

Instinct: A Biologically Inspired Reactive Planner for Embedded Environments

Robert H. Wortham, Swen E. Gaudl, Joanna J. Bryson

Dept of Computer Science, University of Bath
Claverton Down, Bath, BA2 7AY, UK

Abstract

The Instinct Planner is a new biologically inspired reactive planner, based on an established behaviour based robotics methodology and its reactive planner component — the POSH planner implementation. It includes several significant enhancements that facilitate plan design and runtime debugging. It has been specifically designed for low power processors and has a tiny memory footprint. Written in C++, it runs efficiently on both ARDUINO (ATMEL AVR) and MICROSOFT VC++ environments and has been deployed within a low cost maker robot to study AI Transparency. Plans may be authored using a variety of tools including a promising visual design language, currently implemented using the DIA drawing package.

INTRODUCTION

From the 1950's through to the 1980's the study of embodied AI assumed a cognitive symbolic planning model for robotic systems — SMPA (Sense Model Plan Act) — the most well known example of this being the Shakey robot project (Nilsson, 1984). In this model the world is first sensed and a model of the world is constructed within the AI. Based on this model and the objectives of the AI, a plan is constructed to achieve the goals of the robot. Only then does the robot act. Although this idea seemed logical and initially attractive, it was found to be quite inadequate for complex, real world environments. Generally the world cannot be fully modelled until the robot plan is underway, since sensing the world requires moving through it. Also, where environments change faster than the rate at which the robot can complete its SMPA cycle, the planning simply cannot keep up. Brooks (1995) provides a more comprehensive history, which are not repeated here.

In the 1990's Rodney Brooks and others (Breazeal and Scassellati, 2002) introduced the then radical idea that it was possible to have intelligence without representation (Brooks, 1991). Brooks developed his subsumption architecture as a pattern for the design of intelligent embodied systems that have no internal representation of their environment, and minimal internal state. These autonomous agents could traverse difficult terrain on insect-like legs, appear to interact socially with humans through shared attention and gaze tracking, and in many ways appeared to possess behaviours

similar to that observed in animals. However, the systems produced by Brooks and his colleagues could only respond immediately to stimuli from the world. They had no means of focusing attention on a specific goal or of executing complex sequences of actions to achieve more complex behaviours. The original restrictions imposed by Brooks' subsumption architecture were subsequently relaxed with later augmentations such as timers, effectively beginning the transition to systems that used internal state in addition to sensory input in order to determine behaviour.

Following in-depth studies of animals such as gulls in their natural environment, ideas of how animals perform action selection were originally formulated by Nico Tinbergen and other early ethologists (Tinbergen, 1951; Tinbergen and Falkus, 1970). Reactions are based on pre-determined drives and competences, but depend also on the internal state of the organism (Bryson, 2000). Bryson (2001) harnessed these ideas to achieve a major step forwards with the POSH (Parallel Ordered Slipstack Hierarchy) reactive planner and the BOD (Behaviour Oriented Design) methodology, both of which are strongly biologically inspired. It is important to understand the rationale behind biologically inspired reactive planning. It is based on the idea that biological organisms constantly sense the world, and generally react quickly to sensory input, based on a hierarchical set of behaviours structured as Drives, Competences and Action Patterns. Their reactive plan uses a combination of sensory inputs and internal priorities to determine which plan elements to execute, ultimately resulting in the execution of leaf nodes in the plan, which in turn execute real world actions. For further reading see Gurney, Prescott, and Redgrave (1998), Prescott, Bryson, and Seth (2007) and Seth (2007).

At run-time, the reactive plan itself is essentially fixed. Various slower reacting systems may also be used to modify priorities or other parameters within the plan. These slower reacting systems might be compared with emotional or endocrinal states in nature that similarly affect reactive priorities (Gaudl and Bryson, 2014). Similarly the perception of senses can be affected by the internal state of the plan, an example being the latching (or hysteresis) associated with sensing (Rohlfshagen and Bryson, 2010).

In nature, the reactive plan is subject to possible learning that may change the plan parameters or even modify the

structure of the plan itself as new skills and behaviours are learned. This learning may take place ontogenetically, i.e. within the lifetime of an individual, or phylogenetically, by the process of natural selection, across the lifetimes of many individuals. Bryson’s BOD approach suggests that humans provide most of the necessary learning in order to improve the plan over time, in place of natural selection. However, Gaudl (manuscript in preparation) successfully uses genetic algorithms to automate part of this learning process, albeit within a computer game simulation.

A reactive plan is re-evaluated on every plan cycle, usually many times every second, and this requires that the inquiries from the planner to the senses and the invocation of actions should respond quickly. This enables the reactive plan to respond quickly to changes in the external environment, whilst the plan hierarchy allows for complex sequences of behaviours to be executed. Applying these ideas to robots we can see that for senses, this might imply some caching of sense data. For actions, it also implies that long running tasks (relative to the rate of plan execution), need to not only return success or failure, but also another status to indicate that the action is still in progress and the plan must wait at its current execution step before moving on to its next step. The action may be executing on another thread, or may just be being sampled when the call to the action is made. This is implementation specific and does not affect the functioning of the planner itself. If re-invoked before it completes, the action immediately returns an In-Progress response. In this way, longer running action invocations do not block the planner from responding to other stimuli that may still change the focus of attention by, for example, releasing another higher priority Drive.

Each call to the planner within the overall scheduling loop of the robot starts a new plan cycle. In this context an action may be a simple primitive, or may be part of a more complex pre-defined behaviour module, such as a mapping or trajectory calculation subsystem. It is important to note that the BOD methodology does not predicate that all intelligence is concentrated within the planner. Whilst the planner drives action selection, considerable complexity can still exist in sensory, actuation and other probabilistic or state based subsystems within the overall agent (Bryson, 2001).

The computer games industry has advanced the use of AI for the simulation of non player characters (Lim, Baumgarten, and Colton, 2010). Behaviour trees are similarly hierarchical to POSH plans, but have additional elements that more easily allow logical operations such as AND, OR, XOR and NOT to be included within the plan. For example it is possible for a goal to be reached by successfully executing only one of a number of behaviours, trying each in turn until one is successful. Bryson’s original design of POSH does not easily allow for this kind of plan structure.

Behaviour trees are in turn simplifications of Hierarchical Task Network (or HTN) planners (Ghallab, Nau, and Traverso, 2004). Like POSH, HTN planners are able to create and run plans that contain recursive loops, meaning that they can represent any computable algorithm. An interesting parallel can be drawn here with Complexity theory. Holland (2014) argues that a Complex Agent System (CAS) is often

characterized by the fact that it can be decomposed into a set of hierarchical layers, with each layer being Turing complete. For a biological entity Holland identifies these layers as existing at the levels of DNA, organelle, cell, organ, organism and social levels. For an artificial agent we can identify these layers as computer hardware, operating system, application programming language, reactive planner, plan, agent and social levels. Thus we can argue that to create an artificial agent truly capable of emergent implicit behaviour, we should strive to ensure that the Planner on which its behaviour depends should be Turing complete, particularly allowing looping and recursion.

THE INSTINCT PLANNER

The Instinct Planner is a reactive planner based on Bryson’s POSH (Bryson, 2008, 2001). It includes several enhancements taken from more recent papers extending POSH (Rohlfshagen and Bryson, 2010; Gaudl and Bryson, 2014), together with some ideas from other planning approaches, notably Behaviour Trees (BT — Lim, Baumgarten, and Colton, 2010). A POSH plan consists of a *Drive Collection (DC)* containing one or more *Drives*. Each *Drive (D)* has a priority and a releaser. When the *Drive* is released as a result of sensory input, a hierarchical plan of *Competences, Action Patterns* and *Actions* follows.

- *Action Pattern (AP)*: Action patterns are used to reduce the computational complexity of search within the plan space and to allow a coordinated fixed sequential execution of a set of elements. An action pattern— $AP = [\alpha_0, \dots, \alpha_k]$ —is an ordered set of Actions that does not use internal precondition or additional perceptual information. It provides the simplest plan structure in POSH and allows for the optimised execution of behaviours. An example would be an agent that always shouts and moves its hand upwards when touching an hot object. In this case, there is no need for an additional check between the two Action primitives if the agent should always behave in that manner. APs execute all child elements before completing.
- *Competence (C)*: Competences form the core part of POSH plans. A competence $C = [c_0, \dots, c_j]$ is a self-contained basic reactive plan (BRP) where $c_b = [\pi, \rho, \alpha, \eta]$, $b \in [0, \dots, j]$ are tuples containing π , ρ , α and η : the priority, precondition, child node of C and maximum number of retries. The priority determines which of the child elements to execute, selecting the one with the highest priority first. The precondition is a concatenated set of senses that either release or inhibit the child node α . The child node itself can be another Competence or an Action or Action Pattern. To allow for noisy environments a child node can fail a number of times, specified using η , before the Competence ignores the child node for remaining time within the current cycle. A Competence sequentially executes its hierarchically organised child-nodes where the highest priority node is the competence goal. A Competence fails if no child can execute or if an executed child fails.

- *Drive (D)*: A Drive— $D = [\pi, \rho, \alpha, A, v]$ —allows for the design and pursuit of a specific behaviour as it maintains its execution state. The Drive Collection determines which Drive receives attention based on each Drive's π , the associated priority of a Drive. ρ is the *releaser*, a set of preconditions using senses to determine if the drive should be pursued. α is either an Action, Action Pattern or a Competence and A is the root link to the Drive Collection. The last parameter v specifies the execution frequency, allowing POSH to limit the rate at which the Drive can be executed. This allows for coarse grain concurrency of Drive execution (see below).
- *Drive Collection (DC)*: The Drive Collection—DC—is the root node of the plan— $DC = [g, D_0, \dots, D_i]$. It contains a set of Drives $D_a, a \in [0 \dots i]$ and is responsible for giving attention to the highest priority Drive. To allow the agent to shift and focus attention, only one Drive can be active in any given cycle. Due to the parallel hierarchical structure, Drives and their sub-trees can be in different states of execution. This allows for cooperative multitasking and a quasi-parallel pursuit of multiple behaviours at the Drive Collection level.

For a full description of the POSH reactive planner see Bryson (2001).

Enhancements and Innovations

The Instinct Planner includes a full implementation of what we term *Drive Execution Optimisation (DEO)*. DEO avoids a full search of the plan tree at every plan cycle which would be expensive. It also maintains focus on the task at hand. This corresponds loosely to the function of consciousness attention seen in nature (Bryson, 2011). A form of this was in Bryson's original POSH, but has not been fully implemented in subsequent versions. The Drive, Competence and Action Pattern elements each contain a *Runtime Element ID*. These variables are fundamental to the plan operation. Initially they do not point to any plan element. However, when a Drive is released the plan is traversed to the point where either an Action is executed, or the plan fails at some point in the hierarchy. If the plan element is not yet completed it returns a status of In Progress and the IDs of the last successful steps in the plan are stored in Runtime Element ID in the Drive, Competence and Action Pattern elements. If an action or other sub element of the plan returns success, then the next step in the plan is stored. On the next cycle of the drive, the plan hierarchy is traversed again but continues from where it got to last plan cycle, guided by the Runtime Element IDs. A check is made that the releasers are still activated (meaning that the plan steps are still valid for execution), and then the plan steps are executed. If a real world action fails, or the releaser check fails, then the Runtime Element ID is once again cleared. During execution of an Action Pattern (a relatively quick sequence of actions), sensory input is temporarily ignored immediately above the level of the Action Pattern. This more closely corresponds to the reflex behaviour seen in nature. Once the system has started to act, then it continues until the Action Pattern completes, or an element in the Action Pattern explicitly fails. Action

Patterns are therefore not designed to include Actions with long running primitive behaviours.

In addition to these smaller changes there are three major innovations in the Instinct Planner that increase the range of plan design options available to developers:

- Firstly, the idea of runtime alteration of drive priority. This implementation closely follows the RAMP model of Gaudl and Bryson (2014) which in turn is biologically inspired, based on spreading activation in neural networks. Within the Instinct Planner we term this *Dynamic Drive Reprioritisation (DDR)*. DDR is useful to modify the priority of drives based on more slowly changing stimuli, either external or internal. For example, a recharge battery drive might be used to direct a robot back to its charging station when the battery level becomes low. Normally this drive might have a medium priority, such that if only low priority drives are active then it will return when its battery becomes discharged to say 50%. However, if there are constantly high priority drives active, then the battery level might reach a critical level of say 10%. At that point the recharge battery drive must take highest priority. A comparison can be drawn here with the need for an animal to consume food. Once it is starving the drive to eat assumes a much higher priority than when the animal experiences normal levels of hunger. For example, it will take more risks to eat, rather than flee from predators.
- Secondly, the idea of flexible latching provides for a more dynamic form of sense hysteresis, based not only on plan configuration, but also the runtime focus of the plan. This implementation follows the work of Rohlfschagen and Bryson (2010). Within the Instinct Planner we term it *Flexible Sense Hysteresis (FSH)*. This hysteresis primarily allows for noise from sensors and from the world, but Rohlfschagens paper also has some basis in biology to avoid dithering by prolonging behaviours once they have begun. If the Drive is interrupted by one of a higher priority, then when the sense is again checked, it will be the Sense Flex Latch Hysteresis that will be applied, rather than the Sense Hysteresis.
- Thirdly, we enhance the Competences within the plan, such that it is possible to group a number of competence steps by giving them the same priority. We refer to this as a *priority group*. Items within a group have no defined order. Within a priority group, the Competence itself can specify whether the items must all be successfully executed for the Competence to be successful (the AND behaviour), or whether only one item need be successful (the OR behaviour). In the case of the OR behaviour, several items within the group may be attempted and may fail, before one succeeds. At this point the Competence will then move on to higher priority items during subsequent plan cycles. A Competence can have any number of priority groups within it, but all are constrained to be either AND or OR, based on the configuration of the Competence itself. This single enhancement, whilst sounding straightforward, increases the complexity of the planner code significantly, but allows for much more compact plans, with a richer level of functionality achievable within a single

Competence than was provided with the earlier POSH implementations.

Multi Platform

The Instinct planner itself is able to run both within MICROSOFT VISUAL C++ and the ARDUINO development environments (Arduino, 2016) as a C++ library. The ARDUINO uses the ATMEL AVR C++ COMPILER (Atmel Corporation, 2016a) with the AVR LIBC library (Atmel Corporation, 2016b) — a standards based implementation of `gcc` and `libc`. This arrangement harnesses the power of the VISUAL C++ Integrated Development Environment (IDE) and debugger, hugely increasing productivity when developing for the ARDUINO platform, which has no debugger and only a rudimentary IDE. We have a complete implementation of the Instinct Planner on an ARDUINO based robot named R5. The robot runs using various test plans, see figure 1. Due to the very compact memory architecture of Instinct, the planner is able to store plans with up to 255 elements within the very limited 8KB memory (RAM) available on the ARDUINO MEGA (ATMEL AVR ATMEGA2560 MICROCONTROLLER). The 255 element limitation arises from the use of a single byte to store plan element IDs within the ARDUINO environment.

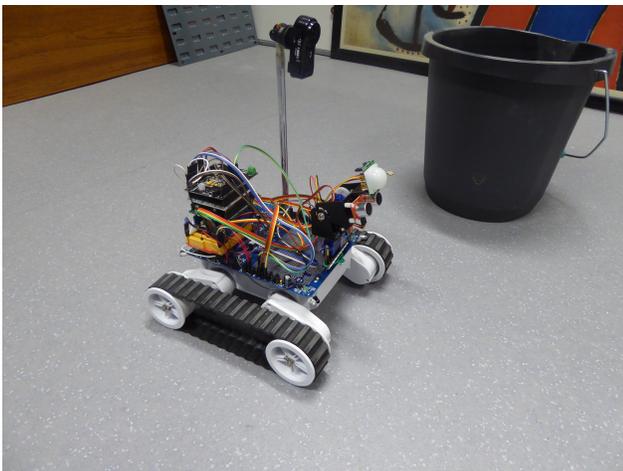


Figure 1: The R5 ARDUINO based Maker Robot in a laboratory test environment. The camera mounted on the robot is used to record robot activity, but is not used by the robot itself.

The robot itself has active infrared and ultrasonic distance sensors, a head capable of scanning its environment, a passive infrared (PIR) sensor to assist in the detection of humans interacting with it, and proprioceptive sensing of odometry (distance travelled) and drive motor current. It has simple and more complex underlying behaviours that can be invoked by the planner, such as the ability to turn in the direction of the most clear pathway ahead, or to use its head to scan for the presence of a human. It also has a multicoloured headlight that may be used for signalling to humans around it. Finally, it has an electronically erasable programmable read only memory (EEPROM) that permanently stores both the robot's configuration parameters and the Instinct plan.

This leverages the planner's ability to serialise plans as a byte stream, and then reconstitute the plan from that stream at startup.

Memory Management

In order to produce a planner that operates effectively in an environment with severely limited working memory resources (RAM), considerable design effort has been applied to the memory management architecture within the planner. There are 6 separate memory buffers, each holding fixed record length elements for each element type in the plan — Drives, Competences, Competence Elements, Action Patterns, Action Pattern Elements and Actions. An instance of Instinct has a single Drive Collection — the root of the plan.

Within each plan element, individual bytes are divided into bit fields for boolean values, and the data is normalised across elements to avoid variable length records. This means, for example, that Competence Elements hold the ID of their parent Competence, but the Competence itself does not hold the IDs of each of its child Competence Elements. At runtime a search must be carried out to identify which Competence Elements belong to a given Competence. Thus, the planner sacrifices some search time in return for a considerably more compact memory representation. Fortunately this search is very fast, since the Competence Elements are stored within a single memory buffer with fixed length records. Testing shows the time taken by this searching was negligible in comparison with the plan cycle rate of the robot.

Plan elements, senses and actions are referenced by unique numeric IDs, rather than by name. The memory storage of these IDs is defined within the code using the C++ `#typedef` preprocessor command, so that the width of these IDs can be configured at compile time, depending on the maximum ID value to be stored. This again saves memory in an environment where every byte counts. Consideration of stack usage is also important, and temporary buffers and similar structures are kept to a minimum to avoid stack overflow.

Fixed strings (for example error messages) and other data defined within programs are usually also stored within working memory. Within a microcontroller environment such as ARDUINO this is wasteful of the limited memory resource. This problem has been eliminated in the Instinct Planner implementation by use of AVR LIBC functions (Atmel Corporation, 2016b) that enable fixed data to be stored in the much larger program (flash) memory. For code compatibility these functions have been replicated in a pass-through library so that the code compiles unaltered on non-microcontroller platforms.

Instinct Testing Environment

As a means to test the functionality of the Instinct Planner within a sophisticated debugging environment, we have an implementation of the planner within MICROSOFT VISUAL C++, and have tested a very simple simulation of a robot within a grid based world. The world allows multiple robots to roam, encountering one another, walls and so on. This could be extended in future, with a graphical user interface

to better show both the world and the real time monitoring available from within the plan. However our current research focusses on the real time debugging of actual robots (Wortham, Theodorou, and Bryson, 2016). Building transparency into robot action selection can help users build a more accurate understanding of the robot, see below.

The Instinct Planner code is not fundamentally limited to 255 plan elements, and will support much larger plans on platforms with more memory. In MICROSOFT VISUAL C++ for example, plans with up to 65,535 nodes are supported, simply by redefining the `instinctID` type from `unsigned char` to `unsigned int`.

Instinct Transparency Enhancements

The planner has the ability to report its activity as it runs, by means of callback functions to a monitor C++ class. There are six separate callbacks monitoring the Execution, Success, Failure, Error and In-Progress status events, and the Sense activity of each plan element. In the VISUAL C++ implementation, these callbacks write log information to files on disk, one per robot instance. This facilitates the testing and debugging of the planner. In the ARDUINO robot, the callbacks write textual data to a TCP/IP stream over a wireless (wifi) link. A JAVA based Instinct Server receives this information, enriches it by replacing element IDs with element names, and logs the data to disk. This communication channel also allows for commands to be sent to the robot while it is running.

With all nodes reporting all monitor events over wifi, a plan cycle rate of 20Hz is sustainable. By reducing the level of monitoring, we reduce the volume of data sent over wifi and plan cycle rates of up to 40Hz are achievable. In practice a slower rate is likely to be adequate to control a robot, and will reduce the volume of data requiring subsequent processing. In our experiments a plan cycle rate of 8Hz was generally used.

Figure 2 shows how the planner sits within the robot software environment and communicates with the Instinct Server.

Instinct Command Set

The robot command set primarily communicates with the planner which in turn has a wide range of commands, allowing the plan to be uploaded and altered in real time, and also controlling the level of activity reporting from each node in the plan. When the robot first connects to the Instinct Server, the plan and monitoring control commands are automatically sent to the robot, and this process can be repeated at any time while the robot is running. This allows plans to be quickly modified without requiring any re-programming or physical interference with the robot.

Creating Reactive Plans with iVDL

POSH plans are written in a LISP like notation, either using a text editor, or the ABODE editor (Brom et al., 2006; Bryson, 2013). However, Instinct plans are written very differently, because they must use a much more compact notation and they use IDs rather than names for plan elements,

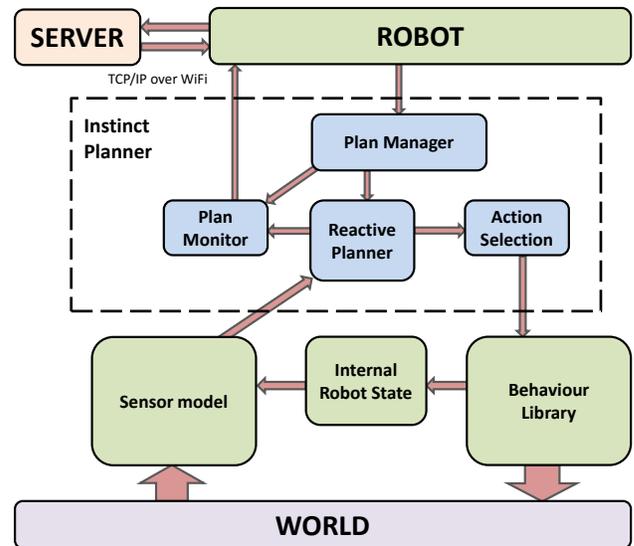


Figure 2: Software Architecture of the R5 Robot showing interfaces with the World and the Instinct Server. The Instinct Planner provides the action selection subsystem of the robot.

senses and actions. We have developed the *Instinct Visual Design Language (iVDL)* based on the ubiquitous Unified Modelling Language (UML) notation. UML is supported by many drawing packages and we have developed a simple PYTHON export script to allow plans to be created graphically within the DIA drawing tool (Macke, 2014). The export script takes care of creating unique IDs and allows the plans to use named elements, thus increasing readability. The names are exported alongside the plan, and whilst they are ignored by the planner itself, the Instinct-Server uses this export to convert IDs back into names within the log files and interactive display.

Figure 3 shows the Instinct plan template within Dia. We use the UML class notation to define classes for the six types of element within the Instinct plan, and also to map the external numerical identifiers (IDs) for senses and robot actions to names. We use the UML aggregation connector to identify the connections between the plan elements. This can be read, for example, as “A Drive can invoke an Action, a Competence or an Action Pattern”.

Figure 4 shows a plan for the R5 robot. At this level of magnification the element details are not legible, but this screen shot gives an impression of how plans can be laid out.

This particular plan searches the robot’s environment, avoiding objects and adjusting its speed according to the space around it. As it moves around it attempts to detect humans within the environment. The robot also temporarily shuts down in the event of motor overload, and it will periodically hibernate when not in open space to conserve battery power. Such a plan might be used to patrol hazardous areas such as industrial food freezers, or nuclear facilities.

The plan was designed and debugged within the space of a week. During the debugging, the availability of the trans-

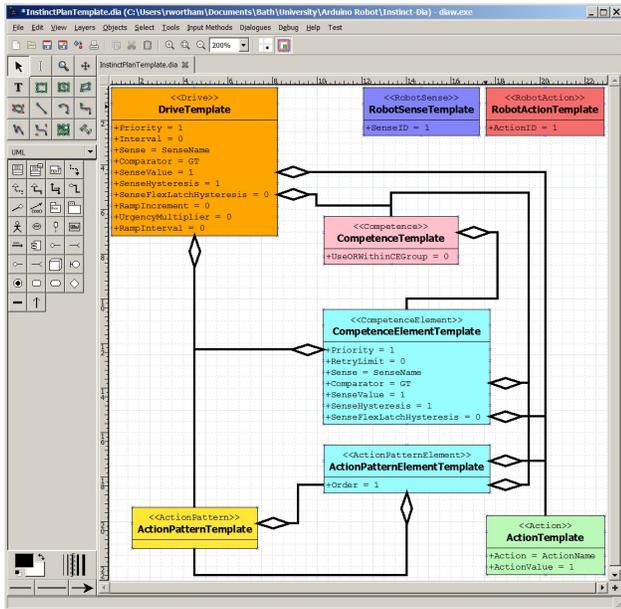


Figure 3: Instinct Plan element types and their relationship, shown within the DIA drawing tool.

parency data logged by the Instinct Server was extremely useful, because mere observation of the robot’s emergent behaviour is frequently insufficient to determine the cause of plan malfunction.

The actual positioning of plan elements within the drawing is entirely up to the plan designer. Since Dia is a general purpose graphical editor, other symbols, text and images can be freely added to the file. This is useful at design time and during the debugging of the robot. It also provides an additional vehicle for the creation of longer term project documentation. We suggest that an in-house standard is developed for the layout of plans within a development group, such that developers can easily read one another’s plans.

Plan Debugging and Transparency

Currently, work is underway within the Artificial Models of Natural Intelligence (AmonI) research group at the University of Bath¹ to create a new version of the ABODE plan editor (Theodorou, Wortham, and Bryson, 2016). This version directly writes Instinct plans, and also reads the real-time transparency data emanating from the Instinct Planner, in order to provide a real-time graphical display of plan execution. In this way we are beginning to explore both runtime debugging and wider issues of AI Transparency.

CONCLUSIONS AND FURTHER WORK

The Instinct planner is the first major re-engineering of Bryson’s original work for several years, and allows deployment in practical real time physical environments such as our ARDUINO based maker robot.

¹AmonI — <http://www.cs.bath.ac.uk/ai/AmonI.html>

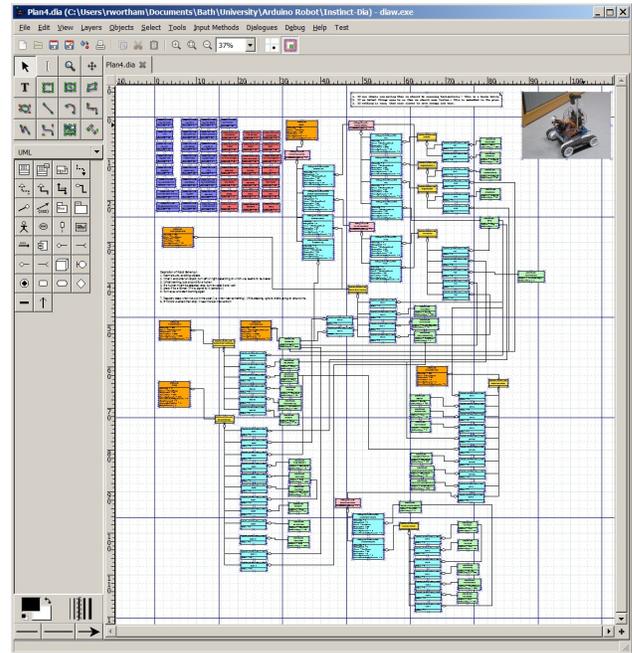


Figure 4: The plan used by the R5 robot to enable it to explore an environment, avoid obstacles, and search for humans. The plan also includes emergency behaviours to detect and avoid excessive motor load, and to conserve battery by sleeping periodically.

By using a very lean coding style and efficient memory management, we maximise the size of plan that can be dynamically loaded and the performance in terms of execution rate.

The transparency capabilities, novel to this implementation of POSH, provides the necessary infrastructure to deliver real time plan debugging. Work is currently underway to leverage this architecture with a real time visual debugging tool, initially to assist the work of reactive plan designers, but also as a research tool for the investigation of wider AI Transparency issues.

The Visual Design Language (iVDL) is a novel representation of reactive plans, and we demonstrate that such plans can be designed using a standard drawing package and exported with a straightforward plug-in script. We envisage the development of similar plug-ins for other drawing tools such as MICROSOFT VISIO.

Although primarily developed for physical robot implementations, the Instinct Planner has obvious applications in teaching, simulation and game AI environments. We envisage extending the current Instinct Testing Environment to provide a richer, GUI based test platform for Instinct, and for use as a teaching tool to teach the concepts of reactive planning in general and the Instinct Planner in particular.

Finally, we would like to see the implementation of Instinct on other embedded and low cost Linux computing environments such as the RASPBERRY PI (Raspberry Pi Foundation, 2016). With more powerful platforms such as the PI, much larger plans can be developed and we can test both the runtime performance of very large plans, and the design

efficiency of iVDL with multi-user teams.

References

- Arduino. 2016. *Arduino Website*. <https://www.arduino.cc/>.
- Atmel Corporation. 2016a. *Atmel Studio Website*. <http://www.atmel.com/Microsite/atmel-studio/>.
- Atmel Corporation. 2016b. *AVR Libc Reference Manual*. <http://www.atmel.com/webdoc/avrlibcreferencemanual/>.
- Breazeal, C., and Scassellati, B. 2002. Robots that imitate humans. *Trends in Cognitive Sciences* 6(11):481–487.
- Brom, C.; Gemrot, J.; Bida, M.; Burkert, O.; Partington, S. J.; and Bryson, J. J. 2006. POSH Tools for Game Agent Development by Students and Non-Programmers. In *The Ninth International Computer Games Conference: AI, Mobile, Educational and Serious Games.*, 1–8. Bath, UK: The University of Bath.
- Brooks, R. a. 1991. Intelligence Without Representation. *Artificial Intelligence* 47(1):139–159.
- Brooks, R. 1995. Intelligence Without Reason. In Steels, L., and Brooks, R. R. A., eds., *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*. Mahwah, New Jersey, USA: L. Erlbaum Associates. 25–81.
- Bryson, J. J. 2000. The study of sequential and hierarchical organisation of behaviour via artificial mechanisms of action selection. M.Phil. Thesis, University of Edinburgh.
- Bryson, J. J. 2001. *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*. Ph.D. Dissertation, MIT, Department of EECS, Cambridge, MA. AI Technical Report 2001-003.
- Bryson, J. 2008. *Parallel-rooted, Ordered Slip-stack Hierarchy*. <http://www.cs.bath.ac.uk/~jjb/web/posh.html>.
- Bryson, J. J. 2011. A Role for Consciousness in Action Selection. In Chrisley, R.; Clowes, R.; and Torrance, S., eds., *Proceedings of the AISB 2011 Symposium: Machine Consciousness*, 15—20. York: SSAISB.
- Bryson, J. J. 2013. *Advanced Behavior Oriented Design Environment (ABODE)*. <http://www.cs.bath.ac.uk/~jjb/web/BOD/abode.html>.
- Gaudl, S., and Bryson, J. J. 2014. The Extended Ramp Goal Module: Low-Cost Behaviour Arbitration for Real-Time Controllers based on Biological Models of Dopamine Cells. *Computational Intelligence in Games 2014*.
- Ghallab, M.; Nau, D. S.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. Morgan Kaufmann Series in Artificial Intelligence. Elsevier/Morgan Kaufmann.
- Gurney, K. N.; Prescott, T. J.; and Redgrave, P. 1998. The Basal Ganglia viewed as an Action Selection Device. In *Eighth International Conference on Artificial Neural Networks*, 1033–1038. London, UK: Springer.
- Holland, J. H. 2014. *Complexity: A Very Short Introduction*. Oxford University Press.
- Lim, C. U.; Baumgarten, R.; and Colton, S. 2010. Evolving behaviour trees for the commercial game DEFCON. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 6024 LNCS(PART 1):100–110.
- Macke, S. 2014. *Dia Diagram Editor*. <http://dia-installer.de/>.
- Nilsson, N. J. 1984. Shakey the Robot. Technical report, SRI International, Technical Note 323.
- Prescott, T. J.; Bryson, J. J.; and Seth, A. K. 2007. Introduction. Modelling natural action selection. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 362(1485):1521–9.
- Raspberry Pi Foundation. 2016. *Raspberry Pi Website*. <https://www.raspberrypi.org/>.
- Rohlfshagen, P., and Bryson, J. J. 2010. Flexible Latching: A Biologically-Inspired Mechanism for Improving the Management of Homeostatic Goals. *Cognitive Computation* 2(3):230–241.
- Seth, A. K. 2007. The ecology of action selection: insights from artificial life. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 362(1485):1545–1558.
- Theodorou, A.; Wortham, R. H.; and Bryson, J. J. 2016. Why is my robot behaving like that ? Designing transparency for real time inspection of autonomous robots. In *EPSRC Principles of Robotics Workshop, Proceedings of the AISB 2016 Annual Conference {accepted for publication}*.
- Tinbergen, N., and Falkus, H. 1970. *Signals for Survival*. Oxford: Clarendon Press.
- Tinbergen, N. 1951. *The Study of Instinct*. Oxford, UK: Oxford University Press.
- Wortham, R. H.; Theodorou, A.; and Bryson, J. J. 2016. What Does the Robot Think ? Transparency as a Fundamental Design Requirement for Intelligent Systems. In *IJCAI-2016 Ethics for Artificial Intelligence Workshop {accepted for publication}*.