

Gauging the Value of Good Data: Informational Embodiment Quantification

Chris Thornton
COGS/Informatics
University of Sussex
Brighton
BN1 9QH
UK
c.thornton@sussex.ac.uk

August 17, 2010

Abstract

While there is increasing recognition of how relations of embodiment can be exploited for achievement of cognitive goals, we still lack any general method for formalizing the benefits that are then obtained, or for quantifying them. The present paper describes a method that can be used to calculate the informational benefit obtained when embodiment becomes a vehicle for generation of ‘good data’, i.e., data exhibiting behaviourally salient correlations.

Keywords: embodiment quantification, information theory, good data

1 Introduction

Cognitive science has often tended towards the ‘input/output’ view of cognition, conceptualizing the process as internal processing of stimuli obtained from a clearly delineated external environment (Hurley, 1998). But in recent years, work in the embodied/embedded approach (Ballard, 1991; Agre and Chapman, 1991; Brooks, 1991; Thelen and Smith, 1993; Beer, 1995; Kirsh, 1995; Clark, 1997; Wheeler, 2005; Pfeifer and Bongard, 2007; Clark, 2008) has highlighted ways in which intelligent behaviour can emerge out of the interplay between brain, body and world. On this view, cognitive functionality may be understood to involve exploitation of physical interaction between agent and environment.

One illustration involves ball catching. On the input/output view, a ball-catcher accommodates visual stimuli to a 3-dimensional model of the environment, prior to plotting out the trajectory of the ball and calculating a sequence of motor outputs that will produce an appropriate catching position. On the embodied/embedded view, the ball-catcher moves physically so as to maintain

the image of the ball at a constant angle and direction. Exploiting a useful body/world relationship reduces the need for information processing, meeting the general principle of ‘cheap design’ (Pfeifer and Bongard, 2007).

The embodied/embedded perspective raises interesting questions about the degree to which cognition should be seen as a process that operates inside the head/body, or whether it is more appropriately seen as spreading across the interface between the two (Beer, 2000). A debate can then proceed on the question of whether we should see the human *mind* as contained somehow within the brain, or as ‘extended’ (Clark and Chalmers, 1998; Adams and Aizawa, 2008). More technical issues are also raised. Key among these is the degree to which exploitation of embodiment has the potential to produce quantifiable advantage. This is the difficult task of ‘measuring exactly *how much* difference embodiment makes with regard to some behavior, capacity or ability’ (Clark, 2008, p. 213).

Work on this topic faces many obstacles. It is a problem in particular that embodiment seems to offer different types of benefit, which often defy quantification (Varela *et al.* 1991). Sometimes, the advantage obtained seems specifically informational in character, however. There then looks to be the possibility of objective measurement. Clark particularly highlights this proposal. The question of how to quantify embodiment advantage, he notes, ‘sounds less peculiar ... once we begin to view embodiment through a broadly speaking information-theoretic lens — that is to say, once we attempt to understand the cognitive roles of body, action, and environment by understanding their roles in the elicitation, storage, transformation, and processing of information, and in securing its poise for use in the control of intelligent action’ (ibid., p. 214).

Clark suggests we can understand such effects in terms of production of ‘good data’. Informational benefits arise, he suggests, from an agent’s capacity ‘to create or elicit appropriate inputs, generating good data (for herself and others) by actively conjuring flows of multimodal, correlated, time-locked stimulation’ (ibid, p. 196). Pfeifer and Bongard (2007) also emphasize the way different kinds of physical interaction can produce differently correlated data. In their view, this is the way ‘the brain gets good — that is, correlated — raw material for further processing’ (p. 360). The benefit that is obtained in this way must correspond to the ‘goodness’ of the good data. Quantifying the effect would then seem to require measuring the quality of correlation achieved. But how is this to be accomplished?

Relevant to this question is work on information self-structuring by Lungarella, Sporns, Pfeifer and colleagues (e.g., Lungarella and Pfeifer, 2001; Lungarella *et al.* 2005; Sporns and Pegors, 2003; Sporns and Pegors, 2004; Lungarella and Sporns, 2005; Lungarella and Sporns, 2006; Pfeifer *et al.* 2007). A key observation in this is that ‘through an agent’s physical interaction with its environment, informative and correlated sensory signals are generated in different sensory channels’ (Pfeifer and Bongard, 2007, p. 360). Various mechanisms of ‘delivery’ are recognized: the approach also suggests ways in which benefits can be quantified. In a recent proposal, agent’s are envisaged to maximize *empowerment* (Klyubin *et al.* 2007). This is a measure of the amount of infor-

mation an agent “injects into” its sensory data through interactions with the environment.¹

Might measuring empowerment be a way of quantifying informational benefits of embodiment? The metric clearly reflects the degree to which agents can create additional sensory information by modulating behaviour. It is not clear that this meets the general requirement, however. The quantity obtained treats the agent as the origin of any changes in physical interaction, and is thus context-sensitive with regard to flow of control. Benefits of embodiment that arise through changes originating independently of the agent cannot be straightforwardly taken into account.

What seems to be required is a specifically *context-free* measure, i.e., a way of gauging the contribution obtained from production of well-correlated data that makes no assumptions about how the deed is done. The measure should be indifferent between processes mediated by agent, environment, or interaction between the two. Applying it, we would then be able to quantify the advantage gained from embodiment by evaluating the degree to which informational gains originate specifically in physical interaction.

What form should this measure take? In the proposal of the present paper, analysis of information refinement is used as a foundation.² Formal modeling of information-increase becomes the means of assessing what happens when data are put into a correlated form. By determining the sequence of optimal refinements that serves to produce a particular set of correlated data from a corresponding set of uncorrelated data, we identify the minimum information change between the two. This then quantifies the informational benefit obtained by any consumer of the correlated data. Where production of the data can be shown to be mediated by physical interaction, the value then determines the information-theoretic advantage obtained from embodiment.

The account is set out in five sections. The next section (Section 2) provides the analysis of information refinement on which the measure is based. Section 3 presents an illustration of information refinement in practice. Section 4 examines the ways in which refinement modeling can be the means of quantifying benefits of embodiment. Section 5 examines related work, and Section 6 presents concluding comments.

2 Analysis of information refinement

For analysis of information refinement we start by envisaging data D to be a set of symbols drawn from an alphabet of n elements. Denoting the number of symbols in the data to be $|D|$, the total information content is then

¹This is one of several measures agent/environment information flow currently under consideration, cf. (Prokopenko *et al.* 2006.; Friston and Stephan, 2007; Friston, 2009).

²Generic and context-free, analysis of information refinement has potential applications in several other areas. Currently under investigation are applications to inductive generalization, grammar induction and data compression.

$$I(D) = |D| \cdot \log n \tag{1}$$

As a simple illustration, D might be the sequence ‘X Y X Y Z’. If symbols are drawn from an alphabet of 26 elements, the total information content is³

$$|D| \cdot \log 26 = 5 \cdot 4.7 = 23.5 \text{ bits}$$

For purposes of this analysis, no constraints are placed on the structure of D or the constructs that it exhibits. Where two or more constructs share the same structure, the union of those constructs can be referenced using certain conventions, however. Specifically, if x represents such a union, $|x|$ is the number of symbols utilized, and x_i is the set representing the choice of symbols for the i 'th element of the common structure.

Continuing to view the data as the sequence ‘X Y X Y Z’, constructs might be taken to be subsequences, such as ‘X Y’. Among the three-element subsequences, we would then have ‘X Y X’ and ‘X Y Z’. These two having the same structure, there is then a valid union. If x represents it, the relevant choices would then be $x_1 = \{X\}$, $x_2 = \{Y\}$ and $x_3 = \{X, Z\}$, with subscripts being indexes into the common structure. The shorthand used to represent the union would be ‘X Y X/Z’

Building on these definitions, it is possible to introduce D' , denoting a *reconstruction* of D . This is a modification of D in which some or its constructs are replaced with symbols representing unions. Replacement is deemed possible just in case the construct is within the represented union.

In the case of the sequence ‘X Y X Y Z’, we might have the reconstruction ‘\$0 X Y Z’, where \$0 is a symbol for the union ‘X Y/Z’. This is possible since the two element construct ‘X Y’ is within the union ‘X Y/Z’ (which combines ‘X Y’ and ‘X Z’). Where replacements introduce choice, there is loss of information, i.e., increase of uncertainty. This can be precisely quantified, however. The information loss (uncertainty) resulting from a replacement involving a particular union x may be calculated as

$$H(x) = \sum_i \log |x_i| \tag{2}$$

Equivalently, it may be calculated as the log of the combinatorial product of x 's choices:

$$H(x) = \log \prod_i |x_i| \tag{3}$$

The total information lost in a reconstruction can then be calculated by summing the information losses associated with its symbols:

$$H(D') = \sum_i H(D'_i) \tag{4}$$

³Logs are taken to base 2 throughout.

Here, $H(D'_i)$ is zero if D'_i is an original symbol, and the information loss of the represented union otherwise.

The total symbol cost of a reconstruction (the total number of symbols used) can also be defined. It is the number of symbols used in the modification itself, added to the total number of symbols used in referenced constructs. This is denoted $c(D')$:

$$c(D') = |D'| + \sum_{x \in D'} |x| \quad (5)$$

Here, $x \in D'$ enumerates the set of unions referenced by D' . As an illustration, the symbol cost of the reconstruction ‘\$0 X Y Z’ is

$$4 + 2 = 6$$

Combining the reconstruction loss with the reconstruction cost, it is then possible to define the informational efficiency of a reconstruction. This is the mean information content of its symbols, i.e., its net information content divided by symbol usage:

$$\bar{I}(D') = \frac{I(D) - H(D')}{c(D')} \quad (6)$$

The informationally optimal reconstruction of D is then that reconstruction that maximizes mean information. This is denoted $r(D)$:

$$r(D) = \operatorname{argmax} \bar{I}(D') \quad (7)$$

Note that the mean information of $r(D)$ can be no less than that of D itself. Were this to be the case, D would be its own optimal reconstruction. Given $r(D) \neq D$, it must be the case that

$$\bar{I}(r(D)) > \bar{I}(D) \quad (8)$$

which further implies that

$$c(r(D)) < |D|. \quad (9)$$

Increasing the mean content of symbols above the level they have in D itself must involve reducing their number. The optimal reconstruction must therefore use a lesser number of symbols than D itself.

The reconstruction $r(D)$ is the most informationally efficient expression of the content of D that can be built using D 's native constructs. This description can be made concrete in various ways, however. One possibility is to focus on how the optimal reconstruction eliminates representational inefficiency. Assuming $r(D) \neq D$, it must be the case that there is a more efficient way of deploying D 's native constructs for purposes of representing D , than D itself. Derivation of the optimal form of D can then be seen as a way of eliminating representational redundancy.

Another possibility is to focus on the way unions effect generalization. The process of generating symbolic labels for unions might be viewed as a form of induction. Use of information theory could then be seen as a learning heuristic. Consider for example the reconstruction ‘X \$0 \$0’ of ‘X Y X Y Z’, where \$0 is defined to represent the union ‘Y X/Z’. Derivation of this might be viewed as the process of learning a generalized representation for two different sequences.

Still another possibility is to focus on the reductive aspect of the process. Here we view what is accomplished as a kind of data compression — the aim is seen to be derivation of a more compact encoding. However, the form of data compression exhibited has an unusual form: it implicitly ‘chooses’ between lossy (i.e., information losing) and non-lossy strategies (Held, 1983) according to which produces the most informational efficient outcome.




Figure 1: The reconstruction tradeoff.

For a more visual interpretation, consider the schematic of Figure 1. Here reconstruction is viewed as the attempt to find a tradeoff between the benefits of increased symbol information, and the costs of symbol elimination. The underlying space is taken to be a graph, with the vertical axis representing mean-information achieved, and the horizontal axis representing the reduction in symbol usage. Treating the information content of original symbols as a baseline (see the dashed line), it is then possible to envisage a locus of points representing all those reconstructions that maximize mean content at a given level of symbol reduction. Derivation of the optimal reconstruction can then be viewed as the process of searching along this curve for its peak.⁴

2.1 Recursive refinement

Finding the informationally optimal reconstruction of a dataset can serve the goals of generalization and/or data compression. The potential of the process

⁴There may be more than one optimal reconstruction and thus more than one peak.

is more fully realized, however, when it is deployed in a *recursive* manner. The end-product of a single reconstruction is a structure of symbols. This is another dataset, for which we can derive a second, optimal reconstruction. The effect is then to exploit generalizations over symbols representing generalizations at a lower level of organization. Taking the process forward recursively, we obtain a *series* of reconstructions, the constructs of which capture generalizations at increasingly coarse-grained levels of organization.

Letting optimal reconstructions now be called ‘refinements’, the optimal reconstruction of an original dataset can be termed the ‘first refinement’, the optimal reconstruction of the reconstruction the ‘second refinement’, and so on. Any dataset then has a first refinement, second refinement, third refinement etc., with the total number depending on the constitution of the data. The requirement for an optimal reconstruction to show an increase of mean information (Equation 8) imposes a limit, however. The entire hierarchy of optimal refinements can then be defined using the following, recursive formula:

$$D^n = r(D^{n-1}) : \bar{I}(D^n) > \bar{I}(D^{n-1}) \quad (10)$$

Labeling the original dataset D^0 , this formula specifies the constitution of the first refinement D^1 , the second refinement D^2 , third refinement D^3 , and so on, up to the n 'th refinement D^n . The value of n is the *last* level at which the optimal reconstruction obtained has mean information greater than that of its source dataset. Beyond this level, further refinement is ruled out. The representation obtained at this level is thus the *root refinement* of the dataset. However, since optimal reconstructions are not necessarily unique, there may be more than one refinement hierarchy and thus more than one root refinement for any given dataset.

2.2 Worked example

Illustrating a complete cycle of recursive refinement is the schematic of Figure 2. This shows derivation of a hierarchy of three levels from the sequence dataset ‘b c d b e d b f d b g d’ (here named ‘eg1’). Possible primitives in this scenario are taken to be the eight characters {a, b, c, d, e, f, g, h}, on which basis the information content of original symbols is 3.0 bits. Constructs are assumed to be subsequences of any length.

The three levels of the tree-structure of Figure 2 correspond to the three levels of refinement obtained, with higher-level refinements appearing higher in the figure. At the bottom, we see the dataset itself, represented as a sequence of 12, oval-shaped nodes. (The numbers underneath are indexes.) Each one of these nodes encloses a representation for a particular element from the original sequence. In general, ovals represent symbols, with their contents representing the referenced construct. Arcs, where shown, identify locations of constituents. With eight primitive symbols in use, each element of the original dataset has an information value of $\log 8 = 3.0$ bits. Comprising 12 symbols in all, the dataset




Figure 2: Refinement hierarchy for the sequence ‘b c d b e d b f d b g d’.

then has a total content of 36.0 bits. (Annotations to this effect appear on the right of the figure.)

At the first level of the hierarchy, we see the initial refinement of the data. As in previous illustrations, this exploits the presence of approximately repeated subsequences. In this case, the repeats commence at indexes 0, 3, 6 and 9; the generalization is ‘b c/e/f/g d’. With the original dataset reconstructed as four repetitions of this construct (here labeled ‘\$0’) an information loss of 8.0 bits is incurred, arising as 2.0 bits per use of the construct. Subtracting this aggregate loss from the original 36.0 bits, we are left with a residual content of 28.0 bits. Dividing by the seven symbols in use (comprised of the three symbols of the construct itself, and the four uses of \$0), we obtain a mean information of 4.0 bits. Precisely one bit above the baseline of 3.0 bits, this is the maximum that can be achieved through native reconstruction with the original ‘eg1’.

At the next highest level level, the refinement obtained embodies no information loss at all. (This exemplifies the point made about optionality of generalization.) The entire 28.0 bits of content from the first refinement is retained. Repetition of the subsequence ‘\$0 \$0’ is exploited, however, through introduction of the construct labeled \$1. A further increase in mean informa-

tion is obtained. Finally, at the highest level, the root refinement is obtained. At 9.33 bits per symbol, mean information at this level is 6.33 bits above the baseline.

Application of refinement to sequential data is particularly convenient for purposes of exposition, due to the ease with which the process can be given a 2-dimensional representation. Choices can be represented as disjunctive strings such as ‘b c/e d’. However, it is important to remember that the model is fully generic and can be applied to data of any structure. It is not limited to sequences.

3 Refinement as a vehicle of measurement

Derivation of refinement hierarchies may serve various ends of an informational or representational nature (see below). For present purposes, however, the model is introduced as a way of measuring the information increase generated when data expressing certain content are put into a correlated form. As noted in the introduction, this can be a means of indirectly quantifying informational benefits obtained from physical interaction.

In the simplest case, the procedure works as follows. Referencing production of particular ‘good data’, we compare these against whatever data are deemed to be originally available to (or instantiated by) the generative process in question. Obtaining the refinement hierarchy for this uncorrelated counterpart, we then look for the refinement level that expresses or approximates the data in its correlated form. The difference in mean content between the original data and the relevant refinement is identified. Refinements being informationally optimal reconstructions, this difference is a lower bound on the increase in symbol information generated when the data are put into the correlated form. It is thus the minimum informational benefit generated for any consumer of the correlated data, given access otherwise limited to the uncorrelated form. If the production of the correlated data is accomplished through physical interaction between agent and environment, the increase is then the direct, informational benefit arising from that interaction.

In principle, this procedure can be used to measure informational benefits of physical embodiment in any scenario. In practice, various difficulties may arise. It may not be possible to give any formal characterization to the ‘good data’ in a particular case. It may not be possible to identify any uncorrelated counterpart. The refinement hierarchy obtained from the uncorrelated counterpart may not contain any representation of the correlated data, or the representation may be insufficiently exact. Last but not least, it may not be possible to form any clear judgement about the degree to which data production is mediated by physical interaction.

Generally in such cases there is the possibility of some form of approximation being applied, for purposes of reaching an estimated result. If the difficulty is that the refinement hierarchy obtained from the uncorrelated data does not explicitly represent the correlated data, there is the possibility of working in

terms of the refinement level that most nearly represents these data (i.e., exhibits the nearest number of distinct symbols). If no clear judgement can be made whether data production is an effect of embodiment, it may be possible to use some measure of the degree to which this is the case.

Consider, for example, a situation in which we want to measure informational benefits of embodiment in the case of a ‘passive-dynamic walker’, such as the Cornell device, built by Steve Collins, following (McGeer, 1993).⁵ Mechanical assemblies that approximately reproduce the skeletal structure of a bipedal agent, these devices exploit the natural dynamics of mechanical linkages in the production of stable walking. No control mechanism or motive power is used, other than the force of gravity.

For purposes of quantifying informational benefits arising from embodiment in such a case, it is necessary, first, to conceptualize the benefit as arising through production of correlated data. There are several ways to proceed in the case of a passive-dynamic walker, depending how we view the constituency of the agent. Particularly critical in the behaviour, is translation of the non-supporting foot in each cycle. The foot must be sufficiently advanced from the current center of gravity, to ensure an effective ‘step over’ and continuation into the next cycle. A plausible approach is thus to treat production of correlated data as a process that transforms data representing the current state of the agent into a signal requesting increase or decrease in forwards translation of the non-supporting foot. On this basis, production of correlated data can be seen as the process of deriving a refinement for a state representation that uses precisely two symbols.

To illustrate how this might work in practice, consider the following listing to be a dynamical systems representation for the state of a passive-dynamic walker at a particular instant in time. The values can be viewed as representing angles of extension (or flex) for knee, hip, shoulder and torso joints, along with relevant parameters of the context.

.9 .6 .7 .7 .2 .6 .5 .4 .1 .7 .5 .9 .3 .2 .8 .2 .7 .1 .3 .2 .7 .6 .1 .4

Effective continuation of the behaviour requires the non-supporting foot to be advanced appropriately in each cycle. In practice this is accomplished through physical effects. Envisaging it to involve production of data embodying one of two calibration signals, however, we can then estimate the informational benefit obtained by finding that level in the relevant refinement hierarchy that deploys just two symbols.

On the assumption that native constructs in this dataset are collections (i.e., subsets) of state values, information refinement should rightly explore reconstructions based on that type of structure. As a simplifying alternative, however, we can view the state description as a sequence, in which the values appear in the the given left-to-right order. This has the advantage of enabling the refinement hierarchy obtained to be represented using the 2-dimensional schematic used previously, as in Figure 3.

⁵In June 2010, details for this walker were available at <http://ruina.tam.cornell.edu/hplab/pdw.html>.

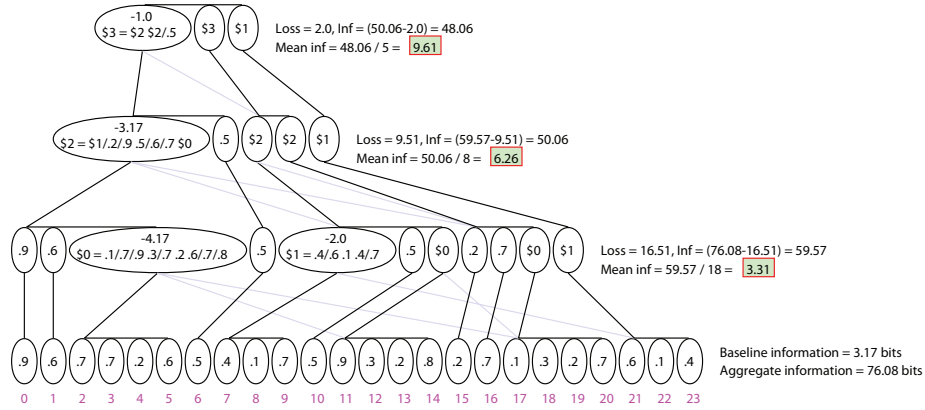


Figure 3: Refinement hierarchy for the passive-dynamic walking data.

Referencing this hierarchy, we find a reconstruction involving use of just two symbols at the third refinement level. Mean information content at this level is 9.61 bits. The mean content of the original data is 3.17 bits. These figures indicate the information increase resulting from putting the data into the required correlated form to be in the region of $9.61 - 3.17 = 6.44$ bits. Production of ‘good data’ in (this interpretation of) passive-dynamic walking being entirely a result of physical interaction, effects of embodiment are then found to produce an informational benefit of 6.44 bits per correlated symbol. This is a more than doubling of the original symbol content of 3.17 bits, reflecting the centrality of physical interaction in this behaviour.

Application of the method to natural contexts works in much the same way. Consider use of pheremone trails in ant foraging. In this example, a critical behavioural function emerges out of a combination of embodiment, scaffolding and multi-agent cooperation (Pfeifer and Bongard, 2007). Ants foraging for food leave pheremone trails marking their route, while always choosing the direction of highest pheremone concentration when encountering any existing trails. Over time, routes to the biggest food concentrations get increasingly marked out, with ants then being led ever more effectively towards sources of food (Deneubourg and Goss, 1989; Goss *et al.* 1989).

In this emergent solution to a navigation problem, exploitation of a bodily property (pheremone production) provides agents with an informational benefit (detection of food sources). Can the refinement method be used to measure the benefit? Taking the correlated data in this scenario to be behavioural cues triggering movement in the direction of ‘most frequently followed route’, the corresponding uncorrelated data might then be seen as the general pattern of sensory stimulation obtained by ants at trail-crossings. Production of ‘good data’, on this view, would be the process of generating appropriate cues from general sensory stimulation.

For purposes of producing this good data, ants might conceivably use internal information processing. In practice, ‘cheap design’ applies — physical interactions serve to produce cues in a easily detected, olfactory form. To calculate the benefit we should then find the level in the originating dataset’s refinement hierarchy offering the most reasonable representation of the correlated data. The informational benefit obtained can then be estimated in terms of the relevant enhancement of symbol content.

Taking this approach to quantification of pheromone-trail benefits, it is natural to equate the uncorrelated data (from which good data are deemed to be derived) with sensory stimulation. This produces an interpretation in which we see relations of embodiment directly substituting for informational processes that might conceivably be performed independently by the agent. This produces a nice result, in which embodiment is seen to replace internal information processing more or less directly. It is not always possible to understand the effect in such neat terms, however.

Requiring a quite different interpretation are the ‘Swiss Robots’ of (Maris and Boekhorst, 1996). On the basis of direct connections from just two infra-red proximity sensors, these 2-wheel robots are able to produce effective clustering of styrofoam blocks. The result seems mysterious given the absence of on-board information processing. However, it can be explained quite easily. Pfeifer and Bongard (2007, p. 74) take up the story. ‘How is it possible that clusters are formed? They come about because the sensors are placed sufficiently far apart so that if a robot encounters a cube head-on, neither sensor fires (provides stimulation), so the robot simply drives forward, thereby pushing the cube. When another cube appears on its side, it turns away, leaving the cube next to the other one. This process, when repeated, leads to clustering.’

Playing a critical role in this scenario is the (body/world) relationship between the width of the styrofoam blocks and the separation of the two IR sensors. Relations of embodiment directly underpin what we interpret as clustering behaviour. But in using the refinement method for purposes of measuring the benefit obtained, care must be taken in conceptualizing production of data. Arguably, the good data in this case must correlate with the decision to ‘turn away’, where stimulation indicates the possibility of a cluster. But it makes no sense to consider the originating data to be the sensory stimulation obtained by the agent. After all, we know that a crucial aspect of the behaviour is executed in the absence of any stimulation at all. In this case, it would be more appropriate to equate the originating data with the general agent/environment state out of which the response emerges.

3.1 Varieties of disembodiment

Referencing information differences in refinement hierarchies can be a way of assessing informational benefits of physical interaction. The natural application of the procedure is to show the extent to which an agent is behaviourally embodied. But it can just as easily be used to quantify the opposite effect. Demonstrating that the informational benefits secured through embodiment are minimal or

non-existent effectively establishes a certain degree of *disembodiment*.

We can use the present framework to divide this attribution into two cases, however, corresponding to the two ways in which informational benefits may be secured *without* involvement of physical interaction. First, we have the familiar scenario in which an agent is deemed to produce correlated data independently, without exploiting physical interaction. This is the ‘classical’ scenario, in which benefits are obtained through internal information processing. From the present perspective, we should call it *local disembodiment*, its key feature being use of processes that are local to the agent.

Contrasting with this, there is what we might term *remote disembodiment*. In this scenario, informational benefits are obtained as a result of environmental dynamics, but with the reservation that these operate independently of the agent. There is production of ‘good data’; the agent secures informational benefits. But the process by which the data are generated does not itself involve agent/environment interaction. This scenario challenges any simplistic dualism, in which exploitation of embodiment is seen to be the opposite of ‘cognitivism’. Strategies of remote disembodiment cannot be grouped in the cognitivist category, since they do not involve agent-based information processing. Neither are they truly embodied, since they do not involve exploitation of agent/environment interaction.

4 Related work

Attention turns now to consideration of connections between the present proposal and other work. Information theory has been applied to the problem of quantifying benefits of embodiment that arise through (potentially implicit) production of ‘good data’, in the manner envisaged by Pfeifer and Bongard (2007), Clark (2008) and others. Pursuing this line, the paper forges a close relationship with the rapidly expanding body of work on ‘information dynamics’ that investigates informational optimizations applied to sensory-motor relationships (e.g., Lungarella and Pfeifer, 2001; Lungarella *et al.* 2005; Sporns and Pegors, 2003; Sporns and Pegors, 2004; Lungarella and Sporns, 2005; Lungarella and Sporns, 2006; Klyubin *et al.* 2007; Prokopenko *et al.* 2006; Friston and Stephan, 2007; Friston, 2009). However, as noted in the introduction, approaches that involve optimizations of ‘information flow’ between environment and agent (such as *empowerment*) cannot easily take account of informational benefits arising independently of the agent. A context-free approach is required to achieve that end.

Perhaps the closest direct counterpart of the present proposal is to be found in the concept of *morphological computation*. This concept gives potentially quantitative evaluation to the benefits that are obtained when agents ‘off-load some neural processing into their morphology’ (Pfeifer and Bongard, 2007, p. 361). However, it remains unformalized at the present time. ‘In spite of its high intuitive appeal,’ Pfeifer and Bongard note, ‘the concept of morphological computation still awaits qualification: how much computation is really done by

a spring in the joint, or a change in morphology? Or perhaps this is not the right question.’ (ibid., p. 361-362)

Approaches focusing specifically on assessment of embodiment advantage have generally adopted an informal or task-specific focus. Investigating ‘ways of using spatial arrangements to informationally jig or structure the environment’, Kirsh (1995, p. 39) has examined cognitively advantageous ways of using location and physical sequencing. At first glance, his results are consistent — even mutually reinforcing — with the present analysis. He proposes, for example, that deliberate manipulation of embodiment can sometimes be understood as an attempt to ‘encode as explicitly as possible ... a key piece of information about the problem state’ (ibid., p. 32). He highlights the critical importance of ‘how the agent organizes the stimulus material internally’ (p. 33) and the way in which structuring actions can ‘serve to reduce the descriptive complexity of our environments’ (p. 65) in terms of creation of ‘critical cues’ (p. 66).

Exploration of information-theoretic possibilities is not pursued, however, and Kirsh ultimately reaches a negative conclusion whether embodiment can be subject to meaningful quantification. ‘Although I have been arguing that the point of informationally structuring space is to reduce the time and memory requirements of cognition,’ he notes, ‘the actual reconstruction in computation achieved by the various methods I shall discuss does not, in general, lend itself to meaningful quantitative estimation’ (p. 41).

Other investigators are more optimistic. In a series of studies Paul Maglio and colleagues have examined exploitation of embodiment in Tetris game play (Maglio and Kirsh, 1996; Maglio and Kirsh, 1996; Maglio and Wenger, 2000; Maglio *et al.* 2003; Maglio *et al.* 2003). As expertise in this game is acquired, players tend to engage in greater amounts of explorative piece-rotation. The hypothesis then put forward (Maglio *et al.* 2003) is that these rotations are *epistemic actions* (Kirsh and Maglio, 1994) pursued for their cognitive reward.

To evaluate the hypothesis, it is necessary to compare the costs of performing rotations against the benefits obtained, a procedure that necessitates mapping both costs and benefits into the same unit. This was done using the notion of response time (RT). The benefits of piece-rotation (measured in terms of impact on RT) were found to far outweigh the costs, as predicted. Other authors have also found cost-benefit analysis to be of use in assessment of embodiment-related advantages (e.g., Tarjan, 1985).

In another application of the cost-benefit approach, Gray and Fu (2004) compared the ‘effort’ involved in memory-accesses against the amount involved in functionally equivalent perceptual inspection of the environment. Their hypothesis is in two parts: ‘1. The lower the effort of sensory-motor access relative to memory retrieval, the greater the reliance on perfect knowledge in-the-world’ and ‘2. Contrariwise, the higher the effort of perceptual-motor access relative to memory retrieval, the greater the reliance on imperfect knowledge in-the-head.’ (ibid., p. 367). In other words, agents will choose the least-cost strategy where there is optionality between a memory based strategy and an environment-based strategy. Like Maglio *et al.*, they found that human agents seem to behave rationally, exploiting available agent/environment dynamics whenever that strategy

offers tangible benefit. Through derived in a completely different way, this result is consistent with the hypothesis of agent/environment interaction being pursued in conformity with its degree of *informational* benefit.

The present proposal also makes connections with several areas of generic information analysis. As noted, information refinement can be viewed as modeling derivation of parsimonious representations, on which basis it is related to the paradigm of Kolmogorov complexity (Kolmogorov, 1965; Li and Vitányi, 1997; Chaitin, 1979). This asserts a formal connection between representation simplification and order, echoing the principle of Occam’s Razor (Tornay, 1938). To a first approximation, the Kolmogorov complexity of a dataset is the size of its smallest computational representation. Viewing optimal refinements as optimally *small* representations, there is then a way of understanding the present proposal as a Kolmogorov method for modeling inductive efficacy (cf. Solomonoff, 1964b; Solomonoff, 1964a), or measuring biological complexity (cf. Chaitin, 1979). In reality, however, refinements are not optimally small representations. They are optimally efficient ones, although one property will often imply the other. Moreover, they are constructed in terms of native constructs of the context, rather than unrestricted computational resources.

Where refinement entails generalization, there is also the potential to interpret the procedure as a method of inductive learning. Here the connection is with compression-mediated learning approaches, such as Minimum Description Length learning (Rissanen, 1985.; Rissanen, 2007) and the complexity oriented approach of (Schmidhuber, 1996). Finally, there is the potential to view refinement as an ordinary method of data compression; although here we must note the way it self-selects between the lossy and lossless strategies in an informationally informed manner.

5 Concluding comments

In the view of Pfeifer and Bongard (2007), informational benefits arising from exploitation of embodiment are mediated by the processes by which ‘the brain gets *good* — that is, correlated — raw material for further processing’ (p. 360, *my emphasis*). Quantifying the benefit obtained then entails assessing the degree to which relevant data present useful correlations. In the proposal of the present paper this is accomplished in a context-free way using formal analysis of information refinement. By discovering the sequence of optimal refinements that serves to put uncorrelated data into a desired, correlated form, it is possible to gauge the increase in symbol information that is thereby accomplished. This then quantifies the benefit obtained. The main contributions of the paper are thus an information-theoretic analysis of the way in which benefits can emerge from embodiment, and a quantitative model for measuring the degree of gain. Applying the method for practical assessment of agent/environment situations is also a possibility. Obstacles involving identification of relevant bodies of data can generally be overcome through some form of approximation or estimation.

The analysis also raises some interesting theoretical questions. One of these

relates to the ontological status of refinement hierarchies, in relation to the information dynamics and IDSO paradigms. In Lungarella and Sporns' view 'the presence of informational structure in the sensory data' is of 'pivotal importance' (Lungarella and Sporns, 2005, p. 25). The same point is also emphasized in the present proposal. But in light of the analysis of Section 2, it is natural to ask how the information structure envisaged by Lungarella and Sporns relates to the information-refinement structure defined by Equation 8.

Lungarella and Sporns take information structure to be something that is mediated by sensory-motor coordination. They argue that 'informational structure in the sensory experience of (natural and artificial) embodied systems is at least in part the result of effectively coordinated motor activity' (ibid., p. 25). Under the present approach information structure is an implicit structure, that is explicated through recursive refinement. On the face of it, these positions cannot be reconciled, suggesting that refinement structure and 'information structure' are fundamentally different. However, it would be interesting to ascertain whether they have ground in common, and if so to what degree. Obtaining further clarification on this point will be a goal for future work.

References

- Adams, F. and Aizawa, K. (2008). Why the mind is still in the head. In M. Aydede and P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press.
- Agre, P. and Chapman, D. (1991). What are plans for?. *Designing Autonomous Agents: From Biology to Engineering and Back*. Cambridge, MA.
- Ballard, D. (1991). Animate vision. *Artificial Intelligence*, 48 (pp. 57-86).
- Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72 (pp. 173-215).
- Beer, R. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4 (pp. 91-99).
- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47 (pp. 139-159).
- Chaitin, G. (1979). Toward a mathematical definition of life. In I.R. Levine and M. Tribus (Eds.), *The Maximum Entropy Formalism* (pp. 477-498). MIT Press,.
- Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action and Cognitive Extension*. Oxford: Oxford University Press.

- Clark, A. and Chalmers, D. (1998). The extended mind. *Analysis*, 58, No. 1.
- Deneubourg, J. and Goss, S. (1989). Collective patterns and decision making. *Ethology, Ecology, and Evolution*, 1 (pp. 295-311).
- Friston, K. (2009). The free-energy principle: a rough guide to the brain. *Trends in Cognitive Sciences*, 13, No. 7 (pp. 279-328).
- Friston, K. and Stephan, K. (2007). Free-energy and the brain. *Synthese*, 159 (pp. 417-458).
- Goss, S., Aaron, S., Deneubourg, J. and Pasteels, J. (1989). Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, 76 (pp. 579-581). Springer-Verlag.
- Gray, W. and Fu, W. (2004). Soft constraints in interactive behavior: the case of ignoring perfect knowledge in the world for imperfect knowledge in the head. *Cognitive Science*, 28, No. 3 (pp. 359-382).
- Held, G. (1983). *Data Compression: Techniques and Applications, Hardware and Software Considerations*. Chichester: Wiley.
- Hurley, S. (1998). Active perception and vehicle externalism. In S.L. Hurley (Ed.), *Consciousness in Action*. Harvard University Press.
- Kirsh, D. (1995). *Artificial Intelligence*, 73 (pp. 31-68).
- Kirsh, D. and Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18 (pp. 513-519).
- Klyubin, A., Polani, D. and Nehaniv, C. (2007). Representations of space and time in the maximization of information flow in the perception-action loop. *Neural Computation*, 19 (pp. 2387-2432).
- Kolmogorov, A. (1965). Three approaches to the quantitative definition of information. *Prob. Inf. Trans.*, 1 (pp. 1-7).
- Li, M. and Vitányi, P. (1997). *An Introduction to Kolmogorov Complexity and Its Applications: Second Edition*. New York: Springer-Verlag.
- Lungarella, M., Pegors, T., Bulwinkle, D. and Sporns, O. (2005). Methods for quantifying the information structure of sensory and motor data. *Neuroinformatics*, 3, No. 3 (pp. 243-262).
- Lungarella, M. and Pfeifer, R. (2001). Robots as cognitive tools: information theoretic analysis of sensory-motor data. *Proceedings of the 2001 IEEE-RAS First International Conference on Humanoid Robots* (pp. 245-252). Piscataway, NJ: IEEE Press.
- Lungarella, M. and Sporns, O. (2005). Information self-structuring: key principles for learning and development. *Proceedings 2005 4th IEEE International Conference on Development and Learning and Learning* (pp. 25-30).

- Lungarella, M. and Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLoS Comput Biol*, 2, No. 10.
- Maglio, P. and Kirsh, D. (1996). Epistemic action increases with skill. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 391-396). Mahwah, NJ: LEA.
- Maglio, P. and Wenger, M. (2000). Two views are better than one: epistemic actions may prime. *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.
- Maglio, P., Wenger, M. and Copeland, A. (2003). The benefits of epistemic action outweigh the costs. *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Maris, M. and Boekhorst, R. (1996). Exploiting physical constraints: heap formation through behavioral error in a group of robots. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1655-1660). Piscataway, NJ: IEEE press.
- McGeer, T. (1993). Dynamics and control of bipedal locomotion. *Journal of Theoretical Biology*, 16, No. 3 (pp. 277-314).
- Pfeifer, R. and Bongard, C. (2007). *How the Body Shapes the Way We Think*. Cambridge, Mass.: MIT Press.
- Pfeifer, R., Lungarella, M. and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318 (pp. 1088-1093).
- Prokopenko, M., Gerasimov, V. and Tanev, I. (2006). Evolving spatiotemporal coordination in a modular robotic system. *From Animals to Animats 9: 9th International Conference on the Simulation of Adaptive Behavior* (SAB 2006). Rome, Italy.
- Rissanen, J. (1985). Minimum-description-length principle. *Ann. Statist*, 6 (pp. 461-464).
- Rissanen, J. (2007). *Information and Complexity in Statistical Modeling*. Springer.
- Schmidhuber, J. (1996). A theoretical foundation for multi-agent learning and incremental self-improvement in unrestricted environments. In X. Yao (Ed.), *Evolutionary Computation: Theory and Applications*. Singapore: Scientific Publishing Company.
- Solomonoff, R. (1964a). A formal theory of inductive inference, part II. *Information and Control*, 7, No. 2 (pp. 224-254).
- Solomonoff, R. (1964b). A formal theory of inductive inference, part i. *Information and Control*, 7, No. 1 (pp. 1-22).

- Sporns, O. and Pegors, T. (2003). Generating structure in sensory data through coordinated motor activity. *Proceedings of the International Conference on Neural Networks* (p. 2796).
- Sporns, O. and Pegors, T. (2004). Information-theoretic aspects of embodied artificial intelligence. *Embodied Artificial Intelligence: LNCS 3139* (pp. 74-85). Heidelberg: Springer-Verlag.
- Tarjan, R. (1985). Amortized computational complexity. *SIAM J. Computing*, *14* (pp. 306-318).
- Thelen, E. and Smith, L. (1993). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Mass.: MIT Press.
- Tornay, S. (1938). *Ockham: Studies and Selections*. La Salle, IL: Open Court.
- Varela, F., Thompson, E. and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Massachusetts: The MIT Press.
- Wheeler, M. (2005). *Reconstructing the Cognitive World: The Next Step*. MIT Press.