

Complexity and Context-Dependency

Bruce Edmonds

Centre for Policy Modelling, Manchester Metropolitan University

Abstract. It is argued that given the “anti-anthropomorphic” principle, that the universe is not structured for our benefit, that modelling trade-offs will necessarily mean that many of our models will be context-specific. The context heuristic, that divides the processing into rich, fuzzy context-recognition and crisp, conscious reasoning and learning is outlined. The consequences of accepting the impact of this human heuristic in the light of the necessity of accepting context-specificity in our modelling of complex systems is examined. In particular the development of “islands” or related model clusters rather than over-arching laws and theories. It is suggested that by accepting and dealing with context (rather than ignoring it) we can push the boundaries of science a little further.

Introduction – *The Exception of Simplicity*

Many times in the past, people have assumed that they occupy a special place in the universe or that the universe is somehow arranged to suit them, only for this to turn out to be wrong. Examples include believing that the Earth was the centre of the Solar System, or that homo sapiens had a special origin, different from than that of other animals. Today many believe the following:

That the universe happens to be structured on grounds of sufficient simplicity that our brains are able to analyze and comprehend it.

Or, to put it another way:

That our brains happen to have evolved so as to be able to analyze/understand models adequate to the phenomena we observe.

We may look down on other animals, perceiving that they have a biased and limited understanding of the world, but somehow assume that we don't have analogous biases or limitations that we cannot somehow overcome. Surely this is merely another example of anthropocentric arrogance. That we have had some notable successes at understanding our world and even a systematic set of approaches that has been shown to be useful is not sufficient evidence to assume a lack of limitations and biases.

This astonishing assumption takes many forms in philosophy and discussions about the scientific method. One such is that somehow simplicity is a guide to truth. That is, that simplicity in a model or theory has advantages other than the obvious pragmatic ones (pragmatic virtues are such as: being able to analyze/solve it; being

2 Bruce Edmonds

able to have good analogies with which to think about it; needing less data in order to parameterize it; and being able to compute it). Elsewhere I have argued that simplicity does not indicate truth [7].

Another version is that everything somehow *must* be simple if only we can find the right way of looking at it, or formalizing it. It is true that frameworks such as Newtonian Physics are relatively simple (though I doubt many in Newton's time would have thought so), and *using this*, many useful models and reliable predictions can be obtained. However, to make it actually cover observed cases and encountered problems one has to complicate it immensely in order to apply it. Even its general scope and simplicity *in theory* are questionable, with the (in hindsight inevitable) complications to cover new situations (Einsteinian, Quantum).

I am not going to spend time arguing the above points here. Rather I will consider the case under the *anti-anthropocentric assumption*, that much of the world around us is organized in a way that is beyond adequate modeling in a sufficiently simple and general manner for us to cope with¹. Where we accept that we have biases and limitations in our abilities to develop our understanding of the phenomena that confront us and, by acknowledging these and thinking about them, consider ways we might extend our scientific understanding to the greatest possible extent. Under this, admittedly pessimistic, view the phenomena that are simple enough for us to understand in a scientific manner are the exception – the exception to be sought and struggled for. Under this view, we should make the greatest use of the strengths we have, and seek to acknowledge and mitigate our limitations. Under this view a “Science of Complexity” makes no more sense than a “Science of Non-Red Things”, since both red objects and simple systems are the exception rather than the rule.

Some Modelling Trade-offs

In any modelling enterprise, when one is trying to represent some observed phenomena in a model (analytic, statistical, or computational) then there are inevitably some compromises and trade-offs one has to make. Figure 1, below, illustrates one possible set, that I will discuss below.

¹ Of course, “adequate” here invites queries as to what this might involve. A model is only adequate or otherwise when compared against the goals for the model. Clearly one of the responses to a overly complex universe is to be less ambitious in our modelling goals.




Fig. 1. Some modelling trade-offs

Clearly the available modelling trade-offs are never as simple as Figure 1 suggests. However, it is rarely the case that one can obtain *all* these four desiderata simultaneously, that is find a model that is: formal, simple enough to analyze and completely understand, that is valid in terms of its specification and results compared to the evidence and is fairly general in its scope. I will briefly discuss each of these four criteria in turn.

Validity. There are many ways in which a model may be judged as valid. This depends somewhat upon the goal for the modelling, that is how good it is for its intended use. In order for a model to be useful for any particular purpose it may have to be adequate in terms of a number of different criteria. For example, to be useful for prediction a model has to be: effectively initialisable in terms of parameters and settings that can be measured or determined from available evidence, computable in that one can somehow derive the predictions from the settings, and the results need to correspond to previously unknown observations to a sufficient degree and to a sufficiently reliable extent (given the purpose of the model). To be useful for establishing a possible explanation for some observed phenomena a model has to be: formulated in terms of the kinds of processes that one wants the explanation to be expressed in; to match the known outcomes that are to be explained; and to computationally trace out how the outcomes arose from the processes. However, whatever the purpose, it is clear that it is fundamental to science that a model is as valid as possible.

Formality. A formal representation is one that can be precisely specified. Thus both mathematical and computational models are formal (the later not being analytic). Not all scientific models (in the broadest sense) are formal, there are plenty of informal theories, descriptions and arguments which are part of science. However formal models play a central role in science, they form the 'backbone' upon which less formal aspects hang, allowing an unambiguous frame of reference for meaning of terms and facilitating a inter-scientist development which would be infeasible if based on less copiable entities. It also provides a basis for the systematic variation of possible explanations of the phenomena we observe. Thus, after validity preserving the formality of our models is of greatest importance.

Simplicity. Simple models are those that are easy to build, test, understand and analyze². If we accept the anti-anthropocentric principle then this implies that we may have to accept that the simplest adequate model (for a given purpose) might be very complex³. There are things one can do about such irreducible complexity but discussing these is not the prime purpose of this paper. Some of the suggestions later in the paper will touch upon ways around complexity. Here we just consider the situation where we have exhausted simplification techniques, and have to look elsewhere to progress our scientific understanding.

Generality. Generality is the extent of the scope of a model, that is, how many different situations or cases a model will be usefully applicable to. If one has a general model one can easily produce less general models by fixing some of its parameters or settings to those appropriate for some specific circumstances, thus making it less abstract. Thus a general model of gravitation might be simplified to one where gravity only works in one direction (downwards) and with one value (9.8 m/s^2), so it is specific to localized areas on the surface of the Earth. However doing the reverse, adding in processes and variables for fixed constants in a model, or by making a model *more* abstract, does not necessarily lead to a more general model, since one does not necessarily know what processes one should add or which abstractions one can make. Indeed, going this way can lead to a loss of generality, since the wrong additions or abstractions can mean it is not applicable anywhere.

Given that validity is essential to science, and formality highly desirable we are left with generality as the “free” aspect we can work on in order to progress scientific modelling in the face of overly complex phenomena. That is, one way we might be able to make do with sufficiently simple models is by reducing the scope of the models – making them more context-specific. In other words, sacrificing generality to gain some simplicity, in the context of obtaining formal models that are sufficiently valid for our needs. This is worth it, for although generality is a highly desirable property of models, like simplicity, it is a pragmatic virtue. Without wide generality we still have science (albeit more fragmented, resembling biology more than physics), but without formality we can not have progressive science as we know it, and without validity we don’t have science at all.

The Context Heuristic

Many aspects of human cognition and behaviour are context-dependent to some degree, including: visual perception, choice making, language comprehension [15], memory, reasoning [16], emotion, judging trustworthiness [21] and assessing reputation [20]. Thus, it appears that the brain uses context-dependency as an effective heuristic for dealing with the complexity of the world around it. In very broad terms, it appears that we recognise context in a rich, automatic, imprecise, unconscious but reliable manner. With each recognised context there is an associated resource of learnt facts, terms, behaviours, norms etc. that are accessible from the

² Oddly, many definitions of simplicity in the philosophical literature have *not* been about a lack of complexity but rather attempt to redefine simplicity in terms of pragmatic virtues [4].

³ Or, to put it another way, that simplicity is not, in itself, a guide to truth [7].

context. This resource makes reasoning, modelling, communicating, understanding etc. much easier since the context allows access to the most relevant information in a structured manner. Within a context we tend to employ more crisp, formal, deliberate and conscious reasoning and learning. Thus context seems to allow the integration of “fuzzy” pattern recognition with “crisp” reasoning mechanisms.

This heuristic is contingent – that is, for it to work, there are a number of conditions that have to hold.

1. The domain has to be usefully divisible into a number of contexts, that is each context must have clusters of related knowledge associated with them, clearly each cognitive context must cover a (maybe small) number of situations which have some underlying commonality [3].
2. The contexts have to be recognisable in a reliable manner (though not necessarily definable in precise terms).
3. Learnt knowledge must be associated with the context so that it is effectively retrievable when the context is encountered again.

For this heuristic to be most effective then the following are should hold.

4. That the contexts can be learnt, so that what is a context can be co-learnt with others so that everybody will recognise the same context.
5. That different contexts can be associated with the same situation, so that the relevant context can be flexibly swapped for another applicable one that might be more helpful (e.g. in case that the wrong context has been chosen or decision making is not feasible [6]).

There are many advantages to this heuristic. Dividing the world up into a number of contexts means that both explicit learning and reasoning happen within constrained sets of knowledge, making both of these processes computationally feasible [6]. In other words, context “solves” the frame problem [18], for without a way of delimiting the possible knowledge one could apply efficient reasoning is impossible [13]. If one is faced with either an under- or over-determination of knowledge with regard to a particular decision then a more or less specific context (respectively) can be selected [6]. The same problem can be thought of from a number of different cognitive contexts, allowing different perspectives to be applied (for example it may help to understand the otherwise surprising reaction of someone if one take the context to be a competition rather than a discussion). Since the heuristic is shared between people then this can facilitate the co-determination of appropriate contexts which aids communication and coordination. Context allows the integration of fuzzy pattern recognition and crisp reasoning processes in a coherent manner [11].

There is a difficulty with talking about or analysing context, due to the rich and unconscious way it can be recognised. Despite the fact that everybody can reliably recognise a context does not mean that the definition of a context is precisely definable or even identifiable. Everybody recognises the difference between the living and the inanimate, but it does not mean that it is easy to precisely define *life*. Context seems to be recognised automatically, so to a large extent, we may be unaware of its identification – it simply may not be easily accessible to as scientists. Thus although I talk about contexts in this paper, as if they are well-defined entities, it may well not be possible to unambiguously talk about *the context* for any particular case being discussed. “*The context*” is an abstraction of the class of situations that

would be recognised as such and to which the same set of existing knowledge would be brought to bear, but this does not mean that context can be reified as a well-defined labelled entity. Thus the statement “*A is true in context C*” is not formalised as $C \rightarrow A$ as suggested in [17] since *C* cannot be usefully reduced to a precise statement.

Given this analysis, it should be clear that context-dependency is not the same as subjectivity. It might well be that all relevant people can correctly identify the correct context and, given that, reliably assess the truth of a statement, where its truth depends on that context. Subjectivity would mean that each person might come up with a different assessment. Thus the *in vitro* vs. *in vivo* distinction is well understood in biology, and which is being assumed is usually abundantly clear, even if not stated explicitly. A statement’s truth may be specific to the *in vitro* context without this either being subjective or necessarily being true generally (e.g. also in the *in vivo* context). Also the meaning of a statement might well be context-dependent, without the meaning being subjective for the same reasons, for example the idea of a “behavioural norm” does not make sense when considering inanimate objects. Thus context-dependent statements are not inherently unscientific. In the case that the same context can not be reliably identified by all participants, but varies with each observer, the context-dependency of a statement may reduce to a form of subjectivity.

However, due to the fact that it can be difficult to identify and talk about *the* appropriate context of any statement this can make statements more difficult to *formally* assess, as the assumptions it depends on in any formal analysis might be “rolled up” into the context. Digging out the assumptions in relevant contexts, can be part of trying to formally assess or test context-dependent statements, or as part of an attempt to generalise out of a specific context to a more general one.

Context and Scientific Modelling

Given the *anti-anthropomorphic* principle we are not going to be able to represent all the possible causal process in our models⁴. In any complex (e.g. biological or social) case there are an indefinite number of possible causes of any event. Indeed elsewhere I have argued that the very notion of a cause only makes sense if there is effectively a context to delimit it [8]. In any case, we have to choose what to explicitly include in our models and what are considered, irrelevant or unchanging [23]. Much of the usefulness of a model comes from being able to use the same model in a related, but different set of circumstances that the one it was initially built for. In order to be able to do this one needs to know that none of those factors that were not represented in the model will make the model invalid in the new circumstances. Sometimes the applicability of a model is described explicitly, but since the set of background assumptions is indefinitely large, this is impossible to do completely (what [22] call “causal spread”). Hence some of these background assumptions have to be ensured simply by the context – that is, there is enough *recognizable* commonality between the situation for which the model was constructed and the situation in which it is later applied, that we do not have to worry about more extreme counterfactuals. A

⁴ It is notable that in the structural model of causation in [19] it is a key assumption that one can list *all possible* causes, which are then sorted using the method described therein.

common way that humans use is the recognition of a common kind of extended (or cognitive) context⁵. This is illustrated in Figure 2.

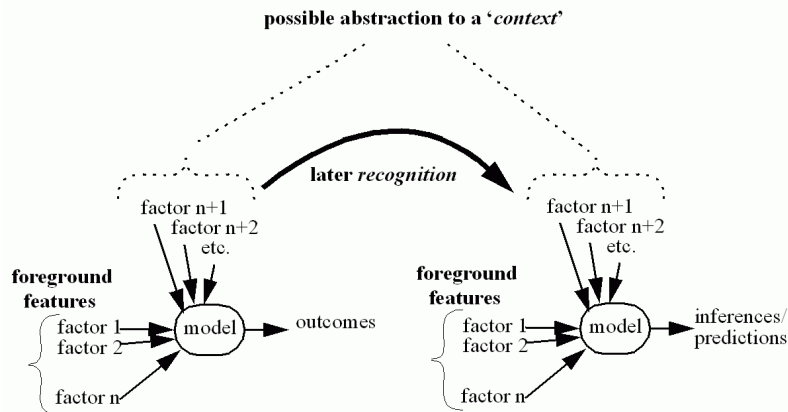


Fig. 2. Using a model to transfer knowledge from one situation to another which share context

As long as we can recognise *when* a model can be applied with a fair degree of reliability this is not a problem. However it is easy to take a model out of context, in which case the model may lose some or all of its validity. A change in context could involve a change in modelling purpose as well as the situation and sometimes that is more difficult to recognise from, say, its description in an academic paper. In this case there is a danger of the modelling being wrongly applied to this different purpose – a tendency that is particularly observed in the social sciences [10]. Sometimes when a model is being used as an analogy [5] (albeit in formal or computational form) this vulnerability to context change is masked, due to the context-sensitive manner in which people apply analogies, adapting meaning and reference almost automatically.

Of course, a model may still be useful when the influence of the unrepresented, background factors effect the results but not systematically or severely enough to completely invalidate them. This sort of “leakage” into a modelling context is often labelled as (exterior) noise [9]. Often we use a random “proxy” to stand in for this interference, for example, to test the sensitivity of the results to that factor. When what was thought to be “noise” turns out to have significant or systematic effects on the significant aspects of the results of a model, then this indicates that either the model needs expanding to include this aspect, or that the context is misidentified.

Regardless of whether one believes that “in principle” that a context-dependent approach to modelling is necessary, with many complex phenomena it is necessary in order to make our modelling feasible. It may be that formal, general and valid models

⁵ The context heuristic implies that similar kinds of situation are recognised as being essentially the same context from the point of view of the individual, e.g. a job interview. The power of the context heuristic results from being able to retrieve the relevant facts, norms, conventions etc. that pertain to that type. This is called the cognitive context, since it is the cognitive correlate of the exterior context, abstracting what is essential from the specific situation.

exist for many phenomena but the models are either too complex for us to deal with, or just be very difficult to discover. The reason we are forced to a context-dependent approach does not ultimately matter – the outcome is the same.

Thus to a large extent, modelling practice and human cognitive bias coincide and for similar reasons – it facilitates representation and reasoning when the phenomena happen to be usefully divisible into contexts. The same underlying trade-offs drive both in circumstances where this is possible.

Some Consequences of Complexity and Context-Dependency

If one accepts that one might have to settle for models that are either too complex to fully understand and/or are limited in scope to a specific context then this will have consequences for how one might have to approach using formal models.

If models are going to be more context-specific then more care will need to be taken to articulate and document as much of the information about the modelling context as possible. This is impossible to do completely, as has been discussed above, but what *is* known should be made explicit. The context the model is developed in should be described, and some indication of the range of contexts in which the model might be applied in, possibly explicating some of the commonality that is hypothesized to underlie its applicability. This could include some positive and negative examples of its scope and maybe a rule of thumb to guide other researchers. Talking about context is far more common in the social and ecological sciences and it may be that physicists and other complexity scientists will have to get used to this too.

It may well be that a set of models is the best way to try and represent some phenomena, rather than any single model [12]. This set might range over different aspects of the phenomena, consider it from different viewpoints, capture it at different granularities, or represent it at different levels of abstraction. This already occurs to some extent, for example in the development of separate explanatory and phenomenological models [2].

An example of this that of an ideal gas. The gas laws successfully model the relation between these triples of data, so that at the same pressure the volume is proportional to the temperature (measured from absolute zero). However the gas laws simply relate the phenomena, they do not explain *why* they are so related. A model of the gas composed of a sparse collection of molecules, can explain the gas laws, when the speed of the particles is related to the temperature. Using some judicious approximations one can derive the gas laws from this model. This model explains why temperature, pressure and volume are related, but does not itself express the relation between them – to infer these relations one needs extra assumptions. The mechanics of the particles in this model are consistent with Newton's laws of motion and the theory of the nature of temperature. Thus, even in this simple case one has, at least, two separate models: the phenomenological model and the explanatory model.

An ideal gas is an *exceptionally* simple case which can be dealt with using judicious assumptions of randomness, due to the lack of any significant interaction between its parts. In many other cases there is no adequate analytic approach to the modelling of molecule interaction – the only way to do this is via extensive

simulation. This approach is discussed in [14] and the relation between the models illustrated in Figure 3.

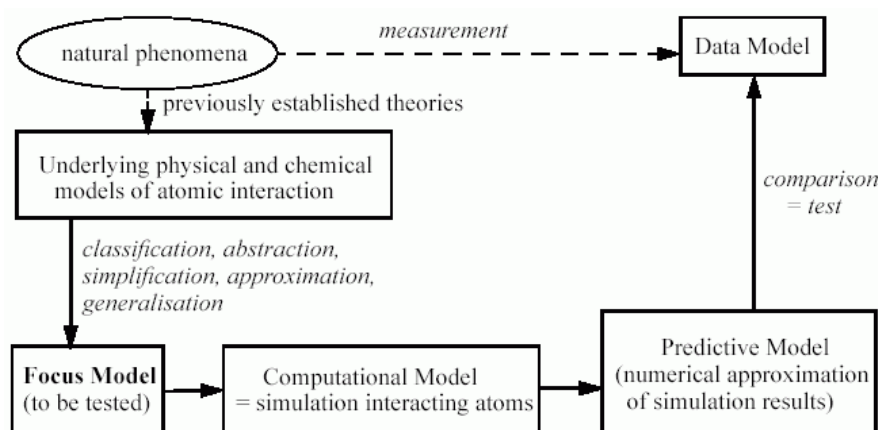


Fig. 3. A layering of related models in chemistry adapted from that in [14].

In this case the focus model concerns how a set of atoms interact (e.g. how the molecules pack or move in a certain state), this can be simulated using a representation of a set of atoms in 3D space using model of the underlying atomic processes. The simulation is run a great many times so that a summary of the results can be obtained by averaging etc. These results are further summarised into an equation that approximates the results, which is then compared to many different sets of measurements of the relevant natural phenomena. Thus here we see a cluster of at least five related models describing different aspects of the interacting molecules, such clusters will become the norm rather than single models as we tackle increasingly complex phenomena. Each member of the cluster might be validated in a different way (e.g. those in [1]) and related to the others in their cluster in a variety of manners. We know little about how to build, test and maintain such model clusters – this is something that needs attention.

Conclusion – Facing up to Complexity and Context-dependency

We may have to get used to the fact that there may not always be a general framework that can be assumed to relate all models together in a formal manner, but rather “islands” of model clusters that relate to each other, where each cluster is related to its own context. In other words, the theoretical landscape might just look a lot more “patchy” with only local areas of coherence, composed of a mess of context-specific models. Where possible, generality is desirable, but it will not be always achievable, but that does not invalidate the models. Context-dependency does not stop it being

science but, acknowledged and dealt with, it may open up new approaches allowing us to push the envelope of scientific knowledge a little further.

Acknowledgements

I would like to thank all those at the Centre for Policy Modelling, those at the International Conferences on Modelling and Using Context, and those who attend the informal Manchester complexity seminar series for many discussions on these issues.

References

- [1] Axtell, R. L. and J. M. Epstein (1994). Agent-based Modelling: Understanding our Creations. *The Bulletin of the Santa Fe Institute* 9: 28-32.
- [2] Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford, Clarendon Pres.
- [3] Edmonds, B. (1999) The Pragmatic Roots of Context. CONTEXT'99, Trento, Italy, September 1999. *Lecture Notes in Artificial Intelligence*, 1688:119-132.
- [4] Edmonds, B. (1999) Syntactic Measures of Complexity. PhD Thesis, University of Manchester, Manchester, UK. <http://bruce.edmonds.name/thesis/>
- [5] Edmonds, B. (2001) The Use of Models - making MABS actually work. In. Moss, S. and Davidsson, P. (eds.), *Multi Agent Based Simulation*, Lecture Notes in Artificial Intelligence, 1979:15-32.
- [6] Edmonds, B. (2002) Learning and Exploiting Context in Agents. *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Bologna, Italy, July 2002. ACM Press, 1231-1238.
- [7] Edmonds, B. (2007) Simplicity is Not Truth-Indicative. In Gershenson, C. et al. (2007) *Philosophy and Complexity*. World Scientific, 65-80.
- [8] Edmonds, B. (2007) The Practical Modelling of Context-Dependent Causal Processes. *Chemistry and Biodiversity*, 4(1):2386-2395.
- [9] Edmonds, B. (2009) The Nature of Noise. In Squazzoni, F. (Ed.) *Epistemological Aspects of Computer Simulation in the Social Sciences*. Lecture Notes in Artificial Intelligence 5466:169-182.
- [10] Edmonds, B. and Hales, D. (2005) Computational Simulation as Theoretical Experiment, *Journal of Mathematical Sociology* 29(3):209-232.
- [11] Edmonds, B. and Norling, E. (2007) Integrating Learning and Inference in Multi-Agent Systems Using Cognitive Context. In Antunes, L. and Takadama, K. (Eds.) *Multi-Agent-Based Simulation VII*, Lecture Notes in Artificial Intelligence 4442:142-155.
- [12] Giere R., N. (1988). Explaining science: a cognitive approach. Chicago ; London, University of Chicago Press.
- [13] Greiner, R., Darken, C. and Santoso, N.I. (2001) Efficient reasoning. *ACM Computing Surveys*, 33 (1). 1-30.
- [14] Gunsteren, W. F; Berendsen, H. J. C. (1990) Computational Simulation of Molecular Dynamics: Methodology, Applications and Perspectives in Chemistry. *Angewandte Chemie - International Edition in English*, 29:992-1023.
- [15] Hayes, P. (1995). Contexts in Context. Context in Knowledge Representation and Natural Language, AAAI Fall Symposium, November 1997, MIT, Cambridge.

- [16] Kokinov, B. and Grinberg, M. (2001) Simulating Context Effects in Problem Solving with AMBR. in Akman, V., Bouquet, P., Thomason, R. and Young, R.A. eds. *Modelling and Using Context*, Springer-Verlag, 2116:221-234.
- [17] McCarthy, J. (1971) Generality in Artificial-Intelligence - Turing Award Lecture. *Communications of the ACM*, 30 (12). 1030-1035.
- [18] McCarthy, J. and Hayes, P.J. (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463-502.
- [19] Pearl, J. (2000) *Causality*. Cambridge University Press.
- [20] Schlosser, A., Voss, M. and Brückner, L. (2005). On the Simulation of Global Reputation Systems. *Journal of Artificial Societies and Social Simulation* 9(1)4 <<http://jasss.soc.surrey.ac.uk/9/1/4.html>>.
- [21] Tykhonov, D, Jonker, C, Meijer, S and Verwaart, T (2008), Agent-Based Simulation of the Trust and Tracing Game for Supply Chains and Networks. *Journal of Artificial Societies and Social Simulation* 11(3) 1 <http://jasss.soc.surrey.ac.uk/11/3/1.html>.
- [22] Wheeler, M. and Clark, A. (1999) Genic representation: reconciling content and causal complexity. *British Journal for the Philosophy of Science*, 50, 1, pp. 103-135.
- [23] Zadrozny, W. (1997). A Pragmatic Approach to Context. *Context in Knowledge Representation and Natural Language*, AAAI Fall Symposium, November 1997, MIT, Cambridge.