

PHASE CHANGES IN META-MODELLING USING THE FRACTAL DIMENSION

James MOFFAT and Susan WITTY

Introduction

In this paper we are concerned with intelligent agent simulation models which can be used to investigate some of the new aspects of “information-age” conflict. Such models consist of a number of agents which interact locally in order to produce global emergent behaviour. One of the simplest, yet very relevant, examples of such a model is ISAAC, developed by the US Marine Corps Combat Development Centre under the “Project Albert” initiative.^{1,2} In this case, modification to the behavioural characteristics of the agents produces significantly different emergent behaviour in terms of the flow of battle and casualties suffered. Some of these behaviours can be surprising, and it is the aim of our work to produce a theory of such processes which helps to explain the types of behaviour to be expected.

In short, understanding the behaviour of such agent based combat models is now becoming more important, especially as the agents gain intelligence and try to outsmart each other, producing potentially very complex behaviour. The principal variables in these models can often be separated out from the rest of the model to produce a meta-model which is aimed at decreasing the run-time of the original model while still retaining the characteristics and arriving at the same final solution as the original model.

As an example, in developing a meta-model we consider the relationship between a key outcome of the model, a , and a set of input variables as follows:

$$a = f(a_1, \dots, a_k, b_1) \quad (1)$$

(This is easily generalised to an arbitrary number of b 's).

The meta-model function can then be shown to have the property: ³

$$a = f(a_1, \dots, a_k, b_1) = a_1^p \dots a_k^r \Phi \left(\frac{b_1}{a_1^{p_1} \dots a_k^{r_1}} \right) \quad (2)$$

In some cases the function Φ can be further simplified using a re-normalisation group.⁴ This approach directs us to search for evidence of power law relationships, re-normalisation groups, and normalised distributions of the form Φ . Such expressions arise naturally in certain types of complex system, particularly where fractal structures are involved, and are referred to as “scaling” relationships, since they have no preferred gauge. Moffat and Passman show that there is clear evidence for such assumptions.⁵ In particular it turns out that the attrition rate for one side in such an agent-based model is correlated to the fractal dimension of the opponent.^{6,7} This fractal dimension is a measure of the ability of a side to cluster and collaborate locally (to produce locally advantageous force ratios).

As an example of how these apparently abstract ideas can give us immediate insight, we illustrate the problem of relating local force clustering and collaboration to the ability of the whole force to control an Area of Operations. Control is defined here in terms of the ability to prevent the opposing force or some third party from being able to move freely across the Area of Operations.

Let a single unit be able to control a patch of the Area of Operations (AO) corresponding to a square of side length l . We assume that the force (from our previous discussions) is clustered fractally with fractal dimension D . It follows from the definition of D that if we cut up the AO into squares of side length l , then the number of occupied squares is

$$N(l) = l^{-D} \quad (3)$$

Thus the probability of a square being under control is given by

$$p = \frac{l^{2-D}}{A}, \quad (4)$$

where A is the area of the AO. This is in fact a standard measure of fractal clustering, and hence is a good measure of the ability of the force to locally cluster and collaborate.⁸

We now apply a re-normalisation group approach. By considering configurations in which 1,2, 3 or 4 of the cells of side l are controlled sub-regions of the square of side $2l$, and assuming that control requires a span of controlled cells stretching from side to side and top to bottom (to prevent flow through the region in any direction), we have the following relationship:

$$p(2l) = 4p^3(1-p) + p^4, \quad (5)$$

where $p(2l)$ is the probability of control of a square of side $2l$. By continually increasing the area considered in this way, through re-normalisation, we have the general iterative relationship for the probability of control at increasing levels of span of the AO:

$$\begin{aligned} p_{n+1} &= 4p_n^3(1-p_n) + p_n^4 \\ &= p_n^3(4-3p_n) \end{aligned} \quad (6)$$

The stable points in this recursive relation then correspond to the intersection between the function

$$g(x) = x^3(4-3x) \quad (7)$$

and the line $y=x$ for values of x in the region between 0 and 1. In this region, $g(x)$ is S-shaped and it thus has three intersections with the line $y=x$. Two of these are stable, at the extreme values 0 and 1, and there is one unstable intermediate point.

Interacting force units should thus polarise to either a very high level of control or a very low level of control. Any particular clustering of the force (corresponding to a particular unit probability of control p) should thus give rise to a level of control which is attracted (by re-normalisation to larger areas of the AO) towards one of the two extremes.

Top Down Planning

All of the above analysis indicates that Fractal Dimension is a key indicator of behaviour in agent based models of conflict. However, it is very much a “bottom up” generated emergent behaviour of the model based on inherent assumptions about local clustering and collaboration. In order to complete the description of a meta-model, we also have to represent the “top down” perspective of a high level plan, and how this interacts with such emergent behaviour generated from below. To do this we have adopted the structure of Bayesian Decision Making, which assumes there is some overall goal or objective (or more generally a set of such objectives) to which the system as a whole evolves, and there is a loss of utility involved in not meeting such objectives. In agent based modelling terms we consider then a set of such desired end states or “attractors” to which the model evolves. We have an understanding of the set of outturns from the simulation model, and an understanding of the loss of utility relating to the gap between each outturn and the set of goals. We then consider the Expected Utility function, which combines together the distribution of outturn, and the associated utility. It turns out that the mathematics of such functions is congruent with that of elementary catastrophe surfaces, and this is discussed below.

In the first of these investigations we look at the robustness of each start point in relation to the end points of the model run. This is in the context of many possible different end states or goals. The second looks at the problem with only a small number of end states or goals.

Robustness Analysis with Many Possible Goals

In a given system, it is possible to divide the parameter space into regions of probability of each end state. In this case the parameter is the fractal dimension of one of the forces at the beginning of a model run. We assume that the probability of any start point of fractal dimension resulting in a particular end state can be approximated by a scaled Normal distribution with a mean value μ_i and variance V_i locally. The probability can be written as:

$$f(D|i) = \frac{h}{\sqrt{2\pi V_i}} \exp\left[-\frac{(D - \mu_i)^2}{2V_i}\right] \quad (8)$$

where, for a two-dimensional problem, $0 \leq D \leq 2$.

However, for any start point D , there is not just a single attractor which the system could move to as its end state. At each fractal dimension there exist j different

attractors which the system could possibly choose as its end state. At this chosen fractal dimension D , each of these end states has a probability $p_j = f(D|j)$. At any given fractal dimension, the sum of the probabilities of reaching all end states must be 1, i.e. $\sum_j f(D|j) = 1$.

Each of the possible end states has a utility attached to it. In this way, some end states or attractors are considered to be “better” outcomes than others. For each fractal dimension value D , with respect to a single end state i , a utility function $U(i, D)$ is assumed, according to the utility of each end state based on the fractal dimension.

To calculate the expected utility for each end state over all possible values of the fractal dimension D , the probability function and the utility function are multiplied together and integrated over all possible values of the fractal dimension:

$$E(i) = \int_{\forall D} f(D|i)U(i, D)dD \quad (9)$$

This indicates the expected utility of end state i . It is a measure of how robust each of the end states is over the range of all possible start points. Alternatively the utility of the point D as a start point can be calculated as a sum over all possible end states:

$$E(D) = \sum_i f(D|i)U(i, D) \quad (10)$$

This is a measure of which fractal dimension to start from to attain the maximum expected utility of outcome.

Catastrophe Theory Approach

In the above discussion, there was no explicit decision to make. The decision was implicitly, “which start point do I choose based on the prior knowledge I have of the model?” or “what is the best end state to aim for based on my prior beliefs?” based on maximum expected utility. By making this decision explicit, it is possible to utilise Bayesian decision theory and the resultant mathematics of catastrophe theory to help in making the decision of which start point to choose for the model. The question we shall ask is: “I want my end state to be attractor i . Which is the best fractal dimension to choose for the start point of the force?”

A probability distribution for the prior belief of the best possible starting values of the fractal dimension can be written down to achieve the end state i .⁹

$$f(D|i) = \frac{h}{\sqrt{2\pi V_i}} \exp\left[\frac{-(D - \mu_i)^2}{2V_i}\right], \quad (11)$$

where the mean of the distribution is μ_i and the variance V_i . The next question we ask is “what would happen if we move the fractal dimension by a mean amount δ in an attempt to improve the likelihood of achieving end state i ?” We assume that the probability of each fractal dimension giving the end state i is now given by

$$f(D|i, \delta) = \frac{h}{\sqrt{2\pi V_i(\delta)}} \exp\left[\frac{-(D - (\mu_i + \delta))^2}{2V_i(\delta)}\right] \quad (12)$$

By changing the starting fractal dimension of the force, we are assuming that the change to the start point will increase the success in outcome in relation to end state i . However, if the starting point is changed from the mean μ_i to the value $\delta + \mu_i$, the benefit of that starting point in terms of achieving goal or end state i is less certain and the variance must increase. Therefore, the variance $V_i(\delta)$ depends on the decision δ made (we assume) according to:

$$V_i(\delta) = (V_i + A\delta^2)^{1/2} \quad (13)$$

A is a scaling factor which relates how much effect the decision δ has on the variance of the distribution. A high value of A means that the user has very little information on what the effect of changing the start point will have on the outcome of the combat model other than the information he has on the original belief of the start point; thus a change will make a large increase to the variance.

It is assumed that there exists a preferred fractal dimension denoted by D_0 . This may be that determined by the doctrine of the force. Alternatively, it may indicate that there is high utility in the end state being achieved relative to some other measure, for example that the goal should be achieved with a high level of control of the battlespace. Therefore, the fractal dimension of D_0 should be the “best” fractal

dimension for any desired outcome or end state and the utility of any starting point can be written as:

$$U(i, D) = U(D) = \exp\left[\frac{-(D - D_0)^2}{2k(\delta)}\right] \quad \forall i \quad (14)$$

Just as the variance of the prior belief function changes with greater decisions, the variance $k(\delta)$ of the utility function can too. As the change δ to the starting value of the fractal dimension is increased, it becomes more important that the starting value of the fractal dimension D is closer to the value D_0 , since it is uncertain what effect other values of D will have on the outcome or end state. So k can be written as

$$k(\delta) = (\eta + B\delta^2)^{-1/2} \quad (15)$$

Here, η is the largest possible variance the utility function may take. B is a scaling factor which represents the effect of a change to the mean fractal dimension value μ_i .

If $U(D)$ and $f(D|i, \delta)$ are multiplied together and integrated over all possible changes to the prior belief, the maximum value corresponds to that decision which will give the highest expected utility to that force given the aim of achieving end state i .

$$E_i(\delta) = \int_{\forall D} U(D) f(D|i, \delta) = h \left(\frac{1}{1 + (\eta + B\delta^2)^{1/2} (V_i + A\delta^2)^{1/2}} \right)^{1/2} \exp \left\{ \frac{-(\delta + \mu_i - D_0)^2}{2 \left(\frac{1}{(\eta + B\delta^2)^{1/2}} + (V_i + A\delta^2)^{1/2} \right)} \right\} \quad (16)$$

To find the points of maximal utility, this function must be differentiated with respect to the decision δ and the points at which $E_i'(\delta) = 0$ are the points of maximum (and minimum) expected utility.

To examine these turning points, we thus plot the function $E_i(\delta)$, the expected utility of the start point given the aim of achieving end state i . For different values of μ , η , V , A and B , different plots are obtained. The following three figures (1, 2 and 3) show the behaviour and the best start point of fractal dimension for increasing values of the parameter A .

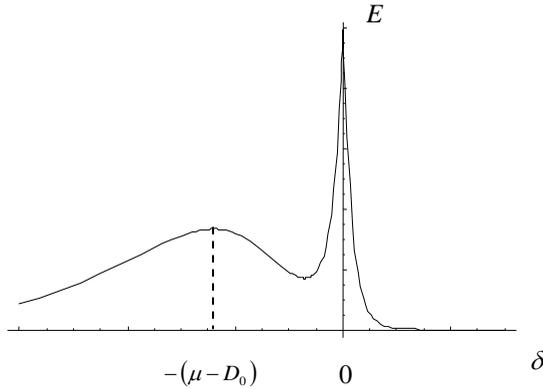


Figure 1: Behaviour of the expected utility for increasing values of A .

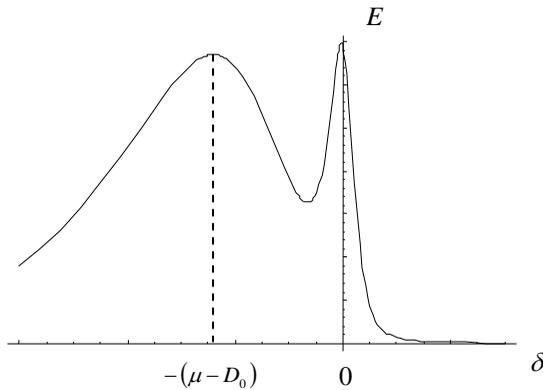


Figure 2: Behaviour of the expected utility for increasing values of A .

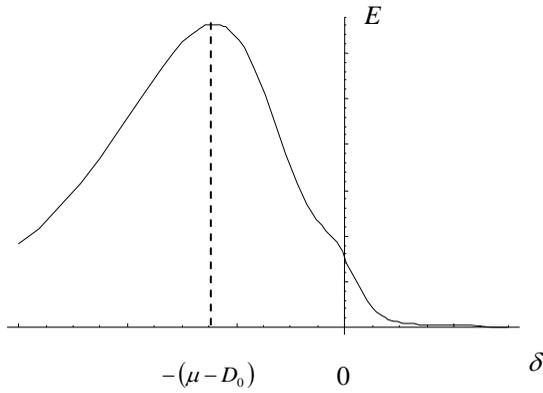


Figure 3: Behaviour of the expected utility for increasing values of A .

The above three figures show that the change δ with the highest utility switches suddenly from 0 to approximately $-(\mu - D_0)$ as the A , the dependence of the variance of the prior belief on the decision change, increases. A value of $\delta = 0$ means the expected starting value of fractal dimension should be μ_i . If $\delta = -(\mu_i - D_0)$, this means that the most appropriate expected initial fractal dimension value is that of D_0 . Similar behaviour can be seen when the value of the parameter μ_i is increased.

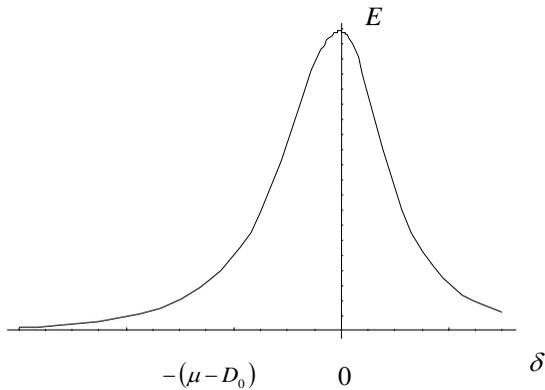


Figure 4: Behaviour of the expected utility for increasing values of μ_i .

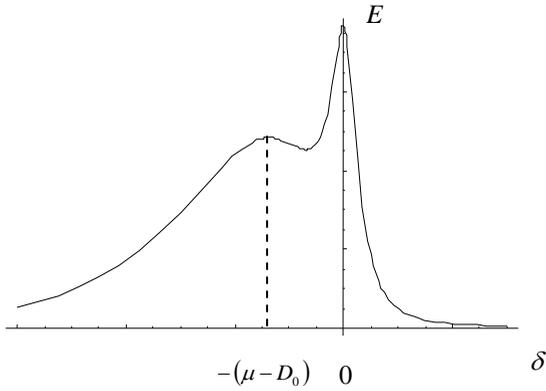


Figure 5: Behaviour of the expected utility for increasing values of μ_i .

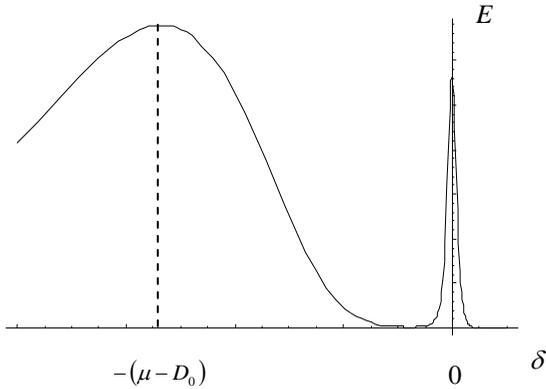


Figure 6: Behaviour of the expected utility for increasing values of μ_i .

These plots thus show a change of phase in expected starting fractal dimension based on small changes to the model parameters. The solution set is restricted to either “no change” ($\delta = 0$) or a significant move to an expected value of D_0 . If we plot the maxima of the expected utility function using just the mean μ_i and A as variables and keeping all other parameters constant, it is possible to build a cusp catastrophe surface of the maxima of the expected utility function as in Figure 7.

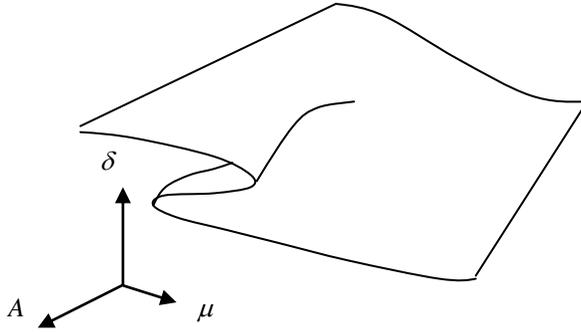


Figure 7: A cusp catastrophe surface

The different “sheets” of the surface correspond to different “best” starting points for the force in terms of the fractal dimension. In the region of the fold, there is more than one maximum for the expected utility function and therefore, at these values of A and μ the global maximum of the expected utility function must be taken. If the other parameters η , V and B are varied, the cusp catastrophe becomes a higher order catastrophe surface in five dimensions.

Small Number of End States

During a single simulation combat model run, the value the fractal dimension takes varies slightly and throughout the run, the fractal dimension takes a range of values. At any point in time in the model run, the probability of the fractal dimension taking a particular value can be approximated by a Normal distribution centred on a mean value of μ with a variance of V as indicated by the equation

$$f(D) = \frac{1}{\sqrt{2\pi V}} \exp\left[-\frac{(D-\mu)}{2V}\right] \quad (17)$$

Changing the starting value of the fractal dimension by an amount δ means that the probability distribution for the fractal dimension throughout the model run can also be approximated by a Normal distribution as in Equation 18. However, in this analysis, the variance of the belief of the fractal dimensions the force takes during the model run will not increase with a change in start point. This is because it is assumed

that the small deviations to the initial fractal dimension value are always of the same magnitude no matter what the starting value is. Therefore, the new probability of each fractal dimension occurring during a model run when the start point is changed is written as:

$$f(D|\delta) = \frac{1}{\sqrt{2\pi V}} \exp\left[-\frac{(D - (\mu + \delta))^2}{2V}\right] \quad (18)$$

In the previous discussion, we assumed that the number of end states or attractors in a system was relatively large. If the set of attractors in the system is small, it is possible to state at each value of the fractal dimension, what the probability of achieving each end state is and that end state's utility. Each end state has a fixed distribution of utility over the fractal dimension and we assume this utility can be approximated as a mixture of Normal distributions. Assuming that there are only two possible end states, the utility function can be written as:

$$U(D) = X \exp\left[-\frac{(D - \mu_1)^2}{2\sigma_1^2}\right] + (1 - X) \exp\left[-\frac{(D - \mu_2)^2}{2\sigma_2^2}\right], \quad (19)$$

where μ_1, μ_2 are the means and σ_1^2, σ_2^2 the variances of each Normal distribution making up the mixture distribution of utilities relating to each of the end states.

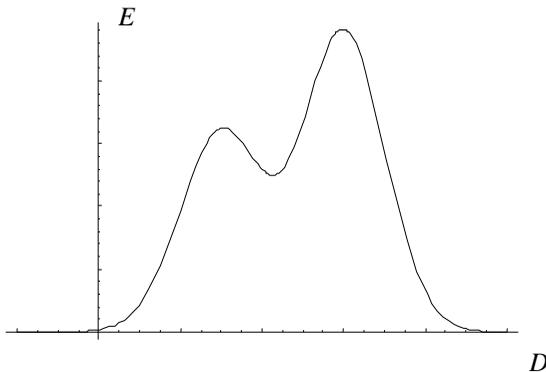


Figure 8: A mixture utility function

Figure 8 shows an example of a mixture utility function. The expected utility $E(\delta)$ from each decision and so starting fractal dimension can be calculated from, as before,

$$E(\delta) = \int_{\forall D} f(D|\delta)U(D)dD = \int_{\forall D} \left(\frac{1}{\sqrt{2\pi V}} \exp\left[-\frac{(D-(\mu+\delta))^2}{2V}\right] \right) \left(X \exp\left[-\frac{(D-\mu_1)^2}{2\sigma_1^2}\right] + (1-X) \exp\left[-\frac{(D-\mu_2)^2}{2\sigma_2^2}\right] \right) dD \quad (20)$$

$$= X \sqrt{\frac{\sigma_1^2}{\sigma_1^2+V}} \exp\left[-\frac{(\delta-(\mu_1-\mu))^2}{2(\sigma_1^2+V)}\right] + (1-X) \sqrt{\frac{\sigma_2^2}{\sigma_2^2+V}} \exp\left[-\frac{(\delta-(\mu_2-\mu))^2}{2(\sigma_2^2+V)}\right]$$

Maximising this function gives the “best” starting points of the fractal dimension relative to the prior knowledge of the model’s behaviour. It can be shown that the function $E(\delta)$ exhibits a catastrophe with a cusp point which can be determined analytically.¹⁰ The global maximum value of the expected utility function depends on the values of μ_1 , μ_2 , σ_1^2 , σ_2^2 and the relative weight, X , of the utilities. At certain points of these parameters, there will only be one local maximum of the expected utility function corresponding to one or other of the utility maxima. At others there will exist two local maxima, for which cases, the best starting value of the fractal dimension will correspond to the global maximum. The following figures show the effect of increasing the relative utilities of each end state.

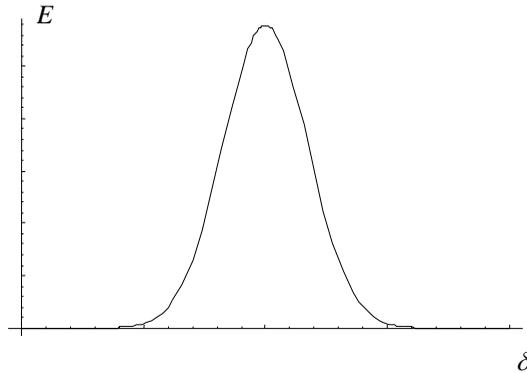
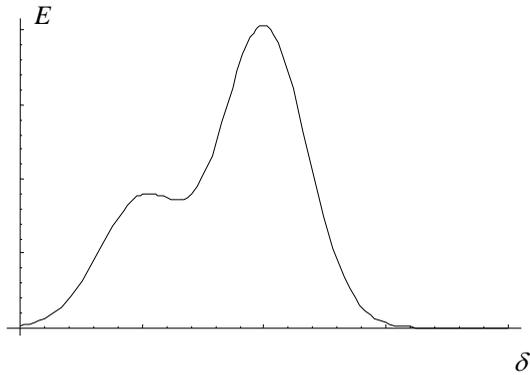
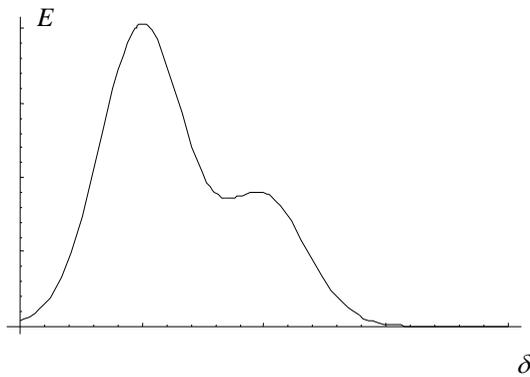


Figure 9: Goal utility – changing μ_i

Figure 10: Goal utility – changing μ_i Figure 11: Goal utility – changing μ_i

The fractal dimensions with greatest utility are always either μ_1 or μ_2 . Similar effects can be seen with changes to the values of the utility variance, σ_1^2 and σ_2^2 . Keeping σ_2^2 fixed and changing σ_1^2 gives the following plots.

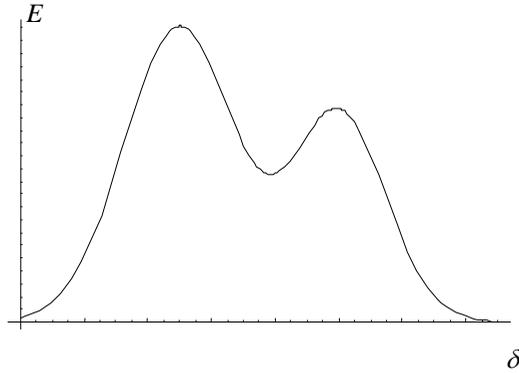


Figure 12: Goal utility – changing utility variance

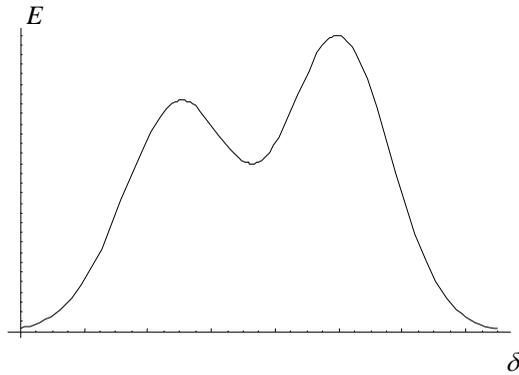


Figure 13: Goal utility – changing utility variance

Again, the only two points of maximum utility (either local or global) are the means μ_1 or μ_2 in the mixture of Normal distributions making up the utility function. The regions of the parameter space $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2, X, V)$ can then be classified according to the starting value of the fractal dimension of the model run with the highest expected utility. This will produce a higher order catastrophe surface.

By including more than two Normal distributions in the mixture making up the utility function, it is possible to examine the effects of more than two end states, or end states which have a more complicated relationship to the starting fractal dimension of the force.

Notes:

- ¹ Andy Illachinski, "Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial Life Approach to Land Warfare," *Research Memorandum CRM 97-61* (Center for Naval Analyses, June 1997).
- ² <www.cna.org/isaac> (June 2002)
- ³ G.I. Barenblatt, *Scaling, Self Similarity and Intermediate Asymptotics*, Cambridge Texts in Applied Mathematics, (Cambridge University Press, 1996).
- ⁴ Barenblatt, *Scaling, Self Similarity and Intermediate Asymptotics*.
- ⁵ Jim Moffat, M. Passman, "Metamodels and Emergent Behaviour in Models of Conflict," in *OR Society Simulation Study Group Workshop Proceedings*, (The Operational Research Society, UK, March 2002).
- ⁶ Jim Moffat, "Command and Control in the Information Age – Representing its Impact," in press, (The Stationery Office, UK, 2002).
- ⁷ M. Lauren, "Firepower Concentration in Cellular Automata Models – An Alternative to the Lanchester Approach", *DOTSE New Zealand Report 172 NR 1350 ISSN 1174-3387*, (2000).
- ⁸ D.L. Turcotte, *Fractals and Chaos in Geology and Geophysics* (Cambridge, UK: Cambridge University Press, 2nd edition, 1997).
- ⁹ J.Q. Smith, P.J. Harrison, E.C. Zeeman, "The Analysis of Some Discontinuous Decision Processes," *European Journal of Operational Research*, 7 (1981): 30-43.
- ¹⁰ J.Q. Smith, "Mixture Catastrophes and Bayes' Decision Theory," *Mathematical Proceedings of the Cambridge Philosophical Society*, 86 (1979): pp.91-101.

Professor **JIM MOFFAT** is a Fellow at the Defence Science and Technology Laboratory, UK, a Fellow of Operational Research and a visiting Professor at Cranfield University, UK. He was awarded the President's medal of the Operational Research Society for the year 2000. He holds a first class honours degree in Mathematics and a PhD in Mathematics, and was awarded the Napier medal in Mathematics by the University of Edinburgh. He has worked for the past 20 years on defence related operational analysis problems and aerospace technology research. His current research interest is in building analysis tools and models which capture the key effects of human decision making and the other aspects of C4ISR.

SUSAN WITTY has an honours degree in Mathematics and a PhD in aero-acoustics research. She has been an analyst in the Defence Science and Technology Laboratory for about three years. Her current research interests are complexity science and Bayesian networks applied to high level data fusion.