Information-Oriented Design and Game AI

Sebastian Welsh
University of Technology, Sydney
Sebastian.J.Welsh@alumni.uts.edu.au

Yusuf Pisan
FIT, University of Technology, Sydney
PO Box 123, Broadway NSW 2007, Australia
Phone: +61 2 9514 4478
Fax: +61 2 9514 4535
ypisan@it.uts.edu.au

ABSTRACT
Information is a resource that every AI relies on to operate effectively. Although information influences the capabilities and performance of AIs, it is not treated as a design issue. Changing the information available to an AI can potentially enhance or cripple its performance. More direct evaluation of issues such as information selection and acquisition can improve the performance of existing AI implementations. The influence of information on AI behaviour is examined with an emphasis on virtual worlds and game AI, as this domain provides both an effective research environment and opportunities for tangible improvements to AI behaviour.

1. INTRODUCTION
For some reason, AI practitioners spend a lot of time worrying about how to make decisions while taking the material decisions are based on for granted. Without information that is relevant and comprehensible even the most sophisticated AI is unable to act effectively. AIs are not just decision systems. They are a fusion of information and information processing resources. Worrying about how to make decisions without considering what information is most appropriate for delivering the desired outcomes is putting the cart before the horse. The starting point of good AI is information, not computation.

AIs are designed. They are built by intentionally selecting resources and tools and combining them to produce a system that generates coherent actions in an environment. The information available to an AI is part of this process of selection. The design process means that AIs reflect many aspects of the designer's own intelligence. The process of selecting and integrating components is a distillation of judgements about salience, utility and significance. Part of this process is designing how inputs to the system are handled. Another part of the process is selecting what information is available to the system.

Thinking about information is not a foreign issue in AI, or to the broader study of intelligent behaviour across the Cognitive Sciences. The problem is that issues of information are often treated implicitly. We contend that AI is as much an issue of information design as it is an issue of technical design, and for this reason it is important to examine information explicitly. Most examinations have focussed on how it is used, manipulated and communicated. While these issues are relevant, the designed nature of AI makes it necessary to consider what information should be generated.

1.1 Why information is important
Intelligent behaviour benefits from good information. There is a relationship between sensory and behavioural sophistication in living creatures. This suggests that getting good information is also an issue for AI. Many creatures that display significant behavioural sophistication possess refined sensory systems that parallel the complexity with which they can interact with the environment. A large part of the human brain is dedicated to sensory mechanisms. The advantage that effective sensory systems offer to help creatures interact effectively with their environment is evident in the repeated, independent evolution of the eye [4]. An effective means of acquiring information is a critical part of the tool-set of an intelligent entity. This is just as true for AI systems as organic creatures. Giving AIs good information helps them respond more appropriately, effectively and hopefully will provide a means of fostering more intelligent behaviour.

The main informational issue in designing AIs is selecting what information they should use. This incorporates other issues such as sensing, but the critical issue is deciding what information needs to fulfil its purpose. Sensing involves collecting individual pieces of information, but individual elements are not always independently meaningful. Wilson [24] describes the problem of perception as *"knowing what in the environment is relevantly the case"* (p5), which encompasses more than acquisition mechanisms. The central issue is designing this perceptual mechanism. This involves identifying information that articulates with what an AI is able to do, what it is supposed to do and when it should do it.

We believe this problem of information-oriented design is underrepresented in AI methodologies. There is an underlying supposition that human designers identify the right information to provide to an AI. This may often be true, but because AIs are ultimately dependent on the information they receive, providing poor information can unnecessarily degrade AI performance. It is possible to improve the quality of information AIs receive through consideration of several issues. These include what types of information exist, what factors influence the efficiency of information, and the theoretical implications of handling information in particular ways.

1.2 Applicability to virtual worlds
AIs in virtual worlds provide an excellent opportunity to investigate these issues. Virtual worlds represent a micro-

cosm of the broader domain of AI research. One advantage of investigating this information relationship in a virtual environment is that both the AI and the environment are explicable. Virtual worlds are self-contained information domains, so the only information available to AIs bounded by the components of the world. Consequently, the entire environment is potentially under experimental control. By focusing on how good information might enhance this particular domain, it is possible to identify some of the issues that exist in the relationship between AIs and the information they receive.

Emphasising information as a design component of AI does not require new decision making methodologies. It is possible to make incremental improvements by enhancing existing approaches. This is desirable for game AI, as implementations frequently use established approaches for practical reasons. Concerns from the game industry are that solutions need to be efficient, usable, and portable between games. Nareyek [12] has argued that game AIs require known capabilities, resource efficiency and predictable behaviour to fit the needs of game designers. Industry preference for established technologies are reflected in efforts to define a standardised AI architecture [13]. These approaches are still dependent on information, so they might be improved by providing practitioners with better tools to assess what information AIs require.

At the same time, Fairclough et al [6] argue that commercial game AI is an under-resourced and under-prioritised component in comparison to graphical richness. It fails to explore the pool of advanced techniques derived from research, and often takes procedural shortcuts such as cheating to obviate the need for complexity in the AI system. These are hardly desirable traits in AI, though they may be necessary in commercial AI. However, it is possible to enhance AI operation without foregoing the control required by commercial implementations. One approach is enhancing the information available to AIs.

Our perspective is that within game AI, performance is ultimately about what the AI does and how easily it does it. The problem domain for a game AI is not necessarily the same as for human players, for the reasons outlined by Nareyek. However the AI still operates in a problem domain: one that is proximate to the domain occupied by a player but also incorporates issues of believability, story line and game play. The AI must be able to act effectively in this environment, so identifying the most appropriate information to provide the AI is an important performance issue.

Despite the need for AI to support wider game play issues, internal behavioural consistency is an important part of AI performance. It is also the 'presentation layer' of the AI, in that a human player observes the actions of the AI in the game environment. Believability in the game world is important, so for that reason we have directed our attention towards how information relates to realistic actions in virtual worlds. We discuss one such approach in Section 5.

Enhancing the information used by AIs has already met with success in both commercial and academic implementations. A range of approaches have manipulated the information resources available to an AI, producing significant changes in behaviour. These include having the AI generate predictive information in order to permit anticipatory planning [9] and imbuing the environment with high level suggestions for actions and responses [5]. Each of these approaches is notable because they fundamentally change the information an AI bases its actions on. These alterations are not limited to academic investigations. Games such as The Sims [25] have used an annotation like approach, while other games have used approaches such as attaching action information to objects and terrain. The diversity of these approaches suggests that altering information not so much a methodology as a way of inviting better methodologies.

While these approaches have been beneficial, they invite deeper examination of how information can improve AI. They emphasise providing particular types of information to AI but represent a subset of a larger problem, which is how to provide information that best supports what an AI is supposed to do. When what matters is how an AI behaves, it is worthwhile considering how to improve how an it can act by improving the information available to it. This is as relevant to established approaches as it is to more novel approaches. Any attempt to investigate how information influences AI raises some interesting theoretical issues. While we will not dwell on these matters here, we describe them briefly as the assumptions about information influence the way it is identified, selected and implemented.

1.3 Theoretical Issues

Trying to identify what information should be available to an AI is complex as solutions that are obvious to the designer may not suit the needs of an AI. Designers select information from many potential sources when making a choice about what information to make available to an AI. The designer also selects from many ways of acquiring that information. The nature of the selection process is irrelevant, no matter whether it is based on a fitness function or explicit inclusion. What matters is that some piece of information is selected in preference over other information that may be available to the system.

Information is functionally separate from its implementation. Implementations that handle a particular piece of information may vary in attributes such as efficiency, comprehensibility and ease of implementation. If these implementations produce the same results, they are instances of a function with multiple realisability. If the two functions are dissimilar and thereby produce differing results, they are still using the same information. When implementations of an information processing functions have differing inputs, the same information is merely encoded differently. Selecting a means of providing or processing information is not the same as choosing the information itself.

Providing information of differing degrees of complexity influences how much intelligence is an AI's and how much intelligence is the designer's. Providing an AI with high level information can create the appearance of behavioural sophistication, but the underlying AI may be rudimentary. In comparison, an AI that uses rudimentary information to achieve the same results is actually generating complex behaviour from simple inputs. The act of selecting what information should be available to an AI is, to some extent, imbuing some of the designer's intelligence in the AI [3]. For the first AI, the high level information is actually a result of the designers intelligence. For the second AI there is still designer influence, but this is less direct. The apparent behavioural sophistication of the first AI is superficial as there is a one to one correspondence between the information and the behaviour, whereas the second AI is generative.

The issues raised here are relevant to general problem of AI design. As our current emphasis is an examination of how information can influence AI performance in games and virtual worlds we cannot do these, and other, issues justice. From a game AI perspective what matters is behavioural and computational performance, rather than the integrity of the AI. Nevertheless, considering the influences on and implications of how information is selected is a part of the design process.

2. A SIMPLE EXAMPLE

The influence of information is easily illustrated by detailing how varying information can change the capabilities of a simple AI. The following scenario explores how information availability influences AIs with rudimentary behavioural capabilities. The relationship between what an AI is able to do and what influence information has on these capabilities is illustrated by evaluating what information is present in this simple environment. This is not a complex scenario, but this simplicity helps illustrate how the selection of information available to an entity can influence its actions.

Consider a 2 dimensional environment which contains obstacles and food sources of varying sizes. Within the environment are two types of animats: predator and prey, which must act, survive and perish according to simple rules. The actions available to all animats are *move*, *bite* and *breed*. Survival depends on health, which can be *replete*, *hungry*, *weak* or *dead*. Animats recover health by feeding, and lose health through conflict with each other. If two animats of the same type are both replete they can *breed* to produce a third animat. The only difference to these conditions is between what predators and prey eat, and when they gain nourishment from eating. For predators, each time a predator *bites* prey, the prey drops one health level immediately. A predator only eats if it kills prey. Prey eat each time they *bite* food objects.

There is a limited amount of information in the environment that they can use to decide what action to take. The only information they can acquire is about themselves and other animats (type, vector distance, health) and objects (type, size, vector distance). Providing animats with differing information can influence when they act, and how effectively. Some combinations of information are shown in Table 1.

Even a small amount of information is enough to equip the animat with basic behaviours. Clearly, no information reduces the animat to acting randomly, but other combination of behaviour show in Table 1 change the the capabilities of animats substantially:

Object or Animat: With this information make decisions about what actions to perform using rough heuristics. Actions are hit and miss. Predators can chase other animats, but risk chasing other predators as well as prey. Prey receive even less useful information. They can only avoid predators by fleeing from every animat, and cannot differentiate between food and non-food objects.

Table 1: Impact of available information

Type	Distance	Self
Nil	Nil	Nil
Object / Animat	Vector	Nil
Predator / Prey Obstacle / Food	Vector	Nil
Predator(+health) / Prey(+health) Obstacle(+size) / Food(+size)	Vector	Health

Animat and Object type: This information makes both categories of animat more effective. Prey need not flee from other prey, and predators need not pursue other predators. Additionally, prey can now distinguish between obstacles and food, providing them with the opportunity to *bite* at solid objects selectively.

Specific type plus health: This information permits increasingly appropriate action selection as animats receive more contextual cues. Adding information about the size of nearby meals enables animats to select larger meals over smaller ones. Adding health information permits animats to decide whether breeding is possible.

These changes in information availability expose two different types of information. Critical information determines whether an action is meaningful. It provides the receiver with insight about whether it *could* perform an action. Contextual information determines whether an action is appropriate, so it helps the receiver decide whether it *should* perform some action. Increasing the information available to animats results in additional refinements to their behavioural repertoire.

More information is not necessarily better. It is only of benefit if it is relevant to what they can or should do. For example, if animats are displayed using two different colours to for presentation purposes there is no benefit for animats to have this information. However if the rules are changes so that breeding requires that parents must be of different colours, colour information becomes valuable.

Animats don't need to directly sense information. Some information can be generated from more basic information in the environment. For example, animats that only sense the animat—object demarcation could use previous experience to generate information about animat and object types. This is suboptimal as it is expensive and fails to capitalise on information in the environment. It is possible to generate more useful information. Animats can pursue moving targets more efficiently by generating a predictive movement model[17]. This internally generated information does not directly exist in the environment. It is built from spatial information in conjunction with time information. Just like more sophisticated behaviours such as selective predation, it relies on basic information to work.

Deciding how much information to give an AI is an important decision. Too little information and the AI must guess what to do, too much and it will struggle to find what is relevant. Deciding what information to provide involves considering what an AI is supposed to do, how it is intended to do it, and how it acquires and encodes that information. Most of the time AIs only need to use a small portion of the range of available information. They occupy an information environment, which encompasses the possible meaningful input to a system, the resources available to acquire it, and the instantiation of it within an AI.

This simple scenario shows how altering the information environment of AIs can affect their behaviour. Varying what information is available doesn't change the actions they can perform: it changes the sophistication of when actions are applied, and the likelihood of success. Altering the information environment substantially changes their opportunities to use those actions. In extreme cases, it determines whether they are able to intentionally perform an action at all. Most of the time, it will change the efficacy and appropriateness of their actions.

3. TYPES OF INFORMATION

AIs must obtain information about differing aspects of their environment. Some things are transient, so information about it is worth having only as the transient event occurs. Other things are persevering, so it is continually useful to the AI. We have categorised this information into *features*, *events*, *state*, *relationships* and *world knowledge*. These demarcations were developed through analysis of specific scenarios [22], so they are prospective rather than prescriptive categories. What is important is that AIs need to acquire information of many kinds, so may benefit from using different approaches to acquire this information.

Feature information reflects the most obvious concept of sensing: identifying what objects exist in the environment. Features include things like entities, food sources, obstacles and other durable attributes of the environment. Identifying features through polling the environment is a common approach in virtual worlds, though not without computational load [16]. The features and objects in the environment constitute the largest part of the virtual world, no matter whether the world is a game or a simulation. Other characters, resources, objects and landscape form the central tableau for interaction. AIs have to know what is out there and how they can interact with it.

Event information is about noticing what is going on in the external environment and what is happening to the AI. Events may be external occurrences such as hearing sounds or internal events such as sensing damage. Event perception is temporally dependent as AIs need immediate notification to get maximum benefit. Event information is about immediate things, so it needs to be delivered quickly, and updated as close to real time as the environment permits. Event oriented sensing has been suggested as a way of reducing the cost of sensing [16]. For certain types of information it provides a useful tool. For other information tasks such as path-finding or looking for features no event perception is involved at all.

State information provides a self monitoring mechanism. An AI may need to know where it is located, its current goals, its inventory or its health. At a first glance this information may seem artificially distinct from events as anything that changes its state is an event! However, this information is potentially useful at times other than when it has immediately changed. In the earlier example, animats using information about their own strength, type or position are using state information. State information is useful, as it can be used to put other decision making mechanisms in context.

Relationship information is an important resource, but one that is difficult to provide. Relationships are the things an AI needs to see that result from interactions between the features of the world. These can include spatial relationships, temporal relationships and causal relationships. In the example prey animats could avoid predators more effectively by identifying line of sight and places of concealment, while predators could hunt more effectively by relating food location to prey location. Because relationships are a complex issue, one approach has been to circumvent the need for them. Often this circumvention involves specifying what to do in potentially interesting areas. For example, waypoints that have cues attached can be used to suggest what sort of actions are appropriate to actors [20].

World information describes elements such as the physics, object properties and success or failure conditions. This sort of information may assist AIs to plan and act appropriately. It is not necessary for AIs to understand their world, as this knowledge can be situated in their inherent capabilities. For example Goal Oriented Action Planning codifies selected world rules into an action framework, effectively ingraining knowledge of the world rules in the AIs planning mechanisms [15]. Other information like cultural knowledge, for example the presence of a secret door in a library, can be conveyed through annotations [5].

The information that an AI needs to receive can vary substantially depending on what it needs to do. Similarly, different mechanisms within the AI may rely on markedly different types of information. An AI may need to act reflexively to protect itself against immediate threats, and it may need to also possess long-term planning and predictive capabilities. Each of these elements reflect information resources that may differ substantially in form, durability, representation and means of acquisition.

4. INFORMATION ACQUISITION

The AI must be able to acquire information effectively and efficiently. How it is acquired influences how useful it will be to an AI. There are several efficiency issues in acquisition. One is the efficiency of the sensor, and the other is the efficiency of the information. In combination, these factors shape how well the information available to an AI fits the problem domain. Determining how an AI acquires information involves making decisions about what tradeoffs are worthwhile. These involve decisions about the accuracy of information and how it is represented as well as the speed, acuity and configuration of sensory mechanisms.

4.1 Information Integrity

AIs that sense efficiently are getting the best information at the lowest cost. Issues of efficiency, and sufficiency, are a powerful determinant in selecting what information should reach an AI as better information is likely to come at a cost. There are a number of elements that influence the efficiency and fidelity of information acquisition.

Cost: Different pieces of information have different levels of utility. Information may be too expensive to acquire, depending on what having the information enables the AI to do compared to what it costs to acquire it [11]. It may also be impossible to acquire desirable resources at a required fidelity. This cost of information includes both sensory and processing mechanisms as differing forms of information may invite different ways and costs of processing.

Another cost is how efficiently information is encoded. There are expensive and inexpensive ways of encoding the same information [18], so the transfer and processing of information from the environment to the AI can vary in efficiency because of the way it is encoded. Coordinating the way information is encoded and the way it is processed reduces the communication overhead for an AI, making information acquisition more efficient.

Fit: Information can vary in how effectively it supports a particular action. Fitness reflects the potential efficiency of information, independent of implementation. In an idealised implementation where the cost of providing and processing any piece of information is equivalent, differing pieces of information will vary in how well they fit the informational needs of an AI. In ideal circumstances the efficiency of information measures how comprehensively it supports a particular action. The more the information is generalised, the less accurately it informs a particular problem. However, these same processes make the information potentially useful to a wider range of actions. Often it is possible to make a piece of information more efficient overall by sacrificing some specific fit.

Accuracy: Accuracy reflects how closely instantiated information reflects the design ideal. The difference between fit and accuracy of information is the difference between how something could work and how it does work. Information that potentially fits the problem well may be difficult to acquire or work with in a practical way. One commonly recognised influence on accuracy is the acuity of the sensor, which is subject to all kinds of variation and inaccuracy in real world implementations [14].

Dynamism: Section 3 proposed different types of information that have variable update rates. Information sources change at different rates, so acquired information remains valid for variable durations. Some information is invariant, such as world rules. Other information is durable, such as the geography. Still more more information is reasonably variable, such as goals or inventory status, while other information changes rapidly, such as visible features. All of this information needs to be accessible to the actor whenever it is making a decision, but the actor should neither re-acquire information unnecessarily nor replicate information gathering across sensory mechanisms. Ensuring that perceptual mechanisms articulate with the mutability of different types of information has potential benefits for the efficiency of AI.

4.2 Sensory Mechanisms

Because sensing is computationally costly, there are advantages in acquiring information efficiently. We have proposed a number of types of information, each of which need to be made available to the AI as efficiently as possible. Different types of information may require different methods for acquisition, so sensory mechanisms need to accommodate a range of information gathering tasks

AIs can acquire the types of information detailed in 3 using a range of established sensing techniques: Feature information is often acquired through polling the environment; Event information frequently uses a notification or interrupt mechanism; State information is available to the AI through introspection; Relationship information can be embedded in the environment and treated as a feature; and world information can be embedded in the design of the decision making mechanisms. These approaches can be enhanced by considering how often particular types of information need to be sensed and sequencing sensory mechanisms appropriately. As information can have very different attributes and levels of dynamism, acquisition mechanisms should parallel these differences.

Using multiple sensory mechanisms to acquire different types of information makes sense from a completely pragmatic standpoint. Information-specific sensory mechanisms may help an AI perceive the right things in order to achieve what it needs to do. Natural organisms use many mechanisms to acquire and process information, leading to the suggestion that perception is a "bag of tricks" [21]. Acquiring desirable input through multiple appropriate mechanisms for acquiring information is a similar approach. Multiple sensory inputs offer substantial advantages to AI systems. Stoffregen and Bardy [19] argue that sensory apparatus provide multiple points of access to a source of information. For AIs, this can provide resiliency. This is one of the arguments made by Brooks, whose subsumption approach [1, 2] couples sensors to specialised autonomous subsystems that interact to create complex behaviours. Sensing mechanisms in AI can benefit from this perspective, as effective means of acquiring information can vary from problem to problem.

Developing more sophisticated sensory processes intended to acquire information in more sophisticated ways may offer further benefits to AI. As simpler, higher order operations can result from more complex underlying mechanisms [7], sensory systems that produce higher level or less noisy information are likely to have more complex underlying mechanisms. Investing effort in enhancing sensing mechanisms is potentially worthwhile as they can reduce the load on decision architectures. Because more sophisticated information would be available to the AI, it can potentially help existing approaches behave more capably and facilitate new approaches to decision making.

5. USING INFORMATION IN GAME AI

As information affects the ability of an AI to perform effectively, the question of how information design can improve AI performance quickly follows. As actions are strongly influenced by information, they are also one of the attributes most likely to see improvement from this process. Action selection in virtual worlds is a worthwhile area of development as AIs and humans in virtual environments exhibit very different capabilities. A human using a virtual environment is constrained by the game world, but AI systems often have difficulty figuring out what to do with it. Differences in how people and actors use environmental features are not directly due to the game engine, as the constraints of the game engine are potentially identical for both AIs and players. If the game permits a human player to perform an actions, it potentially accommodates an AI doing the same thing. This raises the question of whether information can enhance AI action selection in virtual environments.

AIs must act in several different ways, so they require different types of information. The first is schema-appropriate actions, in which they use features of the environment as they are typically used. An example is using a chair as it is intended, which is to sit on. The other is aschematic usage, in which they use features innovatively and adaptively. The same chair can be stepped on for elevation or used to drape clothes over, for example.

Information to support schematic interactions is well served by existing approaches such as annotations [5] which embed information in the environment. Placing knowledge about objects in the environment and mechanisms to use these instructions in AIs gives them the ability to carry out behaviourally appropriate actions. Providing this type of information is vital, as it enables AIs to deal with schema appropriate actions effectively.

Similar approaches do not support aschematic usage so effectively. The nature of aschematic usage is illustrated by Maier's two-string problem [10]. In this problem, a room contains two strings hanging from the ceiling that are just too far apart that a person cannot, while holding onto one cord, tie it to the other. In the room are a few objects, including a pair of pliers. The problem is easily solved, depending on how the objects in the room are perceived. If the pliers are perceived schematically, they are no use. However, their shape and weight permits them to be tied to one cord and swung to the other. To act adaptably, AIs need more information than annotations. This issue reflects what might be called intelligent use of the environment. We consider the problem of aschematic actions to be one of the more interesting problems in AI, as it relates to issues of adaptability, innovation and flexibility.

Approaches that provide information to help generate both these types of action can reduce the complexity of decision making by providing information that helps solve the problem directly. Embedding information in the game world means that AIs only need to filter this information, rather than possess an internal understanding of every object in the game world. However, as Forster argues, to produce inputs to the AI that are simple, complex mechanisms underly the creation of this simplicity [7]. These complex mechanisms need not be computationally expensive. They can draw on other complex mechanisms, such as human analysis of what information should be made available to an AI to solve these problems.

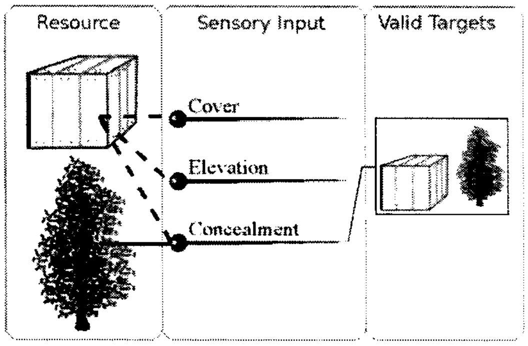

Figure 1: Providing information via annotations

Our research investigates what information supports these capabilities. Information must not only reach the AI, it must be relevant to the capabilities of the AI with regards to its environment. Ideally, the information available to an AI will provide a bridge between the AI's capabilities and what can be done with the environment. One approach with a specific emphasis on this relationship is the concept of affordances, originally proposed by Gibson [8]. Affordance information is relevant to aschematic behaviour due to its emphasis on the relationship between the features of the environment and the capabilities of an entity.

From an affordance-oriented perspective, an entity views the environment in terms of what the environment can offer it, whether harmful or benign. Affordances are properties of the environment that have a relationship with the observers needs and capabilities. An entity sees the environment in the context of what interactions it can have with it. The main emphases of affordances are the focus on the outcomes of the perceptual process, and the emphasis on the intrinsic relationship between perception and action.

It might appear that the amount of information required to provide affordance information is overwhelming, and that an AI would need to know every attribute of every object in its environment. The complexity of the problem is actually far less forbidding than might seem. The information required by AIs in virtual worlds is relatively coarse-grained, translatable and generalised. Despite high levels of visual presentation, many features of virtual worlds are used in common ways.

In game terms, crates in a warehouse may afford an AI cover, elevation or concealment, depending on the requirements of the situation. An approach based on annotations may attach behaviour tags like 'provides cover', 'provides concealment' or 'provides elevation' to crates. This annotation information is shown in Figure 1. However, these attributes are not unique to crates- they may equally match bales of straw, trees, benches or steel drums. Using annotations to inform aschematic behaviour has several limitations. First, if AIs acquire a new capability, each object that supports that capability needs to be updated. Second, if a new AI with different capabilities is introduces, objects need to be coded for both types of AI.

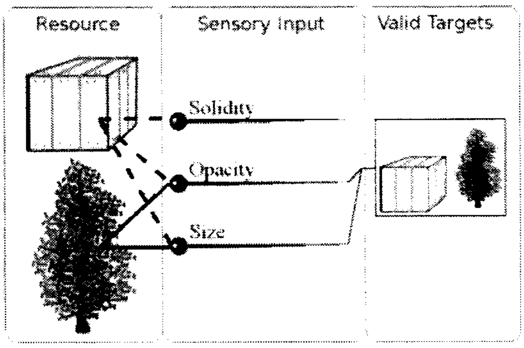

Figure 2: Providing information via affordances

We suggest one way of resolving this problem is to use generalised attributes that describe invariant characteristics of the features in the environment instead [23]. By generalising information it is possible to support a flexible range of actions, and information that can deal with objects at a level that is relevant to the sorts of actions an AI is likely to need to take. Generalised information is potentially more durable and less situationally fragile than scripted uses of objects.

Affordance information provides a way of communicating feature information in a way that avoids the need to explicitly specify matching behaviours. These objects have common characteristics, in that they are solid, opaque and appropriately sized, so they afford the same general actions. Physical cover requires objects have solidity and size, whereas concealment requires objects have opacity and size. These types of attributes can be conveyed to AIs, enabling them to determine what affordances features in the environment possess. This is shown in Figure 2. Affordances provide a way of identifying environmental features that support the general capabilities of AIs.

Sensors can be used to detect the relevant features and assemble affordance information from these inputs. This approach potentially increases the efficiency of information by several means. This approach emphasised sense data can inform more than one action. It also enables AIs with capabilities that vary through influences like power-ups or skills to use the same information, but in different ways. The utility of this information can be seen wen comparing Figures 1 and 2. Using annotations, creatures of different sizes would need separate annotation tags for each object. When general information is used to identify affordances the same information can be reused. This approach illustrates the benefits of matching information to the the problem space an AI is in.

Affordances provide information that let AIs know if particular actions are possible. They identify if the criteria for an action exist in the features of the environment. Marking features of the environment in ways that indicate whether particular actions can be performed gives AIs this affordance information. Affordances provide a way of helping AIs see what they can do in an environment. Other mechanisms can determine whether those actions are desirable.

6. CONCLUSION

The selection of what information is used by an AI pervades the remainder of its design. The technical implementation often has an overriding influence on AI design. While the success of a particular AI comes down to the technical implementation, there are other design issues that shape how well an AI can possibly work. One of the most important issues is ensuring the AI has the right resources, and information is foremost amongst these. Getting the right information to an AI gives the implementation the best opportunity to make effective decisions.

Consideration of the influences on and characteristics of available information offers a means of improving existing approaches in AI. It is also possible to take this further and design AIs around a particular type of information, but a focus on information has benefits for both established and research AI methodologies. The issue of providing information is particularly relevant to virtual worlds as the entire information environment is designed. Consequently, the environment, sensors and decision mechanisms can be implemented in a way that maximises available information and minimises computational load. There are many influences on what information is available to an AI. The primary influence is what information the designer considers relevant. In turn, this is influenced by the way the AI acquires and processes it.

Better information can potentially benefit AIs by enhancing what they can do, and how well they can do it. AIs depend on getting the right information. Designing an AI involves designing an information environment: selecting what information is important, how it is acquired, what form it is presented in and what tradeoffs can be made. Altogether, these elements constitute the entirety of the resources the AI has to operate in its environment.

The design of the information environment determines how well the inputs to an AI articulate with external environment. Various factors influence this articulation. At the simplest level including or omitting particular kinds of information influences what decisions an AI is capable of making. AIs may draw on different types of information, such as features, events, state, relationships and world knowledge, all of which vary in their immediacy, variability and utility. Different elements may benefit from different acquisition processes, whose acuity is a compromise of fit, accuracy and dynamism.

The opportunity to improve AI performance through tuned sensory mechanisms is matched by the potential to improve it through different types of information. Enhancements to the information environment such as annotations have resulted in substantial improvements to both action and planning mechanisms in game AI. However, opportunities to improve AI performance still remain. One candidate for improvement is how AIs interact with environmental features. We have outlined how an approach based on affordance theory could provide AIs with far more adaptive and flexible environmental interactions. Pursuing ways to capitalise on the information already in components of the game engine is one approach.

Our central argument is that AIs should acquire information that matches what they are intended to do. AI uses many established mechanisms of processing and manipulating information, but design considerations have largely emphasised the decision architecture. The information presented to and processed by AIs is as much a part of the system as the tools used to process it, and as such, we believe it should be treated as a design issue.

7. REFERENCES

[1] R. A. Brooks. Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 569–595, Sydney, Australia, 1991. Morgan Kaufmann: San Mateo, CA, USA.

[2] R. A. Brooks. Intelligence without representation. *Artificial Intelligence Journal*, 47:139–159, 1991.

[3] D. Chalmers, R. French, and D. Hofstadter. High-Level Perception, Representation and Analogy: A Critque of Artificial Intelligence Methodology. *Journal of Experimental and Theoretical Artificial Intelligence*, 1991.

[4] R. Dawkins. *Climbing Mount Improbable*. Penguin, London, 1996.

[5] P. Doyle and B. Hayes-Roth. Agents in annotated worlds. In K. P. Sycara and M. Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 173–180, New York, 9–13, 1998. ACM Press.

[6] C. Fairclough, M. Fagan, B. Mac Namee, and P. Cunningham. Research Directions for AI in Computer Games. In D. O'Donoghue, editor, *12th Irish Conference on AI and Cognitive Science*, pages 333–344, NUI, Ireland, 2001.

[7] M. R. Forster. How Do Simple Rules Fit to Reality in a Complex World? *Minds and Machines*, 9:543–564, 1999.

[8] J. J. Gibson. *The Ecological Approach to Visual Perception*. Mifflin, Boston, 1979.

[9] J. Laird. Using a computer game to Develop Advanced AI. *Computer*, 34(7):70–75, 2001.

[10] N. R. F. Maier. Reasoning in humans : II. The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology*, 12:181–194, 1931.

[11] J. E. Matheson. Using influence diagrams to value information and control. In R. M. Oliver and J. Q. Smith, editors, *Influence Diagrams, Belief Nets and Decision Analysis*, chapter 2, pages 25–48. Wiley, 1990.

[12] A. Nareyek. AI in Computer Games. *acm queue*, 1(10):58–65, 2004.

[13] A. Nareyek, B. Knafla, D. Fu, D. Long, C. Reed, A. El Rhalibi, and N. S. Stephens. The 2003 Report of the IGDA's Artificial Intelligence Interface Standards Committee. Technical report, International Game Developers Association, 2003. http://www.igda.org/ai/report-2003/report-2003.html.

[14] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to Evolve Autonomous Robots: Different Approaches in Evolutionary Robotics. In R. A. Brooks and P. Maes, editors, *Proceedings of the IV International Workshop on Artificial Life*, pages 190–197, Cambridge Mass, 1994. MIT Press.

[15] J. Orkin, P. Baillie-de Byl, D. Borrajo, J. Funge, M. Garagnani, P. Goetz, J. J. Kelly III, I. Millington, B. Schwab, and R. M. Young. Working Group on Goal-oriented Action Planning. Technical report, International Game Developers Association, 2004. http://www.igda.org/ai/report-2004/goap.html.

[16] S. Rabin. Stratagies for Optimizing AI. In M. DeLoura, editor, *Game Programming Gems 2*, pages 251–257. Charles River Media, 2001.

[17] D. Sanchez-Crespo Dalmau. *Core Techniques and Algorithms in Game Programming*. New Riders, Indianopolis, 2003.

[18] C. E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 1948.

[19] T. A. Stoffregen and B. G. Bardy. On specification and the senses. *Behavioural and Brain Sciences*, 24:195–261, 2001.

[20] W. van der Sterren. Terrain Reasoning for 3D Action Games. In M. Deloura, editor, *Game Programming Gems 2*, pages 307–316. Charles River Media, 2001.

[21] R. V.S. The neurobiology of perception. *Perception*, 14:97–103, 1985.

[22] S. J. Welsh. Perceiving Information: building better game AI by seeing ideas. Master's thesis, University of Technology, Sydney, 2004.

[23] S. J. Welsh and Y. Pisan. Enhancing agent capabilities in computer games. In *Proceedings of the 2005 International Conference on Artificial Intelligence*, Las Vegas, USA, 2005. CSREA Press.

[24] S. W. Wilson. The animat path to AI. In J.-A. Meyer and S. W. Wilson, editors, *From animals to animats*, pages 15–21, 1991.

[25] W. Wright. *The Sims*. Maxis, Redwood City CA, 2000.