Principled Exploitation of Behavioural Coupling

Chris Thornton
Department of Informatics
University of Sussex
Brighton
BN1 9QH
UK

Email: c.thornton@sussex.ac.uk
Tel: (44)1273 678856

January 18, 2006

Abstract

In robot building, attention tends to focus on internal signal processing and the way desired motor signals are generated. But equally important is the selection and configuration of the robot's sensory resources. The degree to which sensory input *informs* the process of output generation directly impacts internal complexity since the more informative input signals are, the less internal processing is required. In the extreme case, input signals may identify motor signals exactly and internal processing may then be eliminated completely. The paper presents an information theoretic model for measuring this tradeoff being input information and internal processing complexity, and establishes the conditions under which it is feasible to utilise the processing-free strategy of 'direct connection'. The model is then used for purposes of analysing five robots, including two commercial products (Sony AIBO and WowWee's Robosapiens). The principled exploitation of very informative sensory input is then shown to be effective in an unexpectedly wide range of situations.

1 Introduction

The establishment of desired behaviour in a robot is ultimately a question of getting the right command signals to the right effectors at the right time, i.e., it is a question of arranging for the production of *appropriate*, behavioural, outputs. While this may be done on the basis of autonomous processes, more often it will be done contingently with respect to information about the robot's environment. Input signals from sensory devices will be available and these will be taken into account in the generation of action.

Where input signals are thus utilised, there is the question of the degree to which they *inform* the production of output (Thornton, 2000). They may do so to only a moderate degree, e.g., by identifying required output signals approximately or within certain limits. Alternatively, they may identify required output signals perfectly, by virtue of being identical to them.

For the robot builder, it is important to know what the situation is. With relatively informative input signals, less internal processing is required in order to produce desired output; a relatively simple internal architecture may then be all that is required. Where required output signals are *identical* to received input signals, the implications are more far-reaching. In this case, there is *perfect coupling* between the behavioural requirements of the robot and its sensory configuration. By exploiting this, the desired behaviour may be implemented without recourse to internal processing.

The degree to which input signals inform output signals, then, dictates the degree of internal signal processing required for the implementation of behaviour. But how should measurement of this critical quantity be accomplished? The present paper suggests information theory may profitably be brought to bear. It shows that where the output-generation system of a robot is treated as an internal 'receiver' unit, input signals may be treated as encoded action commands allowing information contents to be be derived in the usual way. On the basis of these measurements, deductions may be made regarding the informational deficit that remains following input-signal receipt. The size of the informational deficit is then a quantitative indicator of the load placed on internal architecture.

The paper is organised as follows. Section 2 reviews key principles of information theory. Section 3 demonstrates the use of information formulae for purposes of measuring behavioural uncertainty; section 4 gives definitions for utilised probability distributions. Section 5 presents illustrative scenarios, setting the scene for the introduction of the gain concepts (section 6) which are then used for purposes of defining a measure of behavioural coupling. Sections 8 and 9 deal with the reliability of 'direct connection' as an implementation strategy while section 10 focusses on reviewing a range of illustrative examples. Section 11 is a concluding discussion.

2 Information theoretic uncertainty

In his formulation of information theory, Shannon (1948) envisaged an agent that must choose from a set of alternatives knowing only the probability of individual choices being correct. The agent's level of uncertainty, Shannon observed, depends on the distribution of the probabilities. If one of the choices has a probability of 1, and all the others have a probability of zero, there is no uncertainty: the agent 'knows' the correct choice. If, on the other hand, all the choices have an equal probability, the level of uncertainty is at a maximum for this number of choices. If the number of choices increases while still remaining equiprobable, the level of uncertainty increases further.

As Shannon noted, the only formula which measures uncertainty consistently with these and other reasonable observations[1] is the entropy formula

$$-\sum_i p_i \log p_i. \tag{1}$$

Here, p_i is the probability of the i'th choice being correct.

Entropy is essentially a measure of the degree of choice available (Cover and Thomas, 1991). But Shannon observed that it can also be thought of as the amount of information contained in an event which resolves the uncertainty completely. This dual perspective becomes most intuitively apparent when we calculate entropy using base 2 logarithms; in this case the uncertainty value turns out to be exactly the number of digits in the smallest binary message which could fully resolve the uncertainty. Uncertainty and quantity of information, then, are two sides of the same coin.

For illustration, consider a situation in which there are four, equally probable alternatives. The probability of each alternative is $\frac{1}{4} = 0.25$ and the entropy, using base 2 logs, is

$$0.25 \log 0.25 + 0.25 \log 0.25 + 0.25 \log 0.25 + 0.25 \log 0.25$$

This evaluates to 2, which is the number of digits in the smallest binary message capable of representing the four alternatives we started out with. Equality thus holds between the uncertainty level and the quantity of potential information. With regard to any situation of choice, then, the level of uncertainty is also a potential information gain, and vice versa.[2]

Also of relevance for present purposes is the way information theory deals with information gains between related variables. We have seen that the level of uncertainty $H(X)$ for some variable X is just the entropy of the associated probability distribution. But what is the situation if there is some other variable which has an impact on X? If we have some knowledge of this other variable, it must tell us something about X; the level of uncertainty about X must then be reduced to a degree.

How is this 'conditional' level of uncertainty to be defined? Assuming we know the conditional probability distributions relating values of X to values of the other variable — call it Y — the level of uncertainty about X which remains when there is complete knowledge of Y is

$$H(X|Y) = -\sum_{x \in X} p(x) \sum_{y \in Y} p(y|x) \log p(y|x) \tag{2}$$

This is the *conditional entropy* of X given Y; i.e., the expected uncertainty about X given knowledge about Y.

[1] In particular, the observation that it should be possible to perform consistent summation of uncertainty values.

[2] The effect generalises over all integer base values; i.e., where logs are taken to any integer base n, the uncertainty value is also the number of digits required in base n representation to specify the correct choice.

The gain in information about X which knowledge of Y provides is then the excess of the uncertainty over the conditional uncertainty:

$$H(X) - H(X|Y).$$

However, since the only way that values of Y can say anything about values of X is through specific association, there is always symmetry between the information gain each variable provides about the other: X always says as much about Y as Y says about X. This information gain is thus the *mutual information* of the two variables, denoted

$$I(X;Y).$$

Mutual information may be defined in terms of either information gain:

$$I(X;Y) = H(X)H(X|Y) = H(Y)H(Y|X) \tag{3}$$

The way in which entropy, conditional entropy and mutual information are related is illustrated schematically in Figure 1. In this figure, $H(X,Y)$ is the *joint entropy*. This is the entropy of the joint probability distribution for the two variables X and Y; i.e., it is the total uncertainty relating to the two variables.

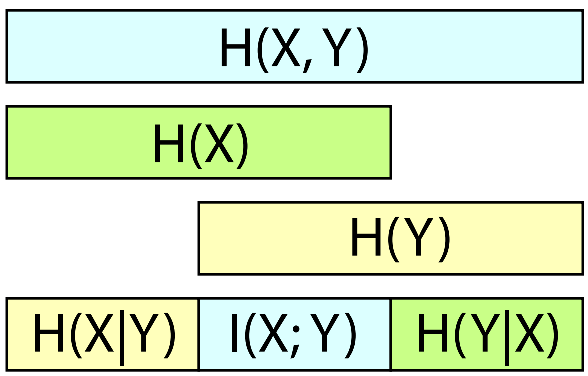

Figure 1: Relationships between entropy, conditional entropy and mutual information.

3 Behavioural uncertainties

The aim of the present analysis is to use the uncertainty formula as a way of measuring the *behavioural* uncertainty of a robot. Provided that the robot's behaviour is made up of discrete actions and that it is possible to ascertain the probability with which specific actions are produced, the robot's behavioural uncertainty may be measured by treating the actions as the outcomes of a discrete, random variable. Entropy may then be calculated in the usual way. Representing the actions of the robot as outcomes of discrete, random variable r, the behavioural uncertainty $H(r)$ of the robot is then

$$H(r) = -\sum_i p_i \log p_i. \tag{4}$$

where p_i is the probability that the robot produces the i'th outcome for r.[3]

As noted, entropy is also the amount of information which is required to resolve the corresponding uncertainty. Once we have defined the robot's behavioural uncertainty, then, we also know the gain in information which is provided by any event which resolves it. Furthermore, where different probability assignments apply at different stages of processing, the robot may be shown to have different uncertainties at different times.

For present purposes, three uncertainties will be distinguished. Taking x to represent the robot's actions prior to any internal processing or assimilation of input,

$$H(x) \tag{5}$$

is the robot's *initial uncertainty*. This reflects the robot's underlying predisposition to act in a certain way regardless of internal processing or sensory input. If the robot has no specific predisposition, i.e., all possible actions are produced with equal probability, the initial uncertainty level is at its theoretical maximum for the given number of choices.

Taking y to represent the robot's actions following receipt of a single input signal,

$$H(y) \tag{6}$$

is the robot's *incremental uncertainty*: This is the level of uncertainty which exists following receipt of an arbitrary *input* signal, i.e., any message, event or datum which is processed by the robot but not itself generated through any sort of processing.[4]

Finally, on the assumption that z represents the robot's actions after all input has been assimilated and all internal processing completed,

$$H(z) \tag{7}$$

is the robot's *final uncertainty*. This is the level of uncertainty that exists *after* the robot has completed all internal processes affecting the given input. As a general rule, we assume that the final uncertainty is zero, which is to say we assume that the robot behaves deterministically. However, the possibility of final uncertainty being greater than zero is accommodated, as we will see.

[3] Logs are taken to base 2 throughout. All information values are therefore expressed in bits.

[4] This definition allows for utilisation of signals drawn from internal state and memory since these do not involve processing.

3.1 Probability distributions

Relevant probability distributions may be defined in terms of the robot's behaviour. Let X be the set of all possible input signals and C be the set of all possible combinations of input signals, i.e.,

$$C = \overline{X}.$$

Let R be the set of all possible actions. The robot's behaviour then expresses a function b from contexts to actions.

$$b : C \rightarrow R \tag{8}$$

On the assumption that the robot produces actions deterministically, b is either one to one or many to one. If the robot responds in different ways (i.e., non deterministically) to the same contexts, the function may be one to many while a many to many function is consistent with any scenario.

The three uncertainties identified above can all be defined in terms of the function b. The initial uncertainty of the robot is written

$$H(r) = -\sum_i p_i \log p_i. \tag{9}$$

We therefore need to know the probability distribution that the robot imposes over actions, i.e., we need to know the p_i values. Taking r_i to represent the i'th action, $P(r_i)$ may be written for p_i. The probability distribution is then

$$P(r_i) = \frac{|\{c \in C : r_i = b(c)\}|}{|C|}. \tag{10}$$

In other words, the unconditional probability that the i'th action is produced is the relative frequency with which this action is seen in the set of all actions.

The definition of incremental uncertainty utilises conditional probabilities of the form $P(r_i x)$, definable as

$$P(r_i \in R | x \in X) = \frac{|\{c \in C : x \in C, r_i = b(c)\}|}{|C|} \tag{11}$$

The form of this definition is essentially the same as in the previous case, except here we are using the frequency with which the action is associated with the relevant input.

Last, the definition of final uncertainty uses probabilities of the form $P(r_i c \in C)$, which may be expressed as

$$P(r_i \in R | c \in C) = \frac{|\{c : r_i = b(c)\}|}{|C|}. \tag{12}$$

This is the frequency with which the action is produced when all input signals are taken into account.

4 Illustrative scenarios

Figure 2 presents seven scenarios illustrating differing uncertainty combinations. Scenarios (a) and (b) are the usual situations regarding final uncertainty. In scenario (a) the final uncertainty is zero. This represents the situation for a *deterministic* robot having no probabilistic or random element in its behaviour.

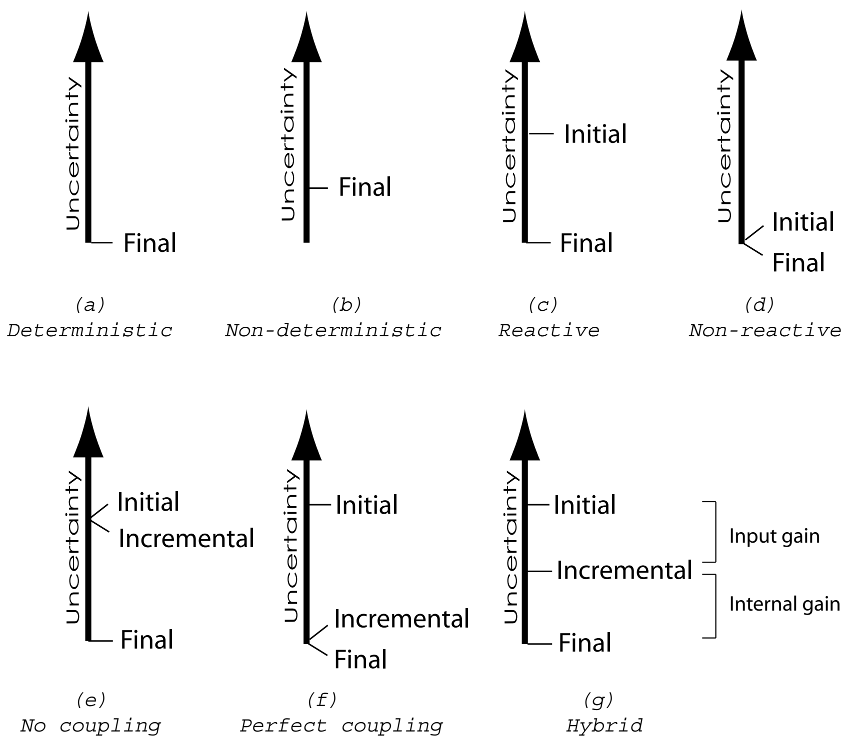

Figure 2: Uncertainty/gain scenarios.

In scenario (b), the final uncertainty has a non zero value. This is the situation for a non-deterministic robot which may produce different responses in the same context. There is a random or probabilistic element impacting behaviour.

In scenarios (c) and (d) the final uncertainty is zero (deterministic robot) but the diagrams now show the initial uncertainty as well. Scenario (c) depicts the case of a 'reactive' robot, i.e., a robot which takes account of input (e.g., sensory) signals in the production of action. The initial uncertainty is the robot's uncertainty prior to the receipt of any input. Scenario (d) is the situation for a non-reactive robot. Here the initial uncertainty is zero indicating that the receipt of input signals makes no difference to the production of behaviour.

In the remaining scenarios, an incremental uncertainty value is shown in addition to the usual values for initial and final uncertainty. In scenario (e) the incremental uncertainty is equal to the initial uncertainty indicating that the receipt of an input signal has no impact on uncertainty. This is the situation in which single inputs are not predictive of output to any degree: there is no coupling between sensory input and behavioural output.

Scenario (f) is the complementary situation in which the incremental uncertainty is zero, indicating that receipt of an input signal resolves *all* the action uncertainty at a stroke. This is the situation of 'perfect coupling' where sensory inputs are so informative that a single case yields perfect confidence with regard to action.

Finally, scenario (g) is the hybrid situation arguably the normal case for robots in practice. Here the incremental uncertainty is greater than zero but still below the initial uncertainty, indicating input signals dictate output to a certain degree but not with certainty.

5 Input and internal gain

As has been noted, the definition of an uncertainty level also defines the amount of information that is generated when the uncertainty is resolved. Any uncertainty level is thus always a potential information gain. However, where different uncertainty levels are identified it is also possible to identify *relative* gains implicit gains that occur when the robot makes a transition from one uncertainty level to another (c.f. Quinlan, 1986).

In the present case, three uncertainty levels have been identified, permitting the identification of two relative gains. The gain achieved when the robot makes the transition from initial to final uncertainty is defined as the total gain:

$$g_{xz} = H(x) - H(z). \tag{13}$$

This is the total information content generated when the robot produces an action. (Gain subscripts identify the chance variables representing relevant 'before' and 'after' situations.)

Contained within the total gain we have, first, the input gain:

$$g_{xy} = H(x) - H(y). \tag{14}$$

This is the mean gain produced by the delivery of a single input signal. It is the expected information content of input signals, i.e., the amount of information the robot acquires from its environment.

Second, we have the internal gain:

$$g_{yz} = H(y) - H(z). \tag{15}$$

This is the information gain that is produced by any internal processing performed by the robot. It is the residual element of gain remaining after input gain has been eliminated from total gain. In all cases total gain is necessarily the sum of input and internal gain:

$$g_{xz} = g_{xy} + g_{yz}. \tag{16}$$

6 Behavioural coupling

The degree to which sensory inputs inform behavioural outputs is related to the level of input gain. However, we cannot use input gain directly as a measure since it is affected by a number of factors including the overall complexity of the behaviour. For purposes of measuring the degree of coupling we must look at the level of input gain *relative* to total gain. If input gain is the dominant factor within total gain, then inputs are relatively informative for production of output. If input gain represents a small component of total gain, we have the opposite case.

In view of this, the level of input gain expressed as a fraction of total gain,

$$\frac{g_{xy}}{g_{xz}}, \tag{17}$$

is the correct measure of *behavioural coupling*. Formally, this quantity is the expected information gain of input signals expressed as a fraction of the robot's total information gain. Informally, it measures the degree to which the robot's inputs dictate its behavioural output.

7 Gain schematic

The generation of gain within the robot may be understood visually in terms of the schematic of Figure 3. The scenario here depicts the processing of a particular set of input signals by an imaginary robot. The robot is represented in terms of its input system (box on the left), processing system (box in the lower middle) and its output system (box on the right). These systems are defined purely in terms of their functional properties and in practice need not be localised in any way: the input system is taken to be the system which produces input signals; the output system is taken to be the system which produces output signals. Labels appearing within the input system represent specific input signals while labels appearing in the output system represent specific output signals (i.e., action commands).

The overall situation depicted is one in which the robot acquires three simultaneous input signals ('Loud noise', 'Left foot touch' and 'High temp') and produces an action signal ('DANCE'). The height of the triangle centred on 'DANCE' represents the total gain in information which is generated when this action is produced. (This is the difference between initial and final uncertainty).

The bar leading from the input 'Left foot touch' to 'DANCE' represents the input signal 'Left foot touch'. The width of the bar is the information content that the signal would have were it to be transmitted direct to the output system.

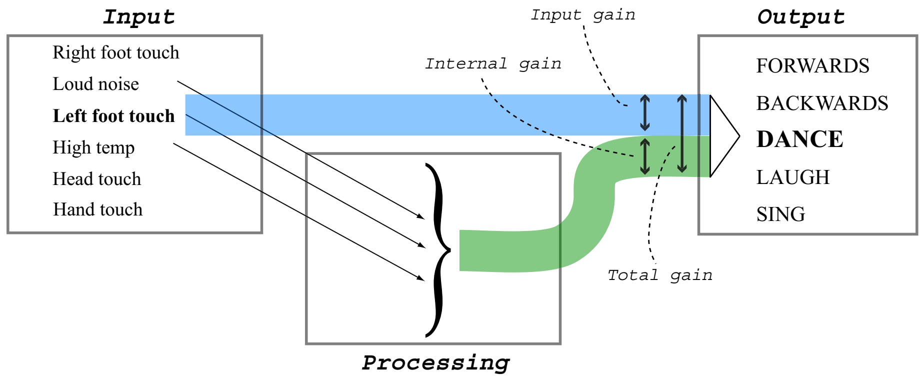

Figure 3: Information gain through behaviour.

On the basis of this, the information content of any signal arriving from the processing system must be the *residual* content remaining when the observed input gain is deducted from total gain, i.e., it must be the informational deficit. Diagrammatically, it is simply the width of the arrowhead, less the width of the input signal bar. Total gain is the width of the arrowhead. Input gain is the width of the bar directly connecting 'Left foot touch' to 'DANCE' and internal gain is the width left over.

8 The reliability of direct connection

From the perspective of effort minimisation, the most favourable situation for the robot builder would seem to be one where behavioural coupling for the intended robot is at its *maximum* (i.e., a coupling value of 1.0). In this case, no internal processing is required. Behaviour may be implemented using *direct connection*, i.e., by re-using input signals directly as output signals. And while this may appear to be a situation that robot builders will encounter infrequently, the generality of direct connection as an implementation strategy may be greater than assumed.

Another way of understanding behavioural coupling is as an indication of achievable *reliability*. The degree of coupling necessarily increases as input signals concentrate *higher* probabilities over *fewer* actions. But this also increases the likelihood that behaviour executed in accordance with the probability distribution will be behaviourally appropriate. The measure of behavioural coupling may also be used therefore as a means of assessing the reliability (i.e., the expected probability of appropriate action) of a robot which implements desired behaviour *without* use of internal processing.

Of course, the relationship between coupling and the reliability is not unexpected given their mathematical origins. Defined in terms of relative input gain, behavioural coupling is essentially a function of the entropy of action probabilities:

$$-\sum_i p_i \log p_i.$$

To calculate the reliability (probability of the robot producing behaviourally correct action), a variation of this is needed. The likelihood of a particular action being appropriate is just the probability value of the action (i.e., the relevant p_i). The *expected* likelihood of the robot producing a correct action is therefore the sum of products of likelihoods and probabilities. But since both are represented by the same p_i value, this is just

$$\sum_i p_i p_i.$$

Formally, this expresses the expected probability of correct action when the robot is behaving in accordance with the probability distribution. Informally it is the reliability with which a robot can produce the specified behaviour without use of internal processing.

For illustration, consider the case of a robot with two, equi-probable actions. Each action is produced with probability 0.5 and the probability of either action being correct is also 0.5. The expected likelihood of the robot producing a behaviourally correct action is then

$$(0.5 \times 0.5) + (0.5 \times 0.5) = 0.5$$

This agrees with intuition since in this scenario the robot always has a 50/50 chance of producing correct action.

An empirical demonstration of the intimate relationship between behavioural coupling and direct connection reliability is provided in Figure 4. Each curve in this graph plots values obtained from 6000 randomly generated probability distributions of a fixed cardinality. Each point shows the correspondence between the normalised entropy[5] of the distribution and the expected reliability of a robot behaving in accordance with the distribution. Strong behavioural coupling is associated with high input gain, i.e., a situation of low uncertainty/entropy. The graph thus plots reliability *inversely* against behavioural coupling.

Curves associated with higher set cardinality (i.e. distribution size) are shorter and displaced downwards and to the right. This is because larger samples necessarily produce higher entropy values and lower reliabilities. The key point to note, however, is the general tendency for reliability to vary *inversely* with entropy, and therefore monotonically with behavioural coupling.

The expectation of reliability may be written as

$$\alpha = 1 - \beta + \frac{\beta}{\gamma} \tag{18}$$

[5]Normalised entropy is the entropy expressed in terms of the maximum entropy value for a distribution of the same size.

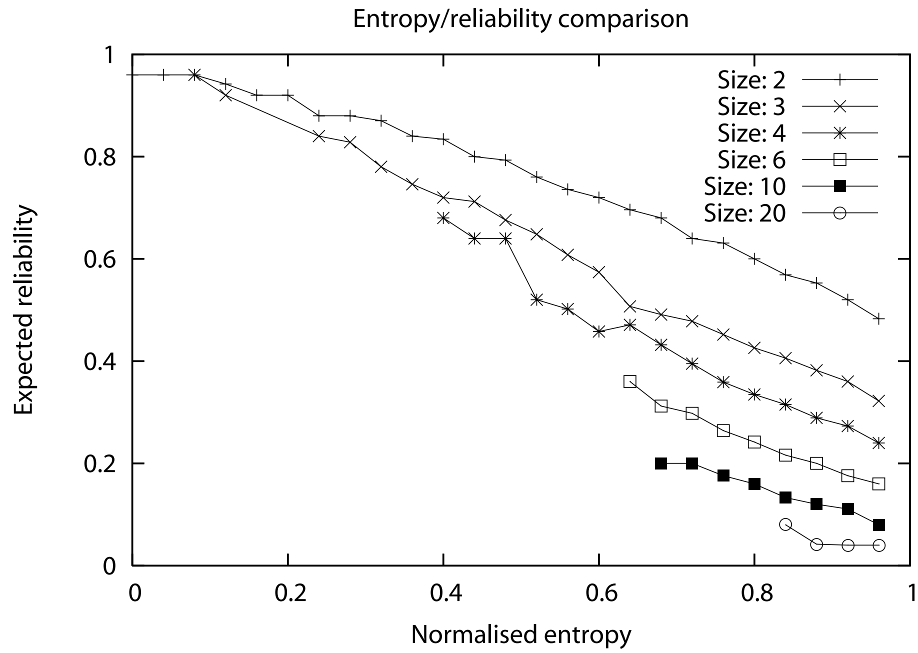

Figure 4: Entropy/reliability curves.

where α is reliability, β is normalised entropy and γ is distribution size. This is a rough approximation of the true reliability value; but it will suffice for most practical purposes.

9 Application examples

The model provides the means of assessing the reliability with which a robot may produce behaviourally appropriate action without recourse to internal processing. Alternatively, it may be used to guage the need for internal processing in a 'perfectly behaved' robot. It can be applied whenever there is a clear specification for (or sample of) the relevant robot's behaviour. Once this has been acquired, probability distributions may be calculated and from these, gain values computed and the degree of behavioural coupling determined. This may then be used to choose an appropriate implementation strategy, either emphasising the use of internal signal processing or minimising it.

9.1 Braitenberg vehicles

To illustrate the way the model may be used in practice, a series of worked examples are now presented. Considered first is the stimulus pursuit robot ('vehicle 2b') described by Braitenberg (1984). Equipped with just two sensors and a differentially steered wheelbase, the behaviour of this robot is to smoothly turn and move towards any source of stimulation, with the intensity of the motion being proportional to the strength of the stimulation. Representing input and output signals as single digits, and tabulating signals in columns depending on their source/destination, a small sample of the behaviour is as shown in Table 1.

x_1	x_2	y_1	y_2
1	4	4	1
2	4	4	2
2	5	5	2
3	5	5	3
4	6	6	4
4	7	7	4
5	7	7	5
6	7	7	6

Table 1: Stimulus pursuit behaviour sample.

Each row in the table represents a specific stimulus/response association: specifically it shows what input signals were present at the production of a specific action, where signals listed in the x_1 and x_2 columns are signals from the two sensors, and signals in the y_1 and y_2 columns are signals sent to the two motors.

Applying the information formulae above to a sample of this type containing 500,000 examples, probability distributions and gain values were derived. Input gain was determined to be 3.47 bits while internal gain was determined to be 0 bits. On the basis of this, the maximum behavioural coupling value of 1.0 was obtained, indicating *perfect* behavioural coupling in this behaviour.

From the theoretical point of view, the analysis shows that a robot may be constructed which will perform this behaviour perfectly without the need for *any* internal processing. Of course, the conclusion comes as no great surprise since it is plain to see that output signal combinations always perfectly reproduce the associated the input signal combination.

The finding that direct connection will perfectly implement stimulus pursuit may also be corroborated in terms of Braitenberg's original design for vehicle 2b (Figure 5). As will be seen, the sensor on the robot's left side is connected so that its signal directly controls acceleration of the motor driving the wheel on the right side, while the sensor on the right controls acceleration of the wheel on the left. The architecture is based exclusively on direct connection. Internal processing does not figure in any way.

For comparison, it is useful to look at Braitenberg's vehicle 3a. This is the robot characterised by Braitenberg as displaying 'love'. Using the same sensory/motor setup as vehicle 2b, it produces much the same behaviour, except

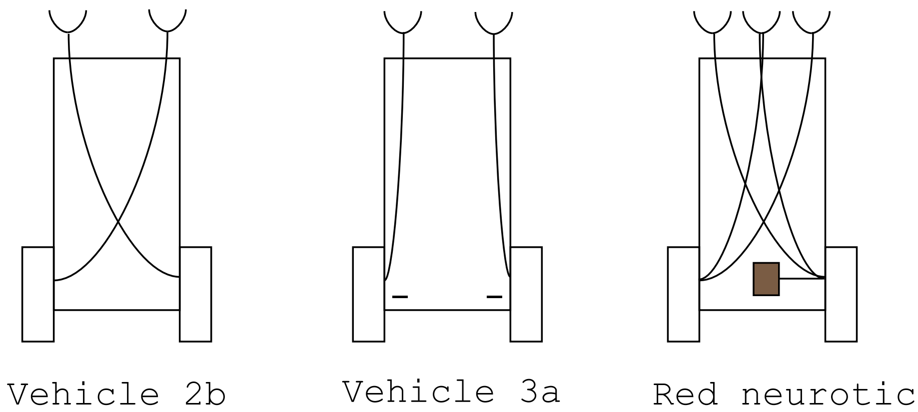

Figure 5: Vehicle architectures, 2b and 3a after (Braitenberg, 1984).

for the fact that it decelerates as it nears the stimulus. A sample of the behaviour (using the same conventions as before) is shown in Table 2.

x_1	x_2	y_1	y_2
7	6	0	0
6	7	0	0
3	7	0	2
6	4	2	0
3	6	0	3
4	7	0	1
5	7	0	1
0	0	6	6
1	0	5	5
0	2	4	6
0	3	3	6

Table 2: Sample of 'love' behaviour.

Using a sample of this type incorporating 500,000 data, an input gain of 1.75 bits and total gain of 2.42 bits were derived, with the behaviour coupling then being 0.72. The lower value derived in this example tells us that reliable execution of this behaviour would appear to require *some* degree of internal processing. Braitenberg's own design for this robot (see figure shown) is similar to that of vehicle 2b and there seems to be no utilisation of internal processing. However, there is a hidden feature of the design which is not captured in the drawing.

Braitenberg's design relies on the assumption that a constant signal is fed to both drive wheels. Signal from the two sensors is then inverted and fed to the drive wheels, with each sensor feeding signal to the wheel on its 'own' side. As the vehicle approaches the source of stimulation, the sensors register increasingly high signal but, due to the inversion, this produces the effect of reversing the rotation of the wheels, effectively slowing the vehicle down. Furthermore, the sensor on the side on which the source appears generates a relatively higher signal, causing greater deceleration of the wheel on that side and therefore a rotation *towards* the source. It is the existence of the constant signal sent to the drive wheels in this example that impacts input gain and behavioural coupling.

As a further illustration of the impact autonomous signal sources may have, consider a self designed variation of vehicle 3a: the 'red neurotic' of figure shown. This is similarly constructed with a differentially steered wheelbase. However, three sensors are now employed instead of one. Objects are the source of stimulation and the behaviour of the robot is, again, to approach all objects but in this case to show 'neurotic' behaviour with respect to red ones, initially accelerating towards them but swerving away just before impact. An illustration of the behaviour is shown in Table 3.

x_1	x_2	x_3	x_4	y_1	y_2
6	0	6	0	6	6
7	0	6	0	7	6
7	0	7	0	7	7
8	0	7	0	8	7
9	0	7	0	9	7
9	0	8	0	9	8
8	0	8	0	8	8
7	0	8	0	7	8
5	0	8	0	5	8
2	8	3	0	7	8

Table 3: Sample of 'red neurotic' behaviour.

On the basis of 500,000 input/output examples, total gain was measured at 4.21 bits with input gain being 2.63 bits. The behaviour coupling value was then 0.622, a still lower value suggesting a still more significant role for internal processing.

In the design for this robot, the two sensors are cross connected to the wheels, producing the usual stimulus pursuit behaviour. The central sensor, however, is sensitive to red objects, producing signal only when the vehicle is 'head on' to a red object. Finally, the processing unit has the effect of delivering an additional input to the right wheel in the case where a red object is immediately adjacent (an effect not seen in the sample shown). There are thus *two* ways in which internally processed signals have an impact in this design and the lower coupling value reflects this.

9.2 Commercial robots

The model is intended for use by robot builders as a means of assessing the degree to which internal processing may be dispensed with in the implementation of behaviour. However, as we have seen with the Braitenberg designs, where the robot (or a design for it) already exists, the model may be used to either guess the character of the design or to decide whether it involves *redundant* signal processing. The latter usage is interesting in the case of commercial robots, where the underlying design may be unknown. Provided that a representative behaviour sample can be obtained, the behavioural coupling can be assessed and, on the basis of this, inferences made about the impact of internal processing.

Unfortunately, behaviour samples for commercial robots are often *difficult* to obtain due to lack of information about implementation details. However, a qualitative application of the analysis may still be made. The essential idea in the model is that where input signals are predictive of action, incremental uncertainty is lower and input gain higher. Behavioural coupling is then specifically proportional to input gain. A viable means of qualitatively assessing a robot is therefore to search for evidence of the *exploitation* of input gain. The more evidence that can be turned up, the greater the indication of behavioural coupling.

9.2.1 The Sony AIBO

One of the better known commercial robots at the present time is the Sony AIBO. This robotic 'dog' (Figure 6) is claimed by Sony to offer 'communication' and 'companionship' to its owner and indeed, in its ERS7 version, it exhibits an undeniably impressive technical specification. Equipped with a 64 bit RISC processor running at 576 MHz and 64 MB of RAM memory, it is capable of processing data from sound and image devices as well a large array of infrared (distance), kinetic, vibration and touch sensors. It can engage in wireless communication, play music, make communicative noises and display a whole range of motor behaviours, including walking and 'dancing'.

The AIBO makes extensive use of software controlled behaviour and according to Sony this is the secret of its 'social' skills. As the publicity material on the AIBO website states, 'In day to day life, this software enables AIBO to entertain and communicate with you. A privileged companionship will flourish with you thanks to AIBO cleverly recognising your face and voice.'

Startling claims indeed. But is there any evidence of input gain usage in the case of the Sony AIBO? At first sight, it seems unlikely. The AIBO engages in both audio and visual processing and, when one considers the sheer volume of processing involved, there seems no possibility that input gain could figure to any significant degree. If each value in an audio stream is regarded as a distinct input signal, and if processing proceeds at consumer rates, then something in the region of 80,000 input signals must be processed *per second*. Some multiple of this number must then be taken into account in the generation of

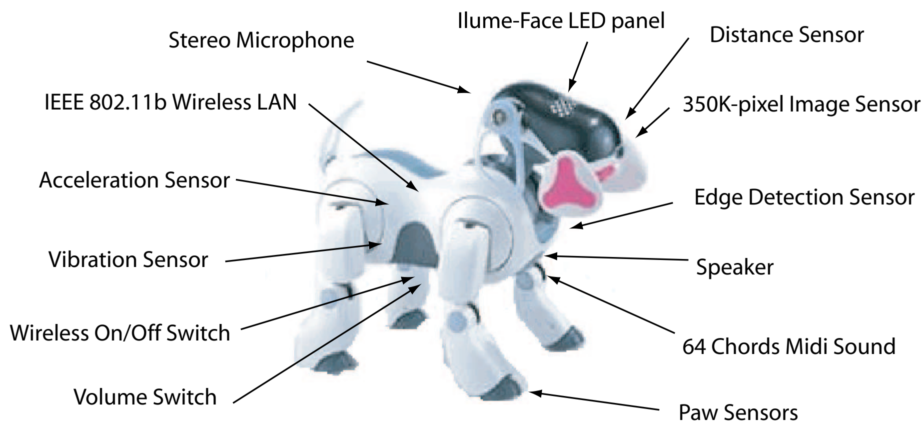

Figure 6: Features of the Sony AIBO viewed from the front.

even a primitive motor response (e.g., turning the head in a certain direction). Presumably, no single input signal could possibly be predictive in this context.

Much the same can be said with regard to the processing performed on visual data. This is carried out with respect to data derived from a 350,000 pixel CMOS image sensor. The visual data stream taken into account is therefore still more voluminous than the auditory one. Again, there seems no possibility that any single input signal could possibly have any predictive value. The idea that there may be exploitation of input gain in visual or auditory modalities seems implausible. But what can be said for other aspects of behaviour?

The AIBO is equipped with a touch sensor on the crown of its head (see figure shown) and, as Sony states, the 'AIBO immediately reacts when this is touched.' On the assumption that this reaction is 'ballistic', i.e,. is completed regardless of ensuing input, signals from the touch sensor are strongly if not perfectly predictive of the output signals utilised in the control of the reaction. With respect to the processing of signals from touch sensors on the head and other significant parts of the body (e.g., the back and chin), use of input gain is apparent with a corresponding diminution in cognitive gain.

Qualitative analysis of the AIBO presents a mixed picture, then. And even with respect to the visual and auditory modalities there are reservations to be entered. Given the temporal nature of the data stream in auditory and visual analysis, its may not be entirely appropriate to equate signals with single data, as is done above. In the case of visual analysis, for instance, where the CMOS sensor is directed towards a particular scene and there is no change in head position or lighting, all the values originating from a particular cell of the sensor will be similar. For the analytic point of view, it may make more sense to take a coarse-grained perspective and view the entire *sequence* of similar values as comprising 'the signal', rather than any one of its elements.

On this basis, the analysis of auditory and visual modalities would need

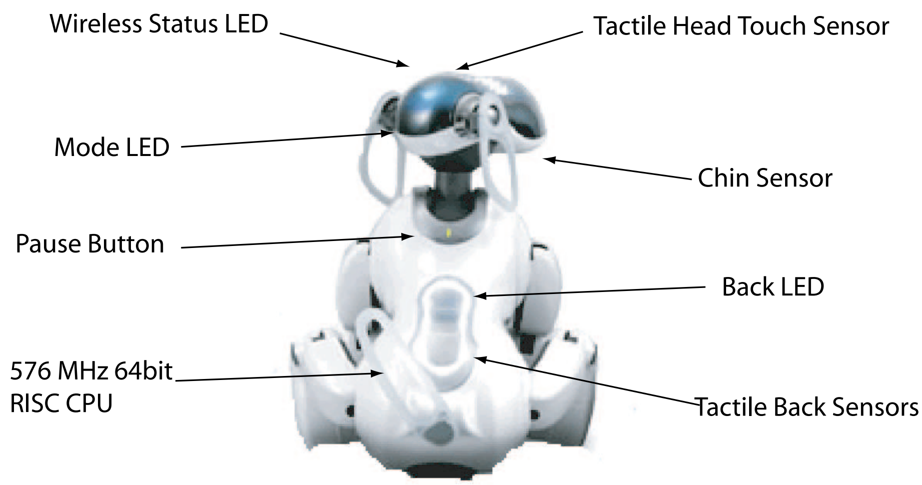

Figure 7: Features of the Sony AIBO viewed from the rear.

revision and we would have to check again for evidence of input gain. On the auditory side, there could be gain utilisation if, for example, voice commands were recognised on the basis of tone, say, i.e., a sequence of data all drawn from a certain range of frequencies. On the visual side, there could be utilisation if, say, behavioural responses were triggered by certain ranges of values registering in certain regions of the CMOS sensor. Is anything along these lines discernible?

One of the advertised skills of the AIBO is its ability to return to its charging station so as to recharge its lithium-ion battery pack whenever the need arises. This in fact is a key element of the vaunted 'autonomy' of the robot. However, Sony's publicity material reveals that the AIBO charging station for this production model has been designed in a very specific way.

Figure 8: The AIBO recharging station target.

As can be seen (figure shown) the charging station is placed in the centre of a two large targets one of which is a large cylindrical object placed immediately behind the charging unit itself. The other target is a disk placed on the floor immediately in front of the unit. Both of these are adorned with a distinctive pattern of interlocking black and white rectangular stripes.

Sony reports that use of this target is essential in order for the AIBO to be able to 'find its way' back to its charging station. But there is another way of understanding what is going on. From the coarse grained perspective, where signals are equated with streams of similar values, the introduction of the charging target creates a situation in which specific input signals are much more likely to be *predictive* of action.

Consider the cylindrical element of the target. Given that the AIBO camera's height over the floor is fixed by the structure of the robot, the positioning and shaping of this part of the target greatly improves the chances that the AIBO will acquire a uniform image of it regardless of the direction from which it is approached. Specific input patterns may thereby be configured to predict specific behavioural responses. The charging target looks to be a vehicle for 'amplifying' input gain in the visual modality.

It is not just with respect to touch triggered behaviours, then, that we find evidence of input gain playing a role in the AIBO. The charging scenario may also fall partly or wholly within this category. Indeed, there are suggestions that exploitation of input gain may feature elsewhere. For example, the arrangement whereby the robot is configured to fetch small objects (such as bones and balls) which are specifically *pink* in colour is a promising candidate for further consideration.

9.2.2 Robosapiens

Produced by toy manufacturer WowWee, Robosapiens is both cheaper than the AIBO (approximately 1/10th the price) and less sophisticated in terms of sensory equipment and on board processing. More of a child's toy than a robotic pet, its behaviour is largely controlled by remote control device. However, it does have some sensory capacity four touch sensors and a 'sharp sound' sensor and a limited capacity for independent action. Like the AIBO it is capable of both walking and 'dancing'.

Robosapiens can be subjected to the same informational analysis as the AIBO. But the results are considerably more decisive. Whereas with the AIBO, evidence of input gain is dimly discernible, in the case of Robosapiens it is readily apparent. In fact, were remote control commands to be treated as bona fide input signals to the robot's cognitive system, every one of the 67 'command functions' would exhibit maximum input gain since every one is perfectly predictive of output. However, it seems more useful to eliminate the remote control aspect from consideration and concentrate on the residual autonomous behaviours.

As far as can be discovered from the user manual, there are just two actions which are built-in. As stated, 'Robosapiens responds to a sound or tap on his

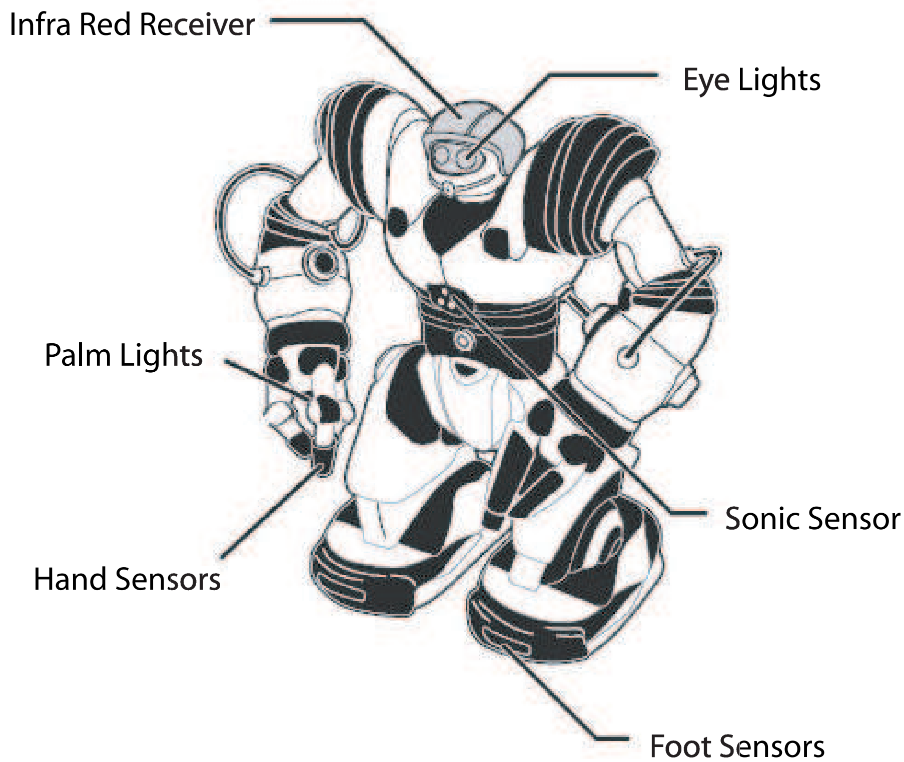

Figure 9: Features of Robosapiens.

body with a default grunt, or a Sonic Sensor Program sequence programmed by you.' There is also a default power off response: 'After approximately 2 hours of uninterrupted sleep, he'll power himself off to save energy.' However, there is no processing of inputs from these sensors. Touch sensors emit a single signal when stimulated and this is used directly to trigger the relevant behaviour. Similarly with the Sonic Sensor: this emits a signal in response to a sharp sound (e.g., a clap) and this is used directly to trigger behaviour.

With regard to the built-in behaviours, then, Robosapiens relies exclusively on perfectly predictive inputs. However, using the remote control, Robosapiens can also be programmed to respond to any particular signal from a touch sensor (on the feet and hands) or from the Sonic sensor with a particular sequence of six actions (e.g., 'forwards step', 'right turn' and 'left hand pick up'). In effect, the robot can be programmed to produce different autonomous behaviours.

This programmability does not affect the analysis of coupling, though. All instances of programmed responses to specific touches will fall into the perfectly predictive category of signal processing since they are all based on direct use of a single signal. The conclusion is that Robosapiens' behaviour is based exclusively on the processing of predictive inputs. Input gain is at a maximum in this case. From the theoretical point of view, Robosapiens utilises perfect, behavioural coupling.

9.3 Coupling values for training agents

A useful source of misc. behaviour samples is the UCI repository of machine learning databases.[6] These databases are actually *training sets* intended for use with machine learning methods such as C4.5 (Quinlan, 1993) and Back propagation (Rumelhart *et al.* 1986). However, being made up of input/output associations, they are also suitable for use as behavioural specifications of 'target agents', an approach in which values of input/output variables are viewed as signals from/to input/output devices.

Measurements of behavioural coupling are straightforwardly calculated for these target agents using the formulae above. The data shown in Table 4 shows coupling values that have been calculated for 12 of the most commonly used training sets in the UCI repository. (These are from the subset used in Holte's study (Holte, 1993)). Abbreviations refer to database names as follows: BC=breast cancer, GL=glass, G2=glass2, HD=heart disease, HE=hepatitis, HO=housing, IR=iris, LA=labor negotiations, LY=lung cancer, SO=soybean, VO=voting records, V1=voting records1.

BC	GL	G2	HD	HE	HO
0.19	0.94	0.83	0.68	0.89	0.99
IR	LA	LY	SO	VO	V1
0.94	0.89	0.37	0.88	0.55	0.55

Table 4: Behavioural coupling of UCI target agents.

On the basis of this array of coupling values, it is clear that internal processing is required in the BC target agent (coupling 0.19) and the LY target agent (coupling 0.37). Elsewhere, however, the relative strength of coupling suggests that direct connection is a feasible strategy in most cases, and particularly so in the cases of GL, G2, HD, HE, HO, IR, LA and SO.

The observation tends to confirm the results of Holte's study, the main conclusion of which was that these training sets do not represent particularly challenging behaviours.[7] Holte showed that rules constructed on the basis of a constraint applied to a *single* input variable produced performance which was almost as good as that produced by unrestricted methods. This is precisely the expected consequence where behavioural coupling is strongly instantiated.

Most of the target agents considered in table shown are associated with real world training sets, i.e., training data which were derived in a non artificial

[6] `http://www.ics.uci.edu/~mlearn/MLRepository.html`.

[7] The title of Holte's paper was 'Very simple classification rules perform well on most commonly used datasets.'

context. However, there is evidence suggesting that specifically *artificial* training sets from the UCI repository may show higher gain values than ones derived from the real world. The 'monks' training sets, devised as challenges in a machine learning competition, are a case in point. All three yield coupling values towards the lower end of the scale and, in one case, a value close to zero, see Table 5. [8]

monks-1	monks-2	monks-3
0.30	0.01	0.32

Table 5: Behavioural coupling for 'monks' training sets.

10 Discussion

The 2004 United Nations World Robotics survey has shown just how quickly use of robots has been increasing in recent years. It reports that in 2004, over 600,000 'household robots' were already in use, with the projected figure tipped to top one million within the next few years. Usage of robots of all categories (including industrial devices) is predicted to grow at an average rate of about 7%.

The two consumer robots considered above the Sony AIBO and Robosapiens are essentially entertainment products. But domestic robots typically have quite practical goals. The Electrolux Trilobyte, for example, is a robot vacuumer, as is the Dyson DCO6 Ultravac. The Husqvarna robot lawn mower will 'constantly and quietly keep the lawn trimmed' while the KUKA KR 40 PA case packing robot offers the consumer the 'highest possible pack rate'.

Robotic toys are also increasingly in evidence. These range from the traditional build it yourself kits (e.g., Lego Mindstorms), to out of the box products such as Robosapiens. And giving a flavour of what may be to come in the future, there are now a number of humanoid robots under development with the major electronics manufacturers. From Sony, for example, there is the 'singing, jogging and dancing' Qrio (Figure 10, right) and from Honda, the miniature humanoid Asimo robot (Figure 10, left).

Clearly, robots are being developed, designed and built in considerable and increasing numbers.[9] Yet the process of behaviour implementation remains a largely unprincipled activity, not guided to any great extent by the application of foundational theory. On the basis of the model presented above, however, robot builders do now have the means of making a principled choice regarding use of direct connection. Whenever behaviour samples are available, the model may also be used to mathematically assess the degree of behavioural coupling evidenced in the relevant scenario. Using the reliability expectation (equation

[8] Interestingly, the very low value for the 'monks-2' agent is consistent with results obtained for these problems in the comparative study (Thrun *et al.* 1991). Although performance varied widely, poor performance on 'monks 2' was frequently observed.

[9] See the 'gizmag' article at `http://www.gizmag.co.uk/go/3500/` for a list of current products.

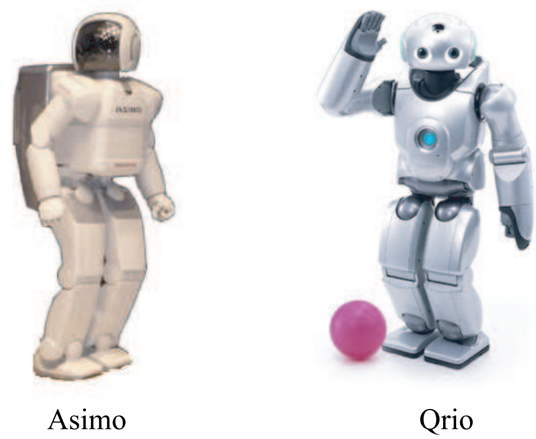

Figure 10: Honda's ASIMO robot with the Sony Qrio.

17), the expected probability of appropriate action in the absence of internal processing can also be estimated. A principled decision may then be made with respect to the use of direct connection for satisficing purposes, i.e. where a degree of behavioural imperfection is tolerable.

While the implementation of desired behaviour on the basis of direct connection is generally associated with toy scenarios such as those described by Braitenberg, the evidence suggests the strategy may be more general. Although, direct connection can only provide a perfect implementation of desired behaviour where the degree of behavioural coupling is maximal, this does not mean that the strategy is a 'laboratory curiosity'. The reliability of behaviour implemented this way is directly related to the degree of coupling. Thus, even a modest degree of coupling a value of 0.75 say still allows for a relatively robust implementation.[10]

The implication is that in those cases where behavioural perfection is surplus to requirement, direct connection may well be effective provided that the level of coupling is sufficient to achieve required levels of reliability. Exactly how generally this may apply in practical robotics is hard to say at present. However, the evidence from the survey of target agents promises a reasonably broad coverage. Results obtained suggest that strong behavioural coupling is a common feature throughout the UCI group and indeed that it may be dominant amount agents which are most directly associated with real-world tasks.

[10]The implementation may be expected to produce appropriate action in 75% of cases.

References

Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. London: The MIT Press.

Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. John Wiley & Sons, Inc.

Holte, R. (1993). Very simple classification rules perform well on most commonly used datasets. *Machine learning*, *3* (pp. 63-91).

Quinlan, J. (1986). Induction of decision trees. *Machine Learning*, *1* (pp. 81-106).

Quinlan, J. (1993). *C4.5: Programs for Machine Learning*. San Mateo, California: Morgan Kaufmann.

Rumelhart, D., Hinton, G. and Williams, R. (1986). Learning representations by back-propagating errors. *Nature*, *323* (pp. 533-6).

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27* (pp. 379-423 and 623-656).

Thornton, C. (2000). *Truth from Trash: How Learning Makes Sense*. MIT Press.

Thrun, S., Bala, J., Bloedorn, E., Bratko, I., Cestnik, B., Cheng, J., De Jong, K., Dzeroski, S., Fisher, D., Fahlman, S., Hamann, R., Kaufman, K., Keller, S., Kononenko, I., Kreuziger, J., Michalski, R., Mitchell, T., Pachowicz, P., Reich, Y., Vafaie, H., Van de Welde, W., Wenzel, W., Wnek, J. and Zhang, J. (1991). The MONK's problems - a performance comparison of different learning algorithms. CMU-CS-91-197, School of Computer Science, Carnegie-Mellon University.