

Jimmy: Searching for Free-Will (A Competition)

Simon EGERTON^{a1}, Marc DAVIES^b, Brian JOHNSON^c, Victor CALLAGHAN^b
^aMonash University (Sunway Campus), ^bEssex University, ³Intel Corp

Abstract. In this invited paper we describe a competition that aims to engage a wide section of society in the exploration of the future nature of technology. In particular the competition will examine one of the most controversial aspects of Artificial Intelligence, specifically the concept of freewill. What is freewill? Can machines achieve it? What are the consequences for society should that happen? These are just some of the tantalising questions that flow from such a vision. For this competition, we roll back to a more basic stage and explore the basic issues of trying to construct a robot controller that displays behaviour that is analogous to freewill. We do this by offering a set of online tools that enable participants in the competition to create a bespoke controller for a simulated robot “Jimmy”; a waiter working in a futuristic bar set in a virtual world. This paper explains the motivation, design, implementation and rules for this competition

Keywords. Science fiction prototyping, futurology, robotics, quantum computing, artificial intelligence, virtual worlds

Introduction - What is a SF Prototype?

A Science Fiction (SF) prototype uses science fiction based explicitly on science fact as a design tool in the development of technology. Through traditional research and development we begin to define and understand *what* a technology is (Figure 1). This is the typical work that is going on in industrial labs and universities all over the world. Usually this work continues iterating itself until the technology is refined to such a point that it can be productized to incorporated into an existing product.

The SF prototyping process adds a step this process. The SF prototype asks *how* this technology will be used. (Figure 1) It provides a virtual reality in which the implications, problems and benefits of the technology can be explored. This exploration could uncover both best case and worst case scenarios but it can also explore the subtleties of how people will use and interact with the technology.

The ultimate goal of a SF prototype is to provide a new perspective on the technology that is feeding into its development. From our use of SF prototypes, as described later in this paper, we have seen that they can also benefit its feeder technology by providing specific experiments and scenarios for the application of the technology, thus illuminating possible benefits or pitfalls. By identifying these elements the development of the actual technology can be broadened and sometimes accelerated.

¹ Corresponding Author: Simon Egerton, Monash University Sunway Campus Email: simon.egerton@monash.edu.my

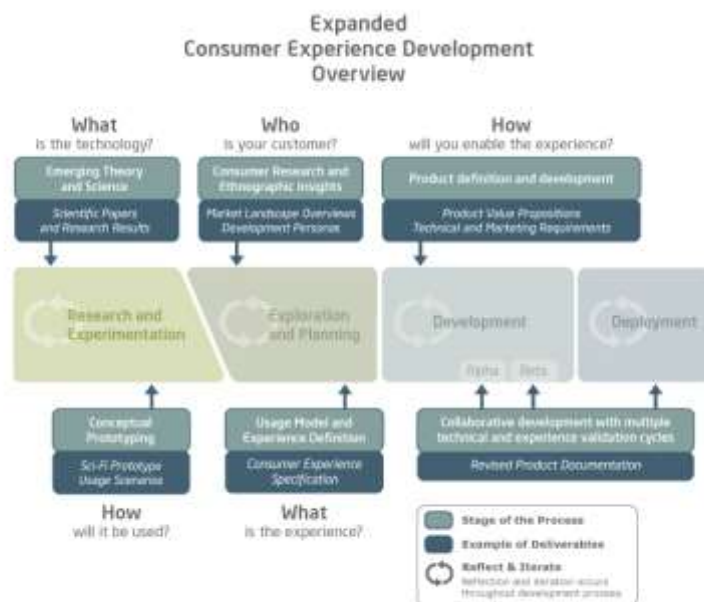


Figure 1. SF Prototyping process applied to a traditional product development process

1. Brain Machines – Birth of the Gin and Tonic Test

SF prototypes can be used in technology and product development. To illustrate this let's take a look at the work that we've been doing on AI for domestic robots in complex environments. The 'Dr. Simon Egerton Stories' are a series of stories that we've been using to develop theories of AI at the University of Essex in the UK and at Monash University in Kuala Lumpur Malaysia. Each of the SF prototypes in the series of stories allows the students, professors, computer scientists and engineers to explore the implications of their work on people, society and broader systems. At the moment we are about halfway through series and the results have been fascinating.

The second SF prototype in the Dr. Simon Egerton Stories was called *Brain Machines*. *Brain Machines* starts with the concept of AIs that could make both rational and irrational decisions and takes this idea a step further, exploring the idea of free will. First it examines the concept of free will as it applies to humans. It turns out that it's hotly debated if we humans even have free will. New neuroscience makes this a hotly debated topic. Second, the SF prototype explores how free will might be beneficial to AIs and robots; a scary notion for some and an intellectually challenging problem for others.

The scientific theories at play in *Brain Machines* come from three recent works. The first is a chapter from Michael Brooks' exceptional book *13 Things That Don't Make Sense* [3]. Chapter 11 is entitled, *Free Will – Your decisions are not your own*. In it Brooks does a brisk work of moving through a history of free will experimentation and the latest advances in neuroscience research. Ultimately he shows that science is proving that humans really don't have free will but that "for all practical purposes, it

makes sense to retain the illusion. Human consciousness, our sense of self and intention, may be nothing more than a by-product of being enormously complex machines that are our big-brained bodies, but it is a useful one, enabling us to deal with a complex environment.”[3]

The second work is a paper from Italian astrophysicist Paola A. Zizzi [17] called *I, Quantum Robot: Quantum Mind control on a Quantum Computer*. In the paper Zizzi explores using quantum metathought and metalanguage as a way to control robots or computers that could become self aware. Simply put, metathought is “the mental process of thinking about our own thought...the process of thinking about thinking.” Zizzi uses metalanguage to keep a robot from attaining free will. “With opportune boundary conditions, an apparently self-aware quantum robot reaches a level of thought. In this case the robot can still be controlled by a metalanguage which prevents him to reach the level of metathought.” [17] The goal of Zizzi’s theory is to keep a robot from attaining free will.

The third and final work that feeds into *Brain Machines* is the continuation of the research the SF prototype *Nebulous Mechanisms*, the paper *Instability and Irrationality: Destructive and Constructive Services within Intelligent Environments* by Simon Egerton, Victor Callaghan, Victor Zamudio and Graham Clarke [11]. This work explores the role of multiple personalities in an artificial intelligence (AI) and both the positive and negative effects of instability and irrationality on the system. The paper asks, “Does chance have a role in intelligent environments? ... chance and non-deterministic behaviour can play a fundamental and important role in intelligent environments...Underpinning our ideas is the view that intelligent environments may be seen as a complex system of interacting services...such complex systems can produce unexpected interactions that cause unplanned and often undesirable instabilities. However, not all instabilities are undesirable and in the second half of this paper, we present a conceptual notion that views system instability as a form of irrationality and propose a quantum control model for service agents within smart environments. We conjecture that irrational control models enable the service agents to perform better than if they were using traditional, rational, control models.” (Egerton, Callaghan, Zamudio and Clarke 2009)[11]

Arguably the greatest creation of the Dr. Simon Egerton Stories of SF prototypes is the character of *Jimmy*. Not only does *Jimmy* embody the three scientific inputs from Brooks, Zizzi, Egerton, Callaghan, Zamudio and Clarke but he also provides the scientist a way to explore the implications of their AI approach.

After the development of *Brain Machines*, the SF prototype was submitted to the development team that was working on Egerton, Callaghan, Zamudio and Clarke’s theories. After reading of the story the team came back to and replied that the scenario where the AI of *Jimmy* is making cocktails is a perfect test for the collection of ideas of the story. The scenario gives us a way to test the quantum transfer block. We need to build *Jimmy*.

And with that *Brain Machines* had achieved its goal as a SF prototype. It had taken the emerging science of Brooks and Zizzi and explores the implications of Egerton, Callaghan, Zamudio and Clarke’s novel approach to AI. In the exploration it has synthesized the ideas into a single experiment that could test the theories expressed in the original paper. The scene where *Jimmy* is challenged to make multiple gin and tonics, provided the scientists with a scenarios that they could build and test. Thus the “Gin and Tonic Test” was born.

2. The Competition Tools – What is an iWorld?

The term iWorld was adopted during another ongoing research project, where several virtual worlds of varying size and complexity have been created to provide infrastructure for the testing and evaluation in the areas of intelligent environments and computer games [7] [8] [9]. To clarify, ‘iWorld’ is a generic term used to describe a virtual world that has been augmented to allow some or all of its contents to possess the functionality commonly found in physical intelligent environments. This could potentially include intelligent devices, linked together on a pseudo-network, that can have their state or settings modified remotely by external computer programs and/or actuators, which can be used as interface tools by users to produce stimuli in an environment. Figure 2 shows some examples of existing iWorlds already in use. Simulated devices in an iWorld are designed to be accessed and/or controlled by computer programs via a set of associated UPnP wrappers, replicating a mechanism used by real intelligent environments [7] [9].



Figure 2. Screenshots of other existing iWorlds

iWorlds are highly adaptable, providing an easily portable and highly visual means of demonstrating the various concepts of intelligent environments research, (including visualising the ‘hidden computer’ premise [10]), which can be difficult to convey when using physical devices and real buildings. Additionally, multi-user iWorlds offer a potential platform for collaborative research. As space is usually only limited by the level of computer resources available to host the virtual environment, an iWorld can potentially include multiple instances of a single building design that can each be assigned to a different research group and customised to suit the individual needs of their particular projects. Each building could be treated independently, so no changes made would cause conflicts with the research being performed by another team. However, as each of the buildings will be located in the same iWorld, different groups could still observe the results being gathered by their peers, with several groups potentially working together on a common project.

Existing iWorlds can so far be placed into one of four categories; a) Two dimensional bespoke virtual intelligent environments, programmed using simple low detail graphics [8]; b) Three dimensional virtual intelligent environments created by modifying existing off-the-shelf simulation software, (i.e. a commercial computer game) [8]; c) Three dimensional virtual intelligent environments built using open source computer games technologies and graphics software [7], and; d) Mixed reality environments produced by augmenting physical devices in a real world intelligent space [5] with virtual counterparts in a category c iWorld [9].

2.2 *Building an iWorld*

Simple two dimensional iWorlds have been produced using the basic graphics toolkit provided by the Java programming language. However the quality of the environments produced is significantly limited, (e.g. different intelligent devices were represented by simple rectangles of different sizes or colours). More realistic visualisations where different objects could easily be identified from each other by users (i.e. at a glance), are provided by more advanced iWorlds, which use three dimensional graphics models to represent individual devices and other content. In addition to the higher quality of material content, three dimensional virtual intelligent environments are also able to represent human avatars in a much more realistic manner, using models that can be tailored to reflect individual appearances via a number of customisable features, (e.g. gender, clothing, height, weight, etc).

To build a bespoke virtual intelligent environment using three dimensional graphics would ordinarily take a large effort, with a significant amount of programming and time required to develop even a simple world. Existing three dimensional iWorlds have been created by using open-source virtual world toolkits such as RealXtend [16] or OpenWonderland [15]. Worlds can be built using the existing architecture as a base and adding additional middleware programming to produce intelligent functionality. Any necessary bespoke graphical content for the virtual world, (e.g. building models) can be produced using open-source graphics editing software, such as Google's SketchUp. However many realistic models, especially for material content such as furniture and devices can be obtained from online repositories, such as Google's 3D Warehouse [12].

Creating multi-user virtual intelligent environments can also be much easier with certain open-source simulation software, due to the inherited features of projects based around Massive Multiuser Online (MMO) virtual worlds. Such technologies are already in common use by the computer games industry for example when creating Massively Multiplayer Online Role Playing Game (MMORPG) titles, such as Blizzard Entertainment's World of Warcraft [2] series, and online communities such as Linden Lab's Second Life [14]. Several of the available open source software packages, enabling the creation of virtual worlds, incorporate MMO technologies. For example, RealXtend was create from a derivative of Open Simulator (a.k.a. Open Sim), an ongoing project which started by modifying an earlier version of the code used for the Second Life online community. Open Wonderland, (originally known as Project Wonderland), incorporates the Darkstar game server, originally developed to provide a graphics and physics engine for online computer games.

2.2 *Building Jimmy*

During the '*Jimmy Competition*', the strategy for evaluating robot controllers submitted by participants will include a performance assessment, which will take place in a bespoke iWorld system, consisting of a three-dimensional virtual intelligent environment with MMO capabilities.

The virtual world itself will consist of a large central area leading to approximately twenty individual rooms, outfitted to resemble a typical bar environment. Each bar

room will include at least, a single *Jimmy* robot, a bar area (similar to that shown in figure 3), with a drink dispenser (replicator) capable of producing several types of beverage and the expected tables and chairs. Figure 4 shows a screenshot of the under construction competition iWorld.




Figure 3. Early Virtual Bar Area Design




Figure 4. Under construction version of the competition iWorld

The replicator and selected tables in the environment are intended to be intelligent devices, which amongst other things provide controllers for the *Jimmy* robot with data relating to the current state of associated features in the world, (e.g. the number of drinks served). Several other intelligent devices will also feature in each bar environment, such as a music player, video-screen, lights and curtains, each of which will be designed to provide additional interactive functionality for users of the virtual world and increase the realism of the bars.

Throughout the competition *Jimmy* controller designs will be submitted for consideration to an online code repository by participants. Periodically a number of these programs will be selected from the repository and loaded into the created iWorld system. Each controller program loaded into the iWorld will be used to control a single *Jimmy* robot in one bar scene. Being an MMO virtual world users will be able to log into the environment via a link on the website of the Creative Science Foundation

(CSf) once they have registered either as a competition participant or an evaluator (i.e. somebody who isn't planning to submit a controller for consideration). Each user who accesses the virtual world first appears in the central arrivals area. From that location they're free to walk around the environment, interacting with any intelligent devices they encounter. The central area provides a means for users to access each of the different bar scenes, where *Jimmy* controllers are being evaluated. In each bar users will either be able to directly command the *Jimmy* robot to perform a series of serving tasks, or watch others who choose to adopt this role. Evaluators will be asked to observe and rate the performance of each *Jimmy* robot they encounter throughout the virtual world. The primary factor the controllers will be evaluated by is how much 'freewill' each program can allow the *Jimmy* robot to exhibit to a user. For example, if a number of pointless or repetitive orders are given to the *Jimmy* robot would it be capable of realising the futility and reacting accordingly?

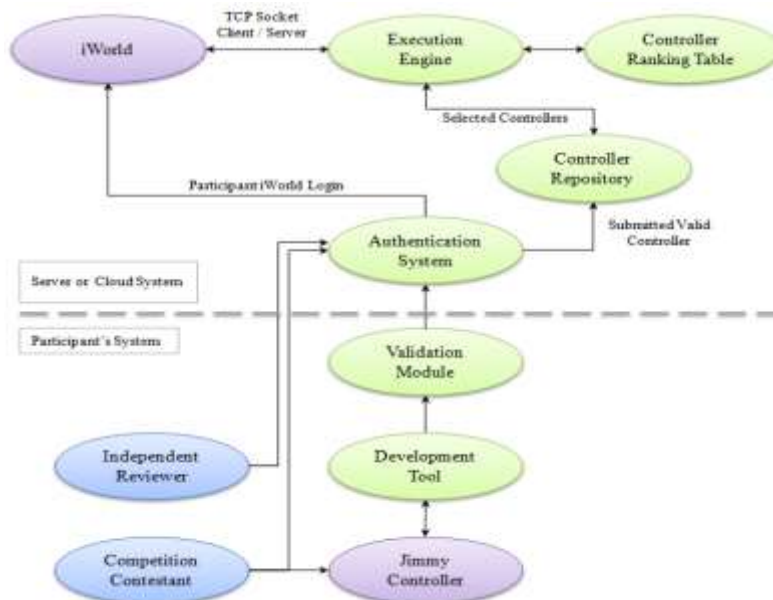


Figure 5. Proposed system architecture.

3. The Competition Server Architecture

Figure 5 illustrates the competition server architecture. The *iWorld* and execution system are supported by significant computing resources, whether a high-end server or from within a cloud based system. Participants make use of a development toolkit to design, test and validate their controllers before submitting them to the competition system through a web-based authentication portal.

The execution-engine currently supports the Java runtime environment. The *Jimmy* controllers are written in Java and interact with the execution engine via a network socket forming a client (the *Jimmy* controller) / server (the execution-engine)

relationship. This client / server arrangement is intended to allow future expansion of the supported runtime environments. Future versions will support C/C++, PROLOG and LISP runtimes.

Controllers submitted to the competition system will be stored in an online repository. At set intervals the execution engine will select a number of controllers from the repository for evaluation. The number of stored controllers selected will be equal to half of the total number of bars featured in the iWorld environment. Each controller is assigned to control a single *Jimmy* robot in one of the iWorld bars. Scores allocated to individual controllers during evaluation are used to rank each program, which is subsequently stored in a ranking table. The other half of the bars featured in the iWorld, are each populated by the best ranked controllers currently available.

The validation module and development tools enable contestants to develop *Jimmy* controllers on their host systems, both can be downloaded from the Creative Science Foundation website. The development tool is a Java API package which implements the Simon *Jimmy* Interface (SJI) used in the competition scenario. The SJI interface is explained further below. The validation module contains a Java applet server which runs on the host system. The server accepts local *Jimmy* controller connections via a socket connection and is designed to receive and exercise all of the messages used in the competition scenario.

3.1. Operation Scenario Example

From the perspective of a contestant, the process of designing a *Jimmy* controller and submitting the controller for evaluation would be seen as the following series of steps;

1. Download the development kit containing the development API and validation software. Both are used to design and implement controller programs.
2. Upon the creation of a valid controller program, (i.e. it works with all the messages sent and received by the validation module), the controller can be submitted to the competition system for consideration by uploading the relevant file(s) via the competition web portal.
 - a. As an additional level of validation, upon making a valid submission, a request will be made to complete a short questionnaire, describing the controller program and its expected functionality. (Note: This step may not be included in the initial release).
 - b. The system will send a notification to each of the authors of the newbie controller requesting their participation in an online evaluation.
3. Contestants login to the iWorld and wander around each the bars, where it is possible to either act in the Simon role, ordering the resident anonymized *Jimmy* controllers around or observe another user performing this function. Contestants are required to either score each of the *Jimmy* controllers independently or select which they feel were the top ten.

To encourage contestants to participate in the ranking, the system may penalise contestants who do not participate in the online evaluation process, either by allocating a point reduction in the score awarded to their controller or removing it from the

competition completely. To facilitate wider participation additional ‘guest’ reviewers may also enter the iWorld and evaluate *Jimmy* controllers using the above strategy.

The competition system will periodically feedback the current ranking table to the contestants, with their controllers ranking revealed. This information can be used to make changes to the *Jimmy* controller, either in an attempt to climb the rankings or maintain a lead.

4. System Communication

The system is architected around a client / server model, this allows it to afford some level of protection against rough controllers, and also enables the system to easily be extended to support other execution environments.

The communication system used within the competition scenario describes the Simon *Jimmy* Interface, specifically the set of messages which are exchanged between the two main actors, Simon and *Jimmy*. Messages can also be passed between these actors and other intelligent devices within an iWorld bar. The messages are arranged into several categories as defined in Table 1 below. A full specification is available from the Creative Science Foundation website.

Simon – to – Jimmy
<ul style="list-style-type: none"> • Request Drink, (if more than one is available specify which type) • Reject Incorrect Drink, (i.e. wrong size or type) • Change Drink
Simon – to – Table
<ul style="list-style-type: none"> • Count Available Delivered Drinks • Count Consumed Drinks • Count All Glasses & Bottles • Consume Drink
Jimmy – to – Replicator
<ul style="list-style-type: none"> • Make Drink • Dispose of Drink • Determine Available Drinks
Jimmy – to – Table
<ul style="list-style-type: none"> • Count Available Delivered Drinks • Count Consumed Drinks • Count All Glasses and Bottles • Place Drink • Remove Empty Glasses and Bottles • Remove Unconsumed Drink
Jimmy – to – Simon
<ul style="list-style-type: none"> • Serve Drink • Determine Number of Drinks Served • Determine Number of Drinks Consumed
Replicator – to – Jimmy
<ul style="list-style-type: none"> • Produced Drink

Table 1. Message categories in the initial release of the Simon *Jimmy* Interface (SJI). Actions within each category are summarised.

The replicator is an intelligent device that can produce several different types of drink, (e.g. gin & tonic, brandy, whisky, etc), and in various quantities, (e.g. small, standard or large). Occasionally the replicator may get an order wrong and produce a different drink to what was requested by *Jimmy*. Another test for the controller could

be to see whether *Jimmy* realises he has been given the wrong drink by the replicator immediately or tries to deliver it to Simon who subsequently rejects it.

Later releases of the system will extend the richness of the messages, to make the environment more interesting from the controllers and competitors perspective. As an example for future improvements, the Simon-to-*Jimmy* category could be extended with messages to “dim/raise the lights”, “open/close window blinds or curtains”, “change a background music track” or “change the video displayed on a wall screen”. The intention is to provide more environment interaction possibilities and to give *Jimmy* a richer vocabulary with which to display a richer set of behaviours, making the experience for the competitors more engaging.

5. Controller Design

The competition allows us to examine one of the more controversial aspects of artificial intelligence, freewill, achieved through the design and implementation of the *Jimmy* controllers. The goal for the *Jimmy* controller is to display behaviours that are indicative of freewill, leading to favourable ranking with the competition when assessed against the subjective, objective and technical assessments of the controller. However, what freewill is, and how artificial intelligence can achieve it are open questions and ones we would like to address. We hope the competition will lead us to some interesting insights or even answers, emerging through the collective efforts of the competition entries.

The controllers will be evaluated using a three-tiered strategy. One of these evaluations objectively measures *Jimmy*’s free-will. We extend the SJI interface with five additional acknowledgements as described in Table 2. This mechanism was inspired by the systems used in some role-playing computer games. Additional levels could be added if more acknowledgements are needed. Objectively, the level of freewill exhibited by *Jimmy* is measured by the range of responses he uses when in conversation with Simon.

<i>Jimmy – to - Simon</i>		
Return Message Acknowledgments	Level of Free Will	Score
“Coming right up!”	None	-2
“Yes sir”, (not quite an enthusiastic as above)	Low	-1
“Null”, (<i>Jimmy</i> says nothing)	Neutral	0
“Apology”	Medium	+1
“Question / Refuse”, (request)	High	+2

Table 2. Enhancements to the SJI enabling an objective measurement of *Jimmy*’s free-will.

The acknowledgments from table 2 could be used in some interesting ways, for example, the apology acknowledgement could be issued in response to the “Reject Drink” action that is performed if *Jimmy* brings Simon a drink different to the one he specified. However, the environment also includes the possibility of certain drinks being unavailable, *Jimmy* could also use the apology acknowledgement if he knows it is impossible for him to fill a specific order rather than robotically trying to complete the request, failing and subsequently returning to Simon with an incorrect drink.

While the above acknowledgements provide a mechanism for quantifying free-will, they potentially reduce the *Jimmy* Controller to a small set of simple rules. An obvious solution example is given in Figure 6 below.

```

function Jimmy-Controller-Trivial
  static: threshold
  repeat
    WAIT(Simon-to-Jimmy: Request Drink)
    if (Jimmy-to-Table: Available Delivered Drinks greater than threshold)
      Jimmy-to-Simon: Question
    else
      Jimmy-to-Simon: Coming Right up
      Jimmy-to-Replicator: Make Drink
      Jimmy-to-Simon: Serve Drink
  
```

Figure 6. A trivial *Jimmy* Controller using set thresholds to simulate free-will against the objective evaluation criteria.

While this trivial controller may evaluate favourably against the objective criteria and the rule set could be extended to score well against the subjective evaluations, it would gain a low score in the technical evaluation, there are far more interesting possibilities. By way of an example, Figure 7 describes a generalised control model developed in [11] which could be applied to the *Jimmy* Controller.

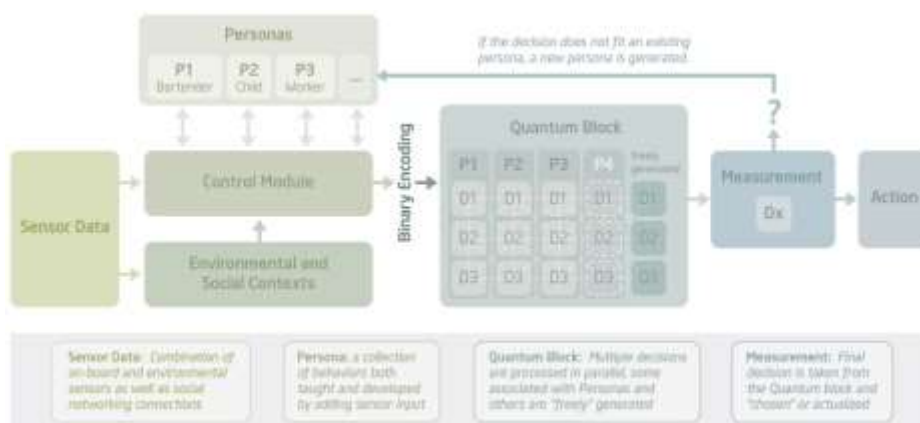


Figure 7. Behavioural control model with persona enhancements and quantum transfer decision logic.

This controller is based on a behaviour based architecture, as originally proposed by Rodney Brookes in 1985 [4]. The basic architecture is extended with ideas borrowed from psychoanalytic theory [6], behaviours are dynamically generated [13] and collected together into distinct personas which are activated by some ‘unseen’ arbiter according to contextual changes. The ‘unseen’ arbiter is implemented here in Quantum logic, the notion is that the effects of quantum entanglement allow for probabilistic interference [18] and also break the link between cause and effect, one of key arguments for free-will [3]. A simplified *Jimmy* Controller based on this architecture is illustrated in Figure 8.

```
function Jimmy-Controller-Quantum
  persona: Bartender, starts a persona process, in this case, Bartender
  persona: Worker
  persona: Child, this persona is responsible for questioning, asking 'why?'
  repeat
    WAIT(Simon-to-Jimmy: Request Drink)
    Context ← UPDATE
    Activate ← QLOGIC(Worker, Bartender, Child, Context)
    ACTIVATE_PERSONA(Activate)
  end
persona Bartender
  repeat
    Jimmy-to-Simon: Coming Right up
    Jimmy-to-Replicator: Make Drink
    Jimmy-to-Simon: Serve Drink
  end
persona Worker
  repeat
    if (Context: waiter)
      Jimmy-to-Simon: Yes Sir
    end
persona Child
  repeat
    if (Context: waiter)
      Jimmy-to-Simon: Question
    end
End
```

Figure 8. A simplified *Jimmy* Controller using quantum logic to decide which persona *Jimmy* should activate, or whether to generate a new persona

The quantum logic implements the persona transfer function, a logic block of quantum gates in some arrangement which take the context vector as input and outputs a persona activation decision. We are currently experimenting with various simulated quantum logic block configurations and will make the controller available when we have a complete design.

6. Controller Evaluation

Controllers will be evaluated using a three-tiered strategy involving a blind review accompanied by independent quantitative and technical evaluations.

- The blind review uses contestants and optionally, sometimes, invited experts to evaluate the controllers in the iWorld environment.
- The quantitative evaluation analyses the performance of the controllers using the evaluation system described above.
- The technical evaluation will be based on the responses to a questionnaire completed by contestants when submitting their controller, which will be analysed by expert reviewers, who allocate a technical merit score.

7. Free Will – A Discussion

The SF prototype *Brian Machines* is all about free will and free will is a tricky thing. For humans it seems it is at the centre of what makes us human. It is fundamental for us to function in our cultures, societies and governments. “Free will goes to the centre of our sense of self, our autonomy as human beings. Strip us of it, and we are nothing more than animals.” [3] But when you apply the notion of free will to machines you get a dramatically different affect; fear. Free will in a robot, computer or a non-human is seen as dangerous. A quantum robot with free will is a threat to humanity. Quantum robots “might even become self-aware, conscious and have *free will*. This will be the sign that the technological singularity has been reached. Such a singularity might be very dangerous if quantum robots decide to act against human beings and take advantage of them.” [17] Obviously when it comes to free will there are two overriding assumptions. The first is that humans must possess free will or at worst they must maintain their delusion of free will to operate in complex societies and environments. “In the illusion of free will, it seems we have been equipped with a neurological slight of hand that, while contra-rational, helps us deal with a complex social and physical environment.”[3] The second assumption is that machines, robots and computers must never develop free will. A machine that can think and act for itself strikes fear into the hearts of many scientists and science fiction fans alike. The big worry is that when machines get smarter than humans then they will take over our role as top dog here on Earth. I’d argue this is more a reflection of human tendencies rather than the ultimate goals of robots. Humans like being in charge and we assume that everybody must want to take it from us.

The American scientific icon Isaac Asimov had an interesting take on this back in 1977 when he was writing his non-fiction robot series for American Airlines magazine. “But if computers become more intelligent than human beings, might they not replace us? Well shouldn’t they? They may be as kind as they are intelligent and just let us dwindle by attrition. They might keep some of us as pets, or on reservations. Then too, consider what we’re doing to ourselves right now – to all living things and to the very planet we live on. Maybe it is *time* we were replaced. Maybe the real danger is that computers won’t be developed to the point of replacing us fast enough.”[1]

8. Summary

This invited paper has been designed to highlight the architectural model for the initial release of an international competition designed to explore and evaluate the concept of ‘freewill’ in artificial intelligence. A rule-set has also been established for the competition, designed to guide participants when creating their respective controller programs, which will be used to control a generic robot waiter, (named *Jimmy*) and allow him to perform his role. Part of the competition procedure takes place in an MMO environment (iWorld), where a number of rooms, each with the guise of a futuristic bar, will be used to host a *Jimmy* robot and evaluate submitted controllers. A number of other ‘intelligent devices’ will also be present in the iWorld, which can each be interacted with by participants accessing the environment. The quest of whether or not free-will exists, or can ever be artificially programmed into machines whilst fascinating is secondary to our main goal in writing this paper which is to engage as wide a section of the public as is possible in thinking about the future and becoming

involved in the intellectually stimulating fields of science and engineering which strives towards explaining the mysteries of our existence and shapes all our lives. The journey to science and engineering, for many of us, started by reading science fiction, so its hardly surprising we find it a powerful tool for driving our research forward, a passion we hope to pass on to others through this effort.

Acknowledgements

We are indebted to Intel Corp and to the Creative Science Foundation for financially supporting this competition.

References

- [1] Asimov I. *Our Intelligent Tools*. American Airlines, (1977)
- [2] Blizzard Entertainment, World of Warcraft: The Burning Crusade, <http://sea.blizzard.com/en-gb/games/burningcrusade/>, Retrieved 22nd March 2011.
- [3] Brooks M, 13 Things That Don't Make Sense, Vintage, (2009)
- [4] Brooks, R. A. "A Robust Layered Control System for a Mobile Robot", IEEE Journal of Robotics and Automation, Vol. 2, No. 1, March 1986, pp. 14–23; also MIT AI Memo 864, September 1985.
- [5] Callaghan V., Clark G., Colley M., Hagrais H., Chin J., Doctor F., "Intelligent inhabited environments," BT Technology Journal , vol. 22, no. 3, Klywer Academic Publishers, Dordrecht, Netherlands, July 2004.
- [6] Clarke GS, Personal Relations Theory: Fairbairn, Macmurray and Suttie. London: Routledge 2006 Chapter 3
- [7] Davies M., Callaghan V., Gardner M., "Towards a mixed reality intelligent campus", Proceedings of the 2008 international conference on Intelligent Environments (IE '08), Seattle, USA.
- [8] Davies M., Callaghan V., Shen L., "Modelling pervasive environments using bespoke & commercial game-based simulators", Proceedings of the 2007 international conference on Life System Modelling and Simulation (LSMS '07), Shanghai China, Springer, September 2007.
- [9] Davies M., Callaghan V., "iWorlds: Building Mixed Reality Intelligent Environments using Customisable 3D Virtual Worlds," The 6th IEEE International Conference on Intelligent Environments (IE '10), Kular Lumpur, Malaysia, 19th – 21st July 2010.
- [10] Dooley J., Davies M., Ball M., Callaghan V., "Declocking Big Brother", The 6th IEEE International Conference on Intelligent Environments (IE '10), Kula Lumpur, Malaysia, 19th – 21st July 2010.
- [11] Egerton S, Callaghan V and Clarke G "Instability and Irrationality: Destructive and Constructive Services within Intelligent Environments", 5th International Conference on Intelligent Environments, (IE'09), Barcelona, Spain, July 2010
- [12] Google, 3D Warehouse, <http://sketchup.google.com/3dwarehouse>, Retrieved: 26th February 2011.
- [13] Hagrais,H.A.K., Callaghan,Vic, Colley,M.J., 'On Line Calibration of the Sensors Fuzzy Membership Functions in Autonomous Mobile Robots', IEEE International Conference on Robotics and Automation, San Francisco, USA, (2000)
- [14] Linden Lab., Second Life, <http://secondlife.com/>, Retrieved: 24th February 2011.
- [15] Open Wonderland, Open source 3D virtual collaboration toolkit | Open Wonderland, "<http://openwonderland.org>", Retrieved: 1st March 2011.
- [16] RealXtend, realXtend – Open source platform for interconnected virtual worlds, <http://www.realxtend.org>, Retrieved: 26th February 2011.
- [17] Zizzi P.A., I, Quantum Robot: Quantum Mind Control on a Quantum Computer, Journal of Physics Conf. Ser., 67, 012045, (2008), p2
- [18] Jerome R. Busemeyer, Zheng Wang, James T., Townsend, "Quantum dynamics of human decision-making", Journal of Mathematical Psychology 50 (2006) 220–241