited above explain why this might be so. This means that learning must be more focused and so, in effect, distill to simply the optimization of a restricted behaviour set or of the gamer experience. This is an important nuance; in the absence of a rich set of actions and precepts, what is learned will be simplistic. This is acceptable if such learning meets the M&S or game requirements.

The question of *what* is to be learned is critical. There are four possibilities, some of which are used in existing games. They are:

1. learn appropriate parametrization - this includes the learning of parameters for steering behaviours, path finding cost functions, probabilities, etc. This is used for tuning neural networks in the driving game, *Forza* [31] [100]
2. learn appropriate decisions for the use of a limited set of actions - including decision trees for the most effective use of, say, attack, retreat, and hold actions.
3. learn to predict opponent actions - such as learning any patterns of behaviour that an opponent might have in order inform decision making. The game *F.E.A.R. 2* did this in testing before deployment. [99].

31. http://forzamotorsport.net

4. learn opponent behaviour - in the context of modulating game difficulty, such as generating more challenges when a player is doing well, in order to improve a player's experience. The game *Left 4 Dead* [32] uses this. [99]

Learning is an appealing concept but it is also difficult, time consuming and unpredictable. The application of learning in M&S and games is still very much evolving and so cannot be regarded as a near-term solution to achieving realistic CGA behaviour. Learning is certainly an area that needs to be further explored and exploited when designing CGA, but realistically, one should not expect learning to be a silver bullet for CGA human-like behaviour.

7.4 Embodiment, Artificial Life, and General AI

The final areas of AI applicability to CGA that were reviewed were embodied AI, artificial life and complexity, and General AI. The interesting aspects of these techniques is that instead of relying on an existing model of the world, a behavioural model, or other forms of hand-crafted rules to describe behaviours and the environment, one lets the model emerge through interactions and sensory inputs. This has the potential to result in interesting, adaptable behaviours for CGA. Nevertheless, there is no indication that any of these approaches will bear fruit in the near future. They are still very much academic endeavours.

32. http://www.ea.com/left-4-dead-2

8 Conclusion

One can identify two main approaches to CGA design with AI. The first is knowledge engineering and the handcrafting of behaviours in production systems, cognitive models or finite state machines. The second approach is the learning of behaviours, which consists of any form of adaptation, including: evolution, emergence, neural networks or reinforcement. The main problems with the first approach are brittle and predictable CGA behaviour. Its main advantages are that these techniques are relatively mature and one can code exactly what is needed in a controlled manner. Learning has the disadvantage of not always resulting in the desired behaviour. The advantage of learning, however, is that CGA can adapt to the environment and/or the user, resulting in more robust, varying and interesting behaviours.

It is likely that a mix of AI techniques, rather than a single one, will be more successful. One approach cannot simply be dismissed over another. Hybrid approaches such as using existing cognitive models or rule bases, coupled with learning, can potentially achieve interesting results, including emergence. One possible approach might be that basic knowledge can be coded while the rest can be evolved, or learned incrementally, during the simulation.

This survey reveals that the quality of believability (Section 1) has been achieved somewhat in existing games. The following shows where gaps exist:

Quality	Achieved?	Comment
consistent, lifelike behaviour	NO	AI behaviour in all games show inconsistencies to varying degrees.
illusion of intelligent life, including personality and emotions	NO	Emotions still not successfully implemented.
social interaction	PARTIALLY	The game *Creatures* had emergent social interaction. Generally, such behaviour is scripted.
adaptation to circumstances, ability to change and evolve	YES	The adaptation is seen in academic simulations or refers to the game adjusting its difficulty to modulate player experience.
self-motivation	PARTIALLY	The game *Creatures* had characters capable of self-motivation, but the authors know of no other implementations.

This survey concludes with following suggestions for future work on the topic of AI for realistic, believable CGA:

1. Develop a general hybrid architecture of AI for CGAs using the concepts identified in this survey. A focus on the inclusion of emotions, self-motivation, and lifelike behaviour would be an important contribution to the body of knowledge.
2. Design a user-based performance evaluation methodology, taking into account factors such as believability and entertainment factors and their impact on the quality of the simulation task (for instance, learning).
3. A build, test and evaluate the architecture in a test bed CGA and simulation world, possibly using off-the-shelf and open-source game environments (some have been mentioned in this paper).
4. Investigate the feasibility of integrating the AI architecture in an actual DND simulation platform. (An alternative may be to develop a new system from basic AI principles since retrofitting an existing system may be too difficult).
5. Design, based on the lessons learned and results from the previous steps, a complete AI-based CGA system or integrate the AI models in to an actual DND simulation platform.

Reproducing intelligence in machines has proven a much harder task than originally anticipated. Although science may not have discovered the fundamental principles that would allow for the creation of human-like intelligence in computers, this survey shows that there are a multitude of interesting techniques available that can help make CGA's more interesting and more believable. If the task seems daunting, the following quote reminds us that there is no single solution one can prejudge as being better than another:

> There have been countless examples of difficult to implement, complex AI that can come out looking stupid. Equally, a very simple technique, used well, can be perfect. [71]

References

[1] Clark, Pongracic Helen, Peter and Chandran, Arvind (2000), Researching the Use of Agent-Based CGF in Human-in-the-Loop Simulation, In *Proc. of the 9th Conf. on Computer Generated Forces and Behavioral Representation (BRIM)*.

[2] Riedl, Mark O. and Stern, Andrew (2006), Believable Agents and Intelligent Scenario Direction for Social and Cultural Leadership Training, In *BRIM 06: 15th Conference on Behavior Representation in Modeling and Simulation*, SISO.

[3] Grainger, B. and Henderson, C. (2003), Battlespace Composition in the OneSAF Objective System, In *Simulation Interoperability Workshops*.

[4] Perla, Markowitz M. Nofi A. Weuve C. Loughran J., P. and Stahl, M. (2000), Technical Report CRM D0002722.A2/Final, Technical Report Centre for Naval Analyses.

[5] Stytz, Martin R. and Banks, Sheila B. (2003), Progress and Prospects for the Development of Computer-Generated Actors for Military Simulation: Part 1-Introduction and Background, *Presence: Teleoperators & Virtual Environments*, 12(3), 311+.

[6] Bates, J. (1997), The Role of Emotion in Believable Agents, *Communications of the ACM*, 37(7), 122–125.

[7] Mateas, M (2002), Interactive Drama, Art and Artificial Intelligence, Ph.D. thesis, School of Computer Science, Carnegie Mellon University.

[8] Bruzzone, A.; Brandolini M. UKSIM 2003. 2005 128-32, A.; Orsoni (2005), Simulating human behaviour by integrating AI techniques, In *Sixth National Conference of the United Kingdom Simulation Society (UKSIM)*.

[9] Bryant, Bobby Don (2006), Evolving visibly intelligent behavior for embedded game agents, Ph.D. thesis, Austin, TX, USA. Adviser-Miikkulainen, Risto.

[10] Stytz, M.R. and Banks, S.B. (2002), Advances in Intelligent Systems for Defence by (Hardcover - Nov 2002), Ch. Chap 3 Computer Generated Forces in Distributed Simulation, World Scientific Pub Co Inc.

[11] Stytz, Martin R. and Banks, Sheila B. (2003), Progress and prospects for the development of computer generated actors for military simulation: part 3–The road ahead, *Presence: Teleoper. Virtual Environ.*, 12(6), 629–643.

[12] Lenat, Douglas B. (1995), CYC: a large-scale investment in knowledge infrastructure, *Commun. ACM*, 38(11), 33–38.

[13] Massey, L. (2007), Contrast Learning for Conceptual Proximity Matching, *Machine Learning and Cybernetics, 2007 International Conference on*, 7, 4044–4049.

[14] Gonzalez, A. J., Castro, J., and Gerber, W. (2005), Automating the Acquisition of Tactical Knowledge for Military Missions, *Journal of Defense Modeling and Simulation*, 2(3), 145–160.

[15] Von der Lippe, & Courtemanche A. J., S. R. (1998), Interim results in the development of user composable behaviors, In *Proceedings of the 7th Conference on Computer Generated Forces and Behavior Representation*, p. 545553.

[16] Henderson, C. and Grainger, B. (2002), Composable behaviors in the OneSAF Objective System, In *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*.

[17] Tollefson, Eric S., Schamburg, Jeffrey B., and Yamauchi, Harold M. (2006), A methodology for conducting composite behavior model verification in a combat simulation, In *WSC '06: Proceedings of the 38th conference on Winter simulation*, pp. 1286–1293, Winter Simulation Conference.

[18] Sui, Qing, Yeo, Ye C., and How, Khee Y. (2008), An Advanced Rule Engine for Computer Generated Forces, In *17th BRIM Conf: Conference on Behavior Representation in Modeling and Simulation*, SISO, US/Curran Assocs Inc.

[19] Landweer, P. (1993), Action/cognition behavior model for computer generated forces, In *Proceedings of the 3rd Conference on Computer Generated Forces and Behavior Representation*, p. 335345.

[20] Weaver, & Mullen W. J., W. B. (1994), Simulating generic military decision making with an empirically-trained neural. network, In *Proceedings of the 4th Conference on Computer Generated Forces and Behavioral Representation*, pp. 531–540.

[21] Busetta, Ronnquist R. Hodgson A. & Lucas A., P. (1999), Light-weight intelligent software agents in simulation, In *Proceedings of the 4th International SIMTECT Conference*, p. 169174.

[22] Dean, D., Syms, P., Hynd, K., Mistry, B., and Vincent, A. (2006), Modelling and Simulation of Combat ID - the INCIDER Model, In *Proceedings of the 2006 IEEE Symposium on Computational Intelligence and Games (CIG06), University of Nevada, Reno, campus in Reno/Lake Tahoe, 22-24 May, 2006*, pp. 156–163, IEEE.

[23] Laird, John E., Newell, Allen, and Rosenbloom, Paul S. (1987), SOAR: an architecture for general intelligence, *Artif. Intell.*, 33(1), 1–64.

[24] Anderson, John R. (1983), The Architecture of Cognition, Cambridge, MA, USA: Harvard University Press.

[25] Byrne, M. D., Kirlik, A., Fleetwood, M. D., Huss, D. G., Kosorukoff, A., Lin, R., and Fick, C. S. (2004), A Closed-Loop, ACT-R Approach to Modeling Approach and Landing With and Without Synthetic Vision System (SVS) Technology, In *Proceedings of the Human Factors and Ergonomics Society 48th Annual Meeting*, Santa Monica, CA: Human Factors and Ergonomics Society.

[26] Best, Lebiere C. & Scarpinatto C., B. (2002), A model of synthetic opponents in MOUT training simulations using the ACT-R cognitive architecture, In *Proceedings of the Eleventh Conference on Computer Generated Forces and Behavior Representation*.

[27] Kennedy, William G., Bugajska, Magdalena D., Adams, William, Schultz, Alan C., and Trafton, Gregory J. (2008), Incorporating Mental Simulation for a More Effective Robotic Teammate, In Fox, Dieter and Gomes, Carla P., (Eds.), *AAAI*, pp. 1300–1305, AAAI Press.

[28] Reece, D. (2001), Notes on Cognitive Limitations of DISAF, Technical Report SAIC Technical Report.

[29] Wedgwood J., Chadha B. Stavash J. Eden J., Cortese A. and Welsh J. (2003). ", Orlando FL September 2003, " (2003), Genetic Programming Approaches to Composable Simulation, In *SISO Simulation Interoperability Workshop*.

[30] Fletcher, M. (2006), A Cognitive Agent-based Approach to Varying Behaviours in Computer Generated Forces Systems to Model Scenarios like Coalitions, In *IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications, 2006. DIS 2006. In Proceedings - Thirteenth International Symposium on Temporal Representation and Reasoning, TIME 2006*, p. 297 284.

[31] Maes, Pattie (1994), Modeling adaptive autonomous agents, *Artificial Life*, 1, 135–162.

[32] Russell, S. J. and Norvig (2003), Artificial Intelligence: A Modern Approach (Second Edition), Prentice Hall.

[33] Franklin, Stan and Graesser, Art (1997), Is It an agent, or just a program?: A taxonomy for autonomous agents, *Intelligent Agents III Agent Theories, Architectures, and Languages*, pp. 21–35.

[34] Huebscher, Markus C. and Mccann, Julie A. (2008), A survey of autonomic computing—degrees, models, and applications, *ACM Comput. Surv.*, 40(3), 1–28.

[35] Aydin, Ali O., Orgun, Mehmet A., and Nayak, Abhaya (2008), The Reactive-Causal Architecture: Towards Development of Believable Agents, In *IVA '08: Proceedings of the 8th international conference on Intelligent Virtual Agents*, pp. 468–469, Berlin, Heidelberg: Springer-Verlag.

[36] Brandolini, Matteo, Rocca, Attilio, Bruzzone, Agostino G., Briano, Chiara, and Petrova, Petranka (2004), Poly-functional intelligent agents for computer generated forces, In *WSC '04: Proceedings of the 36th conference on Winter simulation*, pp. 1045–1053, Winter Simulation Conference.

[37] Huang, Eliens A., Z. and Visser, C. (2002), 3D agent-based virtual communities, In *Proceedings of the Seventh International Conference on 3D Web Technology*, pp. 137–143.

[38] Weiss, Gerhard, (Ed.) (1999), Multiagent systems: a modern approach to distributed artificial intelligence, Cambridge, MA, USA: MIT Press.

[39] Sokolowski, J.A. (2003), Enhanced Decision Modeling Using Multiagent System Simulation, *SIMULATION*, 79(4), 232–242.

[40] Ryan, Malcolm (2008), Constraint-Based Multi-agent Path Planning, In Wobcke, Wayne and Zhang, Mengjie, (Eds.), *Australasian Conference on Artificial Intelligence*, Vol. 5360 of *Lecture Notes in Computer Science*, pp. 116–127, Springer.

[41] Wang, Yingxu (2009), A Cognitive Informatics Reference Model of Autonomous Agent Systems (AAS), *IJCINI*, 3(1), 1–16.

[42] Best, C., B. J. & Lebiere (2008), Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation, Ch. Cognitive agents interacting in real and virtual worlds, pp. 186–218, New York, NY, USA: Cambridge University Press.

[43] Noy, N. F. and McGuiness, D. L. (2001), Ontology Development 101: A Guide to Creating Your First Ontology, *Stanford Knowledge Systems Laboratory Technical Report*, (KSL-01-05).

[44] Davies, John (2006), Semantic Web Technologies, Chichester: John Wiley & Sons.

[45] Grimaldo, Francisco, Lozano, Miguel, Barber, Fernando, and Vigueras, Guillermo (2008), Simulating socially intelligent agents in semantic virtual environments, *Knowledge Eng. Review*, 23(4), 369–388.

[46] Zsambok, C. E. (1997), Naturalistic decision making, Ch. Naturalistic decision making: Where are we now?, pp. 3–16, Mahwah, NJ.: Lawrence Erlbaum.

[47] Minsky, Marvin (1995), A framework for representing knowledge, Menlo Park, CA, USA: American Association for Artificial Intelligence.

[48] Aguirre, Héctor R., Corchado, Félix F., Ruvalcaba, Luis A., Rios, Jaime A., and Thalmann, Daniel (2008), Use of Intelligent Emotional Agents in the Animation of Autonomous Virtual Creatures, In *MICAI '08: Proceedings of the 2008 Seventh Mexican International Conference on Artificial Intelligence*, pp. 420–425, Washington, DC, USA: IEEE Computer Society.

[49] Damasio, Antonio R. (1994), Descartes' Error: Emotion, Reason, and the Human Brain, Avon Books.

[50] Goleman, D. (1995), Emotional Intelligence, Bantam Books.

[51] Lim, Mei Y., Dias, Joao, Aylett, Ruth, and Paiva, Ana (2008), Improving Adaptiveness in Autonomous Characters, In *IVA '08: Proceedings of the 8th international conference on Intelligent Virtual Agents*, pp. 348–355, Berlin, Heidelberg: Springer-Verlag.

[52] Cattinelli, Isabella, Goldwurm, Massimiliano, and Borghese, Alberto N. (2008), Interacting with an artificial partner: modeling the role of emotional aspects, *Biol. Cybern.*, 99(6), 473–489.

[53] Bosse, Tibor, Lam, Dung N., and Barber, Suzanne K. (2008), Tools for analyzing intelligent agent systems, *Web Intelli. and Agent Sys.*, 6(4), 355–371.

[54] Dumitrache, Ioan (2008), From model-based strategies to intelligent control systems, In *ICAI'08: Proceedings of the 9th WSEAS International Conference on International Conference on Automation and Information*, pp. 408–415, Stevens Point, Wisconsin, USA: World Scientific and Engineering Academy and Society (WSEAS).

[55] Camponogara, Eduardo and Kraus, Werner (2003), Distributed Learning Agents in Urban Traffic Control, *Progress in Artificial Intelligence*, pp. 324–335.

[56] Klügl, Franziska, Hatko, Reinhard, and Butz, Martin V. (2008), Agent Learning Instead of Behavior Implementation for Simulations — A Case Study Using Classifier Systems, In *MATES '08: Proceedings of the 6th German conference on Multiagent System Technologies*, pp. 111–122, Berlin, Heidelberg: Springer-Verlag.

[57] Shoham, Y., Powers, R., and Grenager, T. (2003), Multi-agent reinforcement learning: a critical survey, Technical Report Stanford University.

[58] Deutsch, T., Gruber, A., Lang, R., and Velik, R. (2008), Episodic Memory for Autonomous Agents, In *Conference on Human System Interaction (HSI 2008)*, pp. 621–626, Krakau.

[59] Heuvelink, Annerieke, Klein, Michel C.A., and Treur, Jan (2008), An Agent Memory Model Enabling Rational and Biased Reasoning, *Web Intelligence and Intelligent Agent Technology, IEEE/WIC/ACM International Conference on*, 2, 193–199.

[60] Kechagias, D.D., Symeonidis, A.L., Mitkas, P.A., and Alborg, M. (2002), Towards improving multi-agent simulation in safety management and hazards control environments, In *AI, Simulation and Planning in High Autonomous Systems - AIS2002*, Lisbon, Portugal.

[61] Kantardzic, M. (2003), Data Mining: Concepts, Models, Methods, and Algorithms, John Wiley & Sons.

[62] Massey, L. (2005), Real-world text clustering with adaptive resonance theory neural networks, *Neural Networks, 2005. IJCNN '05. Proceedings. 2005 IEEE International Joint Conference on*, Vol. 5.

[63] Witten, Ian and Frank, Eibe (2005), Data Mining, San Francisco: Morgan Kaufman.

[64] Giannikis, Georgios K. and Daskalopulu, Aspassia (2008), How Can Agents Know What to Assume When?, *Web Intelligence and Intelligent Agent Technology, IEEE/WIC/ACM International Conference on*, 2, 538–545.

[65] Brooks, R. A. (1991), Intelligence without Representation., *Artif. Intell.*, 47, 139–159.

[66] Fong, Gwenda (2004), Adapting cots games for military simulation, In *VRCAI '04: Proceedings of the 2004 ACM SIGGRAPH international conference on Virtual Reality continuum and its applications in industry*, pp. 269–272, New York, NY, USA: ACM.

[67] McDowell, Darken R. Johnson & E. Sullivan J., P. (2008), Delta3D: A Complete Open Source Game and Simulation Engine for Building Military Training Systems, *Journal of Defense Modeling and Simulation*.

[68] Rothstein, E. (1994), A New Art Form May Arise From the 'Myst', *The New York Times*, Vol. Arts.

[69] Grand, S. (2000), Creation : life and how to make it, London: Orion House.

[70] Buro, Michael and Furtak, Timothy M. (2004), RTS games and real-time AI research, In *In Proceedings of the Behavior Representation in Modeling and Simulation Conference (BRIMS*, pp. 51–58.

[71] Millington, Ian (2006), Artificial Intelligence for Games, First Edition, San Diego: Morgan Kaufmann.

[72] Rabin, Steve (2008), Ai Game Programming Wisdom 4, London: Charles River Media.

[73] Krajewski, J. (2009), Creating All Humans: A Data-Driven AI Framework for Open Game Worlds. https://skatgame.net/mburo/orts/orts.html (Access date: March 2009).

[74] Champandard, A.J. (2004), AI in Game Development, New Riders Publishing.

[75] Yildirim, Sule and Stene, Sindre B. (2008), A survey on the need and use of AI in game agents, In *SpringSim '08: Proceedings of the 2008 Spring simulation multiconference*, pp. 124–131, New York, NY, USA: ACM.

[76] Schaeffer, J., Bulitko, V., and Buro, M. (2008), Bots Get Smart, *IEEE Spectrum*.

[77] Ulam, Patrick, Jones, Joshua, and Goel, Ashok K. (2008), Combining Model-Based Meta-Reasoning and Reinforcement Learning for Adapting Game-Playing Agents, In Darken, Christian and Mateas, Michael, (Eds.), *AIIDE*, The AAAI Press.

[78] Sailer, F., Buro, M., and Lanctot, M. (2007), Adversarial Planning Through Strategy Simulation, In *Computational Intelligence and Games, 2007. CIG 2007. IEEE Symposium on*, pp. 80–87.

[79] Kelly, John P., Botea, Adi, and Koenig, Sven (2008), Offline Planning with Hierarchical Task Networks in Video Games, In Darken, Christian and Mateas, Michael, (Eds.), *AIIDE*, The AAAI Press.

[80] Miikkulainen, R., Bryant, B. D., Cornelius, R., Karpov, I. V., Stanley, K. O., and Yong, C. H. (2006), Computational intelligence in games, In Yen, G. Y. and Fogel, D. B, (Eds.), *Computational Intelligence: Principles and Practice*, p. 155191, Piscataway, NJ: IEEE Computational Intelligence Society.

[81] Schrum, Jacob and Miikkulainen, Risto (2008), Constructing Complex NPC Behavior via Multi-Objective Neuroevolution, In *Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference, October 22-24, 2008, Stanford, California, USA*, The AAAI Press.

[82] Lucas, S.M. (2008), Investigating learning rates for evolution and temporal difference learning, In *Computational Intelligence and Games, 2008. CIG '09. IEEE Symposium On*, pp. 1–7.

[83] Lucas, S. M. (2008), Computational Intelligence and Games: Challenges and Opportunities, *International Journal of Automation and Computing*, 5, 45 – 57.

[84] Sutton, & Barto A. G., R. S. (1998), Reinforcement learning: An introduction, MIT Press.

[85] Mcpartland, Michelle and Gallagher, Marcus (2008), Learning to be a Bot: Reinforcement Learning in Shooter Games, In *Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference, October 22-24, 2008, Stanford, California, USA*, The AAAI Press.

[86] Bryant, B. D. and Miikkulainen, R. (2007), Acquiring Visibly Intelligent Behavior with Example-Guided Neuroevolution, In *Twenty-Second National Conference on Artificial Intelligence (AAAI 2007)*, Meno Park, CA: AAAI Press.

[87] Reeder, J., Miguez, R., Sparks, J., Georgiopoulos, M., and Anagnostopoulos, G. (2008), Interactively Evolved Modular Neural Networks for Game Agent Control, In *Computational Intelligence and Games, 2008. CIG '09. IEEE Symposium On*, pp. 167 – 174, Perth, Australia.

[88] Yannakakis, G. N. and Hallam, J. (2007), The AIIDE'07 Workshop on Optimizing Player Satisfaction, *AI Magazine*, 28(4), 131–132.

[89] Yannakakis, G. N. (2008), How to Model and Augment Player Satisfaction: A Review Proceedings of the 1st Workshop on Child, Computer and Interaction, ICMI'08, Chania, Crete, October, 2008. ACM Press, In *Proceedings of the 1st Workshop on Child , Computer and Interaction, ICMI'08,*.

[90] Koster, R. (1981) . (1981), A Theory of Fun for Game Design, Paraglyph Press.

[91] Avery, P. (2008), Adapting to Human Game Play, In *IEEE Symposium on Computational Intelligence and Games*, Perth, Australia.

[92] Raybourn, E.M., Deagle, E., Mendini, K., and Heneghan, J. (. . I/ITSEC 2005 Proceedings (2005), Adaptive Thinking & Leadership Simulation Game Training for Special Forces Officers, In *Interservice/ Industry Training, Simulation and Education Conference Proceedings*, Orlando, Florida.

[93] Hladky, S. and Bulitko, V. (2008), An Evaluation of Models for Predicting Opponent Locations in First-Person Shooter Video Games, In *IEEE Symposium on Computational Intelligence and Games.*, Perth, Australia.

[94] van der Heijden, M. (2008), Dynamic formations in real-time strategy games, Masters, Department of Computer Science, Universiteit Maastricht, The Netherlands.

[95] ORTS2010 (2010), Online Realtime Strategy. https://skatgame.net/mburo/orts/orts.html.

[96] Ontan, S. and Ram, A. (2011), AI for Games: State of the Practice, Ch. Case-Based Reasoning and User-Generated AI for Real-Time Strategy Games, Entertainment Intelligence Lab, Georgia Tech University.

[97] Sander Bakkes, Pieter Spronck and van den Herik, Jaap (2009), Rapid and Reliable Adaptation of Video Game AI, *IEEE Transactions on Computational Intelligence and AI in Games*, Volume 1(2).

[98] Mayo, M., Singer, M. J., and Kusumoto, L. (2005), Massively Multi-Player (MMP) Environments for Asymmetric Warfare., In *27th Interservice/Industry Training Systems and Education Conference*, Arlington, VA: National Training Systems Association.

[99] Nutt, C. (2008), The Technology of F.E.A.R. 2: An interview on engine and ai development. http://www.gamasutra.com (Access date: November 2009).

[100] Kline, D. (2008), In-Depth: Stanford Conference Explores the State of AI. http://www.gamasutra.com/ (Access date: November 2009).

[101] Peterson, S. (2011), AI: Building Smarter Games. http://www.industrygamers.com/news/ai-building-smarter-games/ (Access date: December 2011).

[102] Champandard, A.J. (2011), This Year in Game AI: Analysis, Trends from 2010 and Predictions for 2011. http://aigamedev.com/ (Access date: December 2011).

[103] Simon, H. A. (1956), A Comparison of Game Theory and Learning Theory, *Psychometrika*, 21(3), 267–272.

[104] Simon, H. A. (1959), Theories of Decision-Making in Economics and Behavioral Science, *The American Economic Review*, 49, 253–283.

[105] Simon, H. A. (1960), Simulation of Individual and Group Behavior, *American Economic Review*, 50, 920–932.

[106] Simon, Herbert A. (1962), The Architecture of Complexity, *Proceedings of the American Philosophical Society*, 106(6), 467–482.

[107] Kim, Kyung J. and Cho, Sung B. (2006), A Comprehensive Overview of the Applications of Artificial Life, *Artif. Life*, 12(1), 153–182.

[108] Collinot, A., Drogoul, A., and Benhamou, P. (1996), Agent oriented design of robotic soccer team., In *ICMAS96*, p. 4157, American Association for Artificial Intelligence.

[109] Bar-Cohen, Y. (2002), Biologically Inspired Robots as Artificial Inspectors, In *SPIE Smart Structures and Materials Symposium, NDE Conference*, San Diego, CA.

[110] Hawick, K. A. and Scogings, C. J. (2008), Emergent Spatial Agent Segregation, *Web Intelligence and Intelligent Agent Technology, IEEE/WIC/ACM International Conference on*, 2, 34–40.

[111] Sweetser, Penny (2007), Emergence in Games, London: Charles River Media.

[112] Wolfram, S. (2002), A New Kind of Scien, Wolfram Media, Inc.

[113] Sweetser, P. (2003), Strategic Decision-Making with Neural Networks and Influence Maps, In Rabin, S., (Ed.), *AI Game Programming Wisdom 2.*, pp. 439–446, Hingham, MA: Charles River Media, Inc.

[114] Eberhart, R. C. and Kennedy, J. (1995), A new optimizer using particle swarm theory, In *Sixth International Symposium on Micromachine and Human Science*, pp. 39–43, Nagoya, Japan.

[115] Jennings, A.L., Ordonez, R., and Ceccarelli, N. (2008), An Ant Colony Optimization using training data applied to UAV way point path planning in wind, In *Swarm Intelligence Symposium, 2008. SIS 2008. IEEE*, pp. 1–8.

[116] Morihiro, Koichiro, Nishimura, Haruhiko, Isokawa, Teijiro, and Matsui, Nobuyuki (2008), Learning Grouping and Anti-predator Behaviors for Multi-agent Systems, In *KES '08: Proceedings of the 12th international conference on Knowledge-Based Intelligent Information and Engineering Systems, Part II*, pp. 426–433, Berlin, Heidelberg: Springer-Verlag.

[117] Danielsiek, H., Ster, R., Thom, A., Beume, N., Naujoks, B., and Preuss, M. (2008), Intelligent Moving of Groups in Real-Time Strategy Games, In *IEEE Symposium on Computational Intelligence and Games (CIG 2008)*, pp. 63–70, IEEE Press.

[118] Brooks, R. A. (1986), A Robust Layered Control System for a Mobile Robot, *IEEE Journal of Robotics and Automation*, RA-2(1), 14–23.

[119] Brooks, R. A. (1987), Planning is Just a Way of Avoiding Figuring Out What To Do Next, (Working Paper 303) Massachusetts Institute of Technology - Artificial Intelligence Laboratory.

[120] Thompson, T. and Levine, J. (2008), Scaling-up behaviours in EvoTanks: Applying subsumption principles to artificial neural networks, In *Computational Intelligence and Games, 2008. CIG '09. IEEE Symposium On*, pp. 159–166.

[121] Ziemke, T. (2008), Emotion in cognitive systems architectures Artificial Intelligence: Theories, Models and Applications, In *5th Hellenic Conference on AI, SETN 2008*, Syros, Greece.

[122] Goertzel, Ben and Pennachin, Cassio (2007), Artificial General Intelligence (Cognitive Technologies), Secaucus, NJ, USA: Springer-Verlag New York, Inc.

[123] Newell, A., Shaw, J.C., and Simon, H.A. (1959), Report on a general problem-solving program, In *Proceedings of the International Conference on Information Processing*, p. 256264.

[124] Looks, M., Goertzel, B., and Pennachin, C. (2004), Novamente: An integrative architecture for general intelligence, In *In AAAI Fall Symposium, Achieving Human-level intelligence*.

[125] Goertzel, Ben, Pitt, Joel, Vepstas, Linas, Gama, Gustavo, Wigmore, Jared, and Waddington, Trent (2009), OpenCog. `http://mloss.org/software/view/173/` .

[126] Hart, David and Goertzel, Ben (2008), OpenCog: A Software Framework for Integrative Artificial General Intelligence, In *Proceedings of the 2008 conference on Artificial General Intelligence 2008: Proceedings of the First AGI Conference*, pp. 468–472, Amsterdam, The Netherlands, The Netherlands: IOS Press.

[127] Bach, Joscha (2003), The MicroPsi Agent Architecture, Vol. 1, pp. 15–20, Universitts-Verlag.

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) Defence R&D Canada – Ottawa 3701 Carling Avenue, Ottawa, Ontario, Canada K1A 0Z4	2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.) UNCLASSIFIED (NON-CONTROLLED GOODS) DMC A REVIEW: GCEC June 2010

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Artificial Intelligence in Games

4. AUTHORS (Last name, followed by initials – ranks, titles, etc. not to be used.)

Bourassa, M.A.J.; Massey, L.

5. DATE OF PUBLICATION (Month and year of publication of document.) August 2012	6a. NO. OF PAGES (Total containing information. Include Annexes, Appendices, etc.) 66	6b. NO. OF REFS (Total cited in document.) 127

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

Defence R&D Canada – Ottawa
3701 Carling Avenue, Ottawa, Ontario, Canada K1A 0Z4

9a. PROJECT NO. (The applicable research and development project number under which the document was written. Please specify whether project or grant.) 13oc02	9b. GRANT OR CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.) DRDC Ottawa TM 2012-084	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

(X) Unlimited distribution
(   ) Defence departments and defence contractors; further distribution only as approved
(   ) Defence departments and Canadian defence contractors; further distribution only as approved
(   ) Government departments and agencies; further distribution only as approved
(   ) Defence departments; further distribution only as approved
(   ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected.)

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

The use of computer games as a means of training has seen considerable growth in industry and the military. In particular, the military has found computer games to provide a cost-effective, realistic, and safe environment to augment actual field exercises. Computer games are also easily made available to individuals or small teams. A challenge to the use of computer games for training is that computer games must frequently be populated with computer generated allies and enemies known commercially as *non-player characters* (NPCs). The quality of the training or game experience will therefore be dependent on the quality of the non-player characters who should demonstrate realistic human-like behaviour. This paper surveys the current state-of-the-art of the application of artificial intelligence techniques to create realistic behaviour in computer generated characters in games. The survey considers both commercial and military computer games in order to establish the state of the technology. Conclusions are drawn on the maturity and applicability of the surveyed approaches and highlights directions for future research.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Artificial Intelligence
Games
AI
Survey

Defence R&D Canada

Canada's leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE DÉFENSE

www.drdc-rddc.gc.ca