Edmonds, B. (1998). On Modelling in Memetics.
Journal of Memetics - Evolutionary Models of Information Transmission, **2**.
http://cfpm.org/jom-emit/1998/vol2/edmonds_b.html

On Modelling in Memetics

Bruce Edmonds
Centre for Policy Modelling
Manchester Metropolitan University

Abstract

The field of memetics is characterised and two types of memetic model analysed: the *a priori* model and the `black-box' model. These are used to motivate a picture of scientific modelling, whereby chains of models are built from abstract models at the `top' down to data models derived from measurement. These models are linked to each other by *validation* (where the structure of a model comes from another) and *verification* (where the results of a model are checked against another). The role of formal models in clearly establishing these links are discussed. This picture is applied to the process of modelling in memetics. The paper argues that stronger selection criteria are required to enable the field to evolve. In particular it argues that for memetic models the minimum set of such criteria must include validation and verification.

Keywords: meme, modelling, philosophy, science, selection, explanation, constructivism, validation, verification

---

1 Introduction

Memetics is, at the moment, a gloriously diverse field - it ranges all the way from a narrative framework in fields such as history and anthropology to formal predictive theories in biology and computer science. It is, by and large, open to ideas and influences from outside of itself which is a trait to be appreciated and preserved. If this were not so, this would be a very different paper - I might be arguing for more variation and less selection.

In this paper I consider that part of memetics which aims to be a formal science. At this time its development in this direction has just begun. As in similarly young fields there is an initial proliferation of possible models and lots of informal `chat' about the proper designation of its subject matter. In this article I take a step back and apply the evolutionary metaphor to the field of memetics itself. I conclude that in order for the field to evolve as a science, better and more stringent selection criteria need to be applied. In particular I advocate the validation and verification of models with respect to other models, so that they form a complete explanatory chain from an abstrac tion of evolutionary and genetic processes to data from the domains of actual cultural communication.

After briefly discussing the content of the field of memetics and its development in section 2, I consider two types of memetic models that have arisen: the `a priori' (section 3) and the `black-box' model (section 4), with examples examined in each. Then in section 5, I argue that the most basic selection criteria should include verification and validation, as a move towards a minimum standard for explanatory models. In section 6, I consider further selection criteria that might aid the judgement of models in complex situations, using a hard-nosed constructivist stance.

---

2 The Subject Matter of Memetics and its Development

The subject matter of memetics has developed at least partly as a matter of historical accident. Although the etymology of the term `meme' goes back only to Dawkins in [6], the ideas can be traced back much further [11].

Instead of working historically I wish to characterise the field in a functional way, namely:

> the application of models with an evolutionary or genetic *structure* to the *domain* of (cultural) information transmission.

Similar to other developing fields there is a natural progression of questions. An initial question might be *whether such an application is at all useful*. Following on from that is the question of *in what ways is such application useful*. Finally a mature field should be concerned with the question of determining *under exactly what circumstances is such an application useful*. At the moment, although a positive answer to the first question is frequently assumed by academics in the field, the battle for its *general* acceptance is far from won. Work has only just began to touch on the second question.

At the moment there are many memetics models of different kinds. A glance through this journal alone shows many different models of what a meme is or could be. This is healthy, it means that the mechanisms of variation are truly operative, but for the field to *evolve* what we also need good selection criteria for such models. While older fields have had time for appropriate selection mechanisms to be themselves evolved, a new field such as memetics is in the position that there is a great deal of uncertainty about what might constitute a good memetic model.

In fields without selection criteria that are grounded outside the field (e.g. evidence, the theories of other sciences etc.), there is a strong tendency to adopt a sort of default set of selection criteria, namely: the extent to which the model uses established language and techniques from the field and the extent to which the model is embedded in the field itself, i.e. it builds on other models, criticises other models, examines the thought of other models etc.[*1].

Now, while such criteria are fine for maintaining the coherence of a field and for entrenching the position of those people already *in* the field, it is not a good set of criteria for ensuring the long term growth and success *of* the body of knowledge that makes up the field itself. Such criteria are widespread in a field like econ omics. Such long-term success will only come about when people outside the field see the memetic models as useful to them and this will only happen quickly if they feel they can rely on them. The reliability of the models will depend on the extent that the field is grounded in something established (real-world studies and/or accepted models from outside the field). In other words the success of the field (as opposed to the success of academics within the field), will depend largely on the *embedding* of the field within the wider academic landscape. It is notable that there appears to be a strong correlation between the success of a field and its openness to ideas from outside. A su sceptibility to Kuhnian paradigm shifts [14] seems to allow fields to adapt better to possible niches in the wider society[*2].

Of course any model selection criteria must take into account the type of model. There are almost as many types of model as there are reasons for modelling. Two dimensions along which they differ are in their abstractness and their temporal direction. One model may be for the prediction of events (forwards), and another for explanation of events that have already occurred (backwards), some will be metaphors for use in thinking about situations (abstract) and others will be bald data models of some aspect of the real world (concrete). In figure 1, I illustrate this categorization with a couple of models on these two axes.

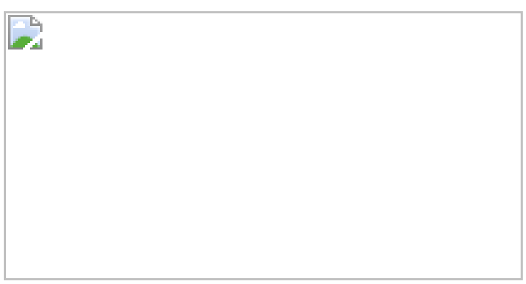

Figure 1. Some different types of memetic models, with their direction of explanation

---

3 *A Priori* Models

The first type of model I will discuss takes a foundationalist approach. That is, it examines the conditions under which an cultural evolutionary process may occur. This is the approach taken by Dawkins in [7], where he states the thesis that in the presence of a new replicator an evolutionary process is "inevitable". A recent example of this type of model is Calvin [4]. Such works attempt to establish that if certain facts are true then a memetic process will (generally) occur. They thus attempt to build forward explanations in a manner that in terms of theory is top-down (i.e. it starting with abstract considerations and concluding with the phenomena of evolutionary process, see figure 1).

One of the weaknesses of such types of models, *when considered in isolation from others*, is that the link to real-world (concrete) models is weak. Typically in such works the examples to support such models are of a very high level, involving such complex phenomena as *history* or *religion*. The complete demonstration that the abstractly motivated, *a priori* model would actually show that such complex phenomena are explained by memetic process can, at best, only be sketched and, at worst, is left solely to the imagination due to the fact that it is merely assumed.

This sort of assumption (that the complete explanatory chain *can* be built from such abstract considerations for a wide range of cultural phenomena) can be taken to ridiculous extremes, so that it may be taken *as read* that almost all cultural phenomena *are* due to memetic processes (I am *not* saying that Calvin isn ecessarily doing so). But the building of such a credible explanatory chain from theory to phenomena is not so easy - this (if and where it is possible) will be long and arduous. It is extremely tempting to attempt to short-cut the process.

One such temptation is to weaken the meaning of key terms in evolutionary models to *ensure* that the phenomena is a `memetic' one. One such weakening is to call any process which includes *some* variation and selection an evolutionary one. This has the effect of losing some of the meaning of memetic models - after all, if one *did* succeed in choosing ones terms such that one could *prove* that many processes were memetic then one would have also proved that such a theory had no empirical foundation; it would be a merely formal proposition. Memetics could become like economics, where certain assumptions, processes and goals are taken as *a priori*[*3] so that the gap between the theory and real-life becomes increasingly wide.

If memetics is to be a substantive theory of cultural processes (as opposed to just one of many equivalent descriptive *frameworks*), it has to be possible that some cultural process is not designated *by the theory* to be of an evolutionary nature. Otherwise we would be left with a non-falisfiable construct since its truth would be independent of the actual processes studied Calvin [19].

Let me propose one such process: *deductive inference*. Here, one or more pieces of information act up on each other to form a new piece of information such that the *status* of this new information is such that it derives from the status of the *old* information (and the process by which it is combined). This is an example of variation that is *not blind*; the variant is produced so that we have some assurance that it will be satisfactory *before* that selection process is actually applied to it. Such blindness is Heylighen's criteria for an evolutionary process [12]. The difference between an evolutionary process and, what I call an `intentional' process is illustrated below in figure 2 and figure 3.

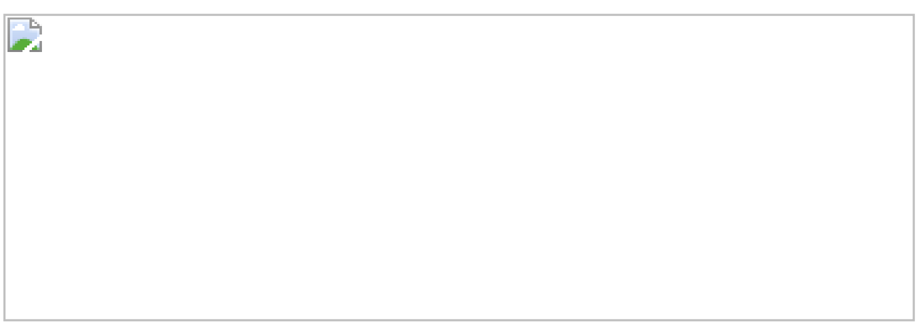

Figure 2. An evolutionary mechanism

Figure 3. An `intentional' mechanism

Thus, an evolutionary mechanism is one where the variations are produced without `consideration' as to whether they are likely to be selected, while an `intentional' one is where the variation is biased so as to produce variations that are likely to be selected. A consequence of this is that in an evolutionary process it is the process of selection that is critical for the process (since the process of variation can of many types as long as it produces sufficient variability and is blind to the selection), while in an intentional process it is frequently the variation process that is critical. A memetic model is an evolutionary model whose subject is the (cultural) transmission of information between peers.

Since I take a constructivist approach to knowledge, I see the distinction between `evolutionary' and `intentional' mechanisms as being primarily a property of the *model* of the process concerned and only by extension (via the success or otherwise of the model) to the process being modelled. Of course, many of the thought processes of humans can be modelled as either evolutionary or `intentional' processes (or, indeed, a mixed of both). I would argue that it would be clearer to say that such processes have memetic *aspects* (to a greater or lesser degree), rather than to simply class them as memetic, since otherwise we are in danger of losing some of the meaning of the term.

Deductive inference is an extreme example of an `intentional' process - in deduction the operators of variation (for example *modus ponens*) are such that they *only* produce the desired results, i.e. logical consequences. Thus such inference (if and where it occurs) is a good candidate for a process that may act upon information but is *not* a memetic one. Of course, many examples that are sensibly categorised as inferential may be seen as a result of an evolutionary processes when the best model of the mechanisms that implement that top-level process have an evolutionary structure. However, the presence of such evolutionary sub-processes do not weaken the claim that top-level process, of which it these are but constituents, is essentially *not* an evolutionary one. The fact that evolutionary processes can implement inferential ones and vice versa does not mean there is not a meaningful and important distinction between them, just as the fact that discrete and continuous mechanisms can implement each other does not mean that there is no difference between them*4.

Unless we want to follow the sterile route of *a priori* economics then we have to ensure that the explanatory chain from abstract model to real-world relevance is really built rather than merely assumed. The *a priori* models should be allowed to survive only if such explanatory chains establishing their relevance are established.

---

4 `Black-Box' Models

The other end of the spectrum of is the `black-box' model. Here some data from a phenomenon is analysed and a memetic model is fitted to it. So one can say that it is *as if* there were a memetic process at work. A good example which illustrates some of the aspects of this approach is Best's [3]. Here the phenomena of the repetition of clusters of words in news groups is analysed and evidence is presented to show that some of these clusters are acting as if they are competing with each other as quasi-species. It is unfortunate (for us readers) that Best did not then go into a detailed case study to see the actual mechanisms that caused this apparent anti-correlation, as a biologist might when faced with an apparent predator-prey cycle in some population statistics*5. Instead in this paper he relies more on reference to middle-range concepts such as the scarcity of resources. Thus his paper can be summarised as: an analysis of data *indicates* that an evolutionary process is taking place.

Of course the weakness of this type of model, *if it is taken in isolation*, is that one might be mistaken. A process may *appear* to be evolutionary but later turn out to be otherwise. If one took the step of defining evolutionary process by its (high-level) phenomena, then one would no longer have a theory but just a description.

To illustrate this case I present a computational model, consisting of a collection of agents which interact and communicate with each other as they learn about what is best to do and say. In this simulation it appears that a memetic process is taking place (if one just considers the output in terms of who said what to whom), but it turns out (if one analyses the internal processes that produce this behaviour) that these results are actually caused by agents coming simultaneously to the same conclusions about what to say. This simulation is analysed elsewhere in detail for some of its other properties in [9, 10]. It requires some detail of the model set-up to be understood for the relevance of the example to be become clear, but I have tried to keep it as simple as possible by relegating much of the technicalities to an appendix (section 7). This model can be seen as an extension of the work in [1], which investigates a three player game.

4.1 The Extended `El Farol Bar' Model

This model is based upon Brian Arthur's `El Farol Bar' model [2], but extended in several respects, principally by introducing learning and communication. There is a fixed population of agents (in this case 10). Each week each agent has to decide whether or not to go to the b ar. Generally, it is advantageous for an agent to go unless it is too crowded, which it is if 67% or more of all the agents go (in this case 7 or more). It is in each agent's interest to go to the bar when most of the others do *not* go and stay at home when they do. Each agent benefits from guessing what (most of) the others will do and th en doing the opposite. Since all the agents are trying to out-guess the others, a sort of modelling `arms-race' develops.

Each week, before making their decision, agents have a chance to communicate with each other. They indicate to each other whether they will go or not, but they do not have to tell the truth. The messages they send can be a simple `yes' or `no' or a more complicated expression.

4.1.1 The environment

Each agent gets the most utility if it goes to the bar when it is not too crowded (i.e. less than 7 of the other agents go), gets a medium utility if they stay at home and the lowest utility if they go when it is crowded. In this way there is no fixed reward for any particular action because the utility gained from going depends on whether too many other agents also go. In this way there is no fixed goal for the agent's learning, but only one which is relative to the other agent's behaviour.

4.1.2 The agents

Each agent learns its own set of strategies of what to do and say in a flexible and open-ended way. It can base its strategy of what to *say* upon what other agents did or said in previous weeks as well as on trends in past attendance statistics, random inputs and what it said or did itself in precious weeks. It c an base its strategy of what to *do* on what the other agents just said as well as what it just said itself, what it did the previous week and a random decision. In either case it can build up mixtures of strategies using logical or numerical operators, without limit as to their complexity. It so happens that the learning mechanism is *i mplemented* by an evolutionary mechanism but this is entirely internal to each agent. I have no reason to suppose that the learning mechanism is critical to the results of this model (as long as it is open-ended, creative and context-sensitive).

An example strategy that an agent learned is shown in figure 4. This would cause the agent to say "`[IPredictedLastWeek]`" if the attendance predicted by the trend in attendances over the last 2 weeks was greater than 8/3 and `false` otherwise, but it would only actually go if either it said it would or if `barGoer-3` said it would.

Figure 4. An example model

A second example (figure 5) shows the operation of a memetic process. In this example the agent says whatever the expression last uttered by `barGoer-5` evaluates to, but goes if it went last time and a random coin-flip turns up heads. In this case this agent would propagate what ever `barGoer-5` said.

Figure 5. A second example model

A full indication of the sort of building-blocks that the agents can use to make these models is illustrated in figure 11.

4.1.3 Communication

Each agent can communicate with any of the others once a week, immediately before they all decide whether to go to the bar or not. A section of the transcript from the model is shown in figure 6, which shows the messages sent in one particular week (early in the simulation).

Figure 6. Part of the model transcript showing some messages (week 14 of the run)

4.1.4 Results, analysis and discussion

If one looks at the occurrence of different types of utterance by the agent and graphs the frequency of different message components it looks distinctly as if a memetic process is taking place. Consider, for example the occurrence of "`IPredictedLastWeek`" over the course of the first 40 week of the simulation (figure 7).

Figure 7. Number of occurrences of "`IPredictedLastWeek`" in agent utterances

The occurrences occur in distinct patterns - connected sequences which rise chaotically to a peak over several weeks but which then peters out. Futhermore it not generated by only one (or a a few) agents, but a substantial proportion of the agents participate, one after the other. The distribution of who says "`IPredictedLastWeek`", when is shown below in figure 8 (`barGoer-3` has not been included n this figure since it never utters this particular message part).

Figure 8. The distribution of utterances of "`IPredictedLastWeek`" by agent

If one kept one's analysis at the level of the emergent communication it would appear that an utterance is started by one agent which is then spread by the rest of the population until it losses the ability to impart comparative advantage on its speakers and so is selected out. This sort of emergent pattern of communication occurred with a variety of different messages at different stages in the simulation. Add to these observations the knowledge that are mechanisms built into the agents that would allow the selective propagation of utterances from agent to agent, so that strategies that supported a memetic process between agents could evolve and one might feel quite safe in concluding it is a memetic process.

However, if one analyses the actual strategies that the agents usedby "looking inside the agent's head", then one gets a different picture. The similar utterances that were observed were not generally a result of imitative strategies learned by the agents (as in figure 5). In figure 9, I exhibit a similar graph to that of figure 7, but separated into those utterances of "`IPredictedLastWeek`" that were as a result of imitation by an agent and those that are the result of other mechanisms. I compiled this by analysing the actual strategies for communication learned by the agents at each date where these utterances occurred to see how they were derived. These other utterances arose because they were *simultaneously* determined as relevant by the learning of each agent. What memetic imitation there was is attributable to chance copying - a sort of weak `echo' of other utterances. These utterances with `` `IPredictedLastWeek' in them `` were not *selected for* by these agents. In particular one does not find any branching chains supporting the presence of memetic copying of messages. In fact there is only *one instance* of such a message being passed down the generations for more than one notional week (this occurs in the middle of the simulation, `barGoer-3` at week 16 is copied by `barGoer-2` at week 17 which is copied by `barGoer-7` at week 18 and then by `barGoer-8` at week 19), and even then there is no *replication* of the meme into more than one strand.

Figure 9. Number of occurrences of "`IPredictedLastWeek`" in agent utterances analysed by generating process

The fact that the learning processes in this simulation are achieved via an evolutionary mechanisms is beside the point. It does not make the process of messages between agents into a memetic one. A memetic process could have co-evolved if the agents evolved models to do so, but this did not happen here. The agents did sometimes copy from each other but not so as these copying processes where chained into a memetic process so that we could meaningfully say the messages themselves underwent an evolutionary process.

If this sort of pseudo-memetic process can occur in an artificial simulation where mechanisms are explicitly built in a mechanism for the expression, propagation and selection of memes by agents, then it is possible that similar situation could occur elsewhere. Thus it is wrong to imply the existence of memetic process, purely as a result of a post-hoc analysis of data and the mere presence of possible mechanisms. Such `black-box' models are insufficient on their own.

---

5 Basic Selection Criteria for Memetic Models

It should be clear that what is needed is for such models to be chained into a complete explanation so that the presence of evolutionary mechanisms can be shown to *cause* the evolutionary phenomena that results. This would then justify the models that composed that chain. Note that th is chain can (and probably will) be composed of more than one model. Even if one did have one model which seems to bridge the gap there would always be, at the top, the metaphor of evolution (that I talked about in section 1) that it must relate to in order to justify its categorisation as a memetic model and the data model at the bottom.

It would be silly to insist that every paper must, by itself, bridge this gap, but what each model should (at a minimum) do is say how it explicitly relates to the models above and below itself in the chain. In other words, each model should be well *verified* and *validated*, as Scott Moss and I suggested for models in a different context [17]. The validation of a model is where the structure of the model is based on the *results* of another model (even if it is fairly abstract like the idea of evolution). The verification of a model is where the results of the model are checked against something more definite (either the structure of other models that are themselves verified or a data model).

To make this clear let me take an example from a more established science: chemistry. Physics has models of the processes of elementary particle interaction[*6]; chemistry builds on this by positing mechanisms of how chemicals interact that use these physical processes (i.e validated) and verify these models against data models of real processes. Sometimes they then use these basic models of how chemicals interact in simulations which involve many chemicals when analytically solving to predict their interaction is too difficult. They may then make approximate numerical models of the results of these simulation and verify these against data models of real processes. The whole chain is illustrated in figure 10.

Figure 10. Validation and verification between layers of models in chemistry

Of course, beyond the mere presence of the validation and verification links, the *strength* of those links is also important. One method for helping to establish strong links is the use of formal models. Even if these formal models are not used to infer anything new they can be immensely useful in pinning down the meaning of terms and so reduce the potential for ambiguity and confusion. It is the bane of the social sciences, that everyone has slightly different meanings for key terms and concepts, and memetics is obviously no exception to this! A formal model allows the clarification of these meanings, so that the exact relationship between different models can be established. Another way of looking at this is that the links are themselves a sort of model, that are verified against the structure of the model they link: formal models thus allow the validation of these links against formal processes.

My characterisation of memetics in section 2, allows the categorisation a model as *memetic*, by the ultimate validation of its structure of against the evolutionary/genetic metaphor and the its ultimate verification against data from a process of cultural communication. These processes can be indirect using ot her intermediate models, as discussed above.

I wish to make clear that I am *not* criticising Calvin, Best or their work. A field will and must have specialists in different sorts of model. However the utility of their work *will ultimately depend* upon whether their models are included in a complete explanatory chain from theory to phenomena. Since memetics has only just began as a serious enterprise it is unsurprising that there is a lack of other convenient intermediate models with which to build these links, but this does not relieve us of the obligation to do so - a situation like economics where the excuse "its still a young science" is used 300 years after its inception should not be allowed. All models that are not part of a full explanatory chain should be mentally tagged as *provisional* and (at least temporarily) discarded if it is not justified in the *near* future.

This is, of course, a sort of bootstrapping process. In order to progress into a "normal science" (in Kuhn's terminology [14]) then we need to establish *some* complete explanatory chains as quickly as possible, so that more elaborate chains may be built with reference to these. To this end it would help if researchers in the field concentrated on domains where both the ultimate validation against the evolutionary metaphor and the ultimate verification against evidence of memes was as straight-forward as possible. Thus (at least until the shape and limitations of memetics becomes clearer) it would be far more helpful for studies in domains such as: bird song, nursery rhymes, tunes, and legal phrases in contracts where the mapping between the form and meaning of potential memes is clearly evident. Domains which involve very high level social constructs such as *history* or *religion* should be generally avoided since in these cases the construction of a convincing explanatory chain is the most onerous, and hence most suspect if used as support for a model.

---

6 Further Methods to Ensure Rigorous Selection

Validation and verification are the basic requirements for a model. In themselves they are only sufficient if one is modelling relatively simple processes, where it is possible to effectively and clearly test models by carefully designed experiment. In modelling cultural transmission, we are not (on the whole) so lucky. In approaching our subject matter we are forced to make so many choices regarding the selection of data, the modelling framework, modelling goals, etc. that there is a danger of these two selection criteria being insufficient to ensure the continued evolution of our models. It is well known that in evolutionary processes if the mutation rate is too high compared to the rest of the process then this can swamp the process so that it does not learn to successfully adapt; this may occur to the field of memetics - if we let it.

Hence for memetics (and other similar fields), the difficulties of verification mean that we need more stringent selection criteria. We may have to accept that we are *constructing* our memetic reality rather than merely passively *reflecting* it by our modelling choices, since the complexity of our chosen phenomena forces this on us . But just because we take such a constructivist stance[*7] we do not have to be vague - we can respond to our chosen challenge by seeking to retain rigour.

These other criteria can be such as: its similarity to other accepted models (a lateral coherence criteria); the complexity of the model; the meaningfulness of the model in human terms; the ease with which the model can be used to calculate its predictions; and how general (or generalisable) it is.

Let me stress again that these are not optional extras - by applying these and making them explicit in our work we are only accepting their inevitability. When dealing with the complexity of phenomena such as social interaction we have no choice but to accept trade-offs in terms of the complexity of our models and (for example) their error-rate when verified against data [8].

---

7 Conclusion

The complexity of social interaction and the difficulty in directly falsifying memetic models means that unless we apply strong selection criteria the field of memetics will be swamped by the forces of variation and hence not successfully evolve. As a result of this memetics as a *field* may well lose the race to others in academia and society. For this reason the application of stringent selection processes (that are relevant to the survival of the field as opposed to relevant to only the survival of academics and their constructs) is vital.

This can be done in a number of ways:

1. (Use of formal models) insisting on formal or computational models to pin down the meaning of our theories. This makes the relationships between our theories easier to assess since the formal models will help specify the reference of terms in an unambiguous way;
2. (Validation) make clear the assumptions behind our models in terms of the modelling language, validating theories, modelling goals etc.;
3. (Verification) explicitly state how we are verifying our model, including: how we have selected the model or information to verify against and the precise method of this verification;
4. (Other criteria) in addition to the above, argue for our models in terms of additional criteria such as: simplicity; ease of computation of predictions; etc.
5. (Domains) initially choosing domains where the meaning of our models are clear - i.e. a concentration of phenomena where the identity of memes is transparent due to the clear mapping between their forms and their content (e.g. birdsong, tunes, legal phrases in contracts, nursery rhymes) and treating examples concerning high-level constructs (e .g. religion, history, organisational behaviour) with great care.

Appendix - Technical details of the extended `El Farol Bar' model outlined in section 4

8.1 The General Set-up

There are 10 agents. Each notional week (out of a total of 100): the agents construct their new set of models, and evaluated them against the average past utility they would have resulted in if they had been used during the last 10 time periods (and all the other agents had acted as they actually did in the past); they choose the best such model; using this model they decide what to say (possibly referencing what others did and said in previous weeks); then they decide what to do (possibly referencing what the agents have just said as well as past actions etc.). The actions are then recorded.

8.2 The Agent's Learning Algorithm

Each agent has a population of (pairs of) expressions (as in figure 4 and figure 5) that represent possible behaviours in terms of what to say and what to do. The size of the population is fixed but not content. These expressions are taken from a strongly typed formal language which is specified by the programmer, but the expression can be of any structure and depth. Each agent does not `know' the meaning or utility of any expression, it can only evaluate each expression as to the utility that would have resulted in if it had used it in the past to determine its action. Each week each agent takes the best such pair of expressions (in terms of its present evaluation against the recent past history) and uses them to determine its communication and action.

Each agent has a fairly small population of such models (in this case 30). This population of expressions is generated initially according to the specified language at random. In subsequent generations the population of expressions is developed by a genetic programming [12] algorithm with a lot of propagation and only a little cross-over.

The formal language, that these expressions are examples of, is quite expressive. The primitive nodes and terminals allowed are shown in figure 11. It includes: logical operators, arithmetic, stochastic elements, self-referential operations, listening operations, elements to copy the action of others, statistical summaries of past numbers attending, operations for looking back in time, comparisons and the quote operator.

Figure 11. Possible nodes and terminals of the tree-structured genes

The primitives are typed (boolean, name or number) so that the algorithm is strictly a strongly-typed genetic program following [15].

8.3 Communication

The communication is determined by the evaluation of the talk expression and is usually either `true' or `false'. The presence of a quoting operator (`quote`) in the formal language of the talk expression allows subtrees of the talk expression to be the content of the message. If a quote node is reached in the evaluation of the talk expression then the contents of the subtree are passed down verbatim rather than evaluated. If a quoted tree is returned as the result of an evaluation of the talk expression then this is the message that is communicated.

The content of the messages can be used by agents via `saidBy` and `saidByLast` nodes in the action and talk expressions. If the message is just composed of a boolean value then the `saidBy` node is just evaluated as this value, but if it is a more complex expression (as a result of a `quote` node in the sending agents talk expression) then the whole expression will be substituted instead of the `saidBy` (or `saidByLast`) node and evaluated as such. The agent can use the output of its own messages by use of other nodes (`IPredictedLastWeek` and `ISaidYesterDay`).

The agents are limited in those from whom they can `listen to', in the sense that at the beginning of the run an arbitrary `friendship structure' (as in) is imposed and agents can only use `saidBy` (or `saidByLast`) nodes referring to agents who are their friends in their models.

In this model `social imitation' is also enabled. this means that other agents can introduce the content of any message (which is not a mere boolean value) into their own (action) model pool, this would correspond to agents taking the message as a suggestion for an expression to determine their own action. In subsequent generation this expression can be crossed with other expressions in its population of constructs so that its content is mixed in with other *action* models of that agent. Note that social imitation does not incorporate the content of such messages into the models that determine what is said by each agent.

8.4 Memetic mechanisms

The expression that controls what the agent says can be of any form allowed in the designated formal language. In determining the causes of utterances in order to compile figure 9, I designated as memetic, any expression whose evaluation could possibly be effected by the utterance of another agent. For example if there were no `saidByLast' nodes in the chosen talk expression, one could safely say it was not memetic.

8.5 Implementation

The model was implemented in a language called SDML (strictly declarative modelling language), which has been developed at the Centre for Policy Modelling specifically for social modelling [18].

Acknowledgements

My thanks to Mark Mills, for his kind comments on a previous paper [9], which provided the impetus for this paper, to Michael Best for his reactions, to the referees whose incisive and insightful comments I really enjoyed reading, to my colleague, Scott Moss, for many discussions about the process of modelling, and, of course, to Steve Wallis for writing SDML.

SDML has been developed in VisualWorks 2.5.1, the Smalltalk-80 environment produced by ObjectShare (formerly ParcPlace-Digitalk). Free distribution of SDML for use in academic research is made possible by the sponsorship of ObjectShare (UK) Ltd. The research re ported here was funded by the Economic and Social Research Council of the United Kingdom under contract number R000236179 and by the Faculty of Management and Business, Manchester Metropolitan University.

Notes

1. For a more formal definition of such embedding see [9].

2. Readers who wonder at my focus on the success of the field, rather than knowledge itself will need to take into account my constructionist and evolutionary stance, because I think that in order to advance knowledge on needs something like an evolutionary process acting upon a population of hypotheses.

3. Equivalently the field of economics can be defined by the study of such situations where they are true so that the field is no longer designated by the relevant subject matter itself.

4. In fact it is very difficult to efficiently implement deductive inference by a purely evolutionary mechanism.

5. Although Michael Best has indicated to me that he is in the process of doing this.

6. Some of the models in physics are of a slightly different character, because they form the ultimate top-level model for many such chains. Thus, despite the fact that these models are often taken as the archetypal model for science they are, in fact, quite special cases.

7. For a good introduction to constructivism and the importance of selection criteria see the pages at the Principia Cybernetica Project at URL: http://pespmc1.vub.ac.be/construc.html

---

References

[1] Akiyama, E. and K. Kaneko, (1996). Evolution of Cooperation, Differentiation, Complexity, and Diversity in an Iterated Three-person Game, *Artificial Life*, **2**:293-304.

[2] Arthur, B. (1994). Inductive Reasoning and Bounded Rationality. *American Economic Association Papers,* **84**: 406-411.

[3] Best, M., L. (1997); Models for Interacting Populations of Memes: Competition and Niche Behavior. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://cfpm.org/jom-emit/1997/vol1/best_ml.html

[4] Calvin, W.H., 1997; The Six Essentials? Minimal Requirements for the Darwinian Bootstrapping of Quality. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://cfpm.org/jom-emit/1997/vol1/calvin_wh.html

[5] Campbell, D. T. (1960). Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes. *The Psychological Review*, **67**: 380--300.

[6] Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.

[7] Dawkins, R. (1982). *The extended phenotype*. Oxford University Press.

[8] Edmonds, B. (1997). Complexity and Scientific Modelling. 20th International Wittgenstein Symposium, Kirchberg am Wechsel, Austria, August 1997. Also available as a CPM Report 97-23, MMU, Manchester, UK. http://cfpm.org/cpmrep23.html

[9] Edmonds. B. (1998). Capturing Social Embeddedness: a constructivist approach. CPM Report 98-34, MMU, Manchester, UK. http://cfpm.org/cpmrep34.html

[10] Edmonds, B. (forthcoming in 1998). Modelling Socially Intelligent Agents. *Applied Artificial Intelligence*, 12(7). http://cfpm.org/cpmrep26.html

[11] Gatherer, D. (1997). Macromemetics: Towards a Framework for the Re-unification of Philosophy. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://cfpm.org/jom-emit/1997/vol1/gatherer_dg.html

[12] Heylighen F. (1992): Principles of Systems and Cybernetics: an evolutionary perspective, in *Cybernetics and Systems '92*, R. Trappl (ed.), Singapore: World Science, 3-10. ftp://ftp.vub.ac.be/pub/projects/Principia_Cybernetica/Papers_Heylighen/Systems_Principles.txt

[13] Koza, J.R. (1992). *Genetic Programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.

[14] Kuhn, T.S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

[15] Montana, D. J. (1995). Strongly Typed Genetic Programming, *Evolutionary Computation*, **3**:199-230.

[16] Moss, S.J. and Edmonds, B. (1998). Modelling Economic Learning as Modelling. *Cybernetics and Systems*, **29**:5-37. http://cfpm.org/cpmrep03.html

[17] Moss, S. J., Edmonds, B. and Wallis, S. (1997). Validation and Verification of Computational Models with Multiple Cognitive Agents. CPM Report 97-25, MMU, 1997. http://cfpm.org/cpmrep25.html

[18] Moss, S. J., H. Gaylard, S. Wallis, and B. Edmonds, (1998). SDML: A Multi-Agent Language for

Organizational Modelling. *Computational and Mathematical Organization Theory*, **4**:43-69.
http://cfpm.org/cpmrep16.html

[19] Popper, K. R. (1968). *The Logic of Scientific Discover*y. London: Hutchinson.

Back to Issue 2 Volume 2