

ORGANISATIONAL ADAPTATION ON RUGGED

LANDSCAPES

Tony Brabazon¹

Faculty of Commerce
University College Dublin

Robin Matthews

Centre for International Business Policy
Kingston Business School
London

ABSTRACT

Organisational adaptation, defined in broad terms as ‘a change in a significant attribute of the organisation’ (Levinthal, 1997, p. 934), is a recurrent theme in organisation science. Adopting a similar perspective to Levinthal and March (1981), that the process of organisational adaptation can be conceptualised as a search process, this paper argues that search and landscape metaphors are closely related and understanding of each is illuminated by an understanding of the other. It is also argued that meaningful applications of these metaphors in organisational science can only arise from a clear definition of both a search process and the related landscape. Formal definitions of each, and of a ‘rugged landscape’, are presented and some of the potential implications stemming from these and from the no-free-lunch theorem are discussed. It is noted that the formal definitions permit the linking of several important concepts in organisational science which arise from the common root of bounded rationality, those of myopic search, founding imprint, competency trap and the choice of problem representation.

¹ Corresponding author: Tel: +353 1 7168705 e-mail:anthony.brabazon@ucd.ie

INTRODUCTION

The landscape metaphor (*'surfaces of selective value'*) was first introduced by Wright (1932). The metaphor sought to explain Darwinian evolution (Darwin, 1859) as a search on a landscape, where the base of the landscape is defined by a species' genetic composition and the height corresponds to a measure of the 'fitness' of the species. In this framework, biological evolution represents a search over a genotypic space, in an effort to enhance phenotypic fitness. Abstracting the idea of a landscape from an evolutionary setting, it can be applied whenever the outcome from a process is dependent on several inputs. For example, consider a system which is fully specified by two variables, each of which is represented by a single axis, with the output of the system being plotted on a third axis. If all possible input-output vectors are plotted in three-dimensional space, a landscape (surface) is produced. A landscape can therefore be considered as a graphical representation of a payoff matrix (Matthews, 1999a, 1999b). The concept can be extended beyond three dimensions to encompass more complex systems. Thus, a common (but as shall be seen, incomplete) definition of a landscape is a mapping from a sequence (input) space into the real numbers, assigning to each sequence a particular fitness value (Eigen, 1971). This can be concisely stated as $f: X \rightarrow \mathfrak{R}$, where f is a fitness function, defined for all $x \in X$, of a configuration set X . By changing the definition of the system under consideration, the landscape metaphor can be extended to applications including combinatorial chemistry (Stadler and Stadler, 2001; Gillespie, 1983), biological evolution (Kauffman, 1993; Macken and Perelson, 1989), physics (Weinberger, 1991) and computer science (Jones, 1995). Applications of the metaphor in organisational science are outlined in the next section.

The motivation for this paper is the widespread use of search and landscape metaphors in organisational science. It is argued that a trans-metaphorical application of these concepts requires that they be rigorously defined. This paper integrates literature from computer science, organisational learning, and biological evolution, to provide a conceptual framework. The aim is to assist the understanding of organisational search processes by analysing how these processes operating on suitably defined landscapes. In the context of organisational science, the paper addresses the following:

- ◆ What is a landscape?
- ◆ How are search and landscape metaphors linked?
- ◆ What is a rugged landscape?
- ◆ Can the degree of ruggedness be measured?
- ◆ Does the no-free-lunch theorem have implications for simulation studies of organisational adaptation?

The remainder of this paper is organised as follows. Initially, a general introduction to the application of the landscape metaphor in organisational science is provided, followed by a formal definition of both a landscape, and a search process. A discussion of landscape topology or '*ruggedness*' follows. Implications of landscape definitions and search processes for organisational adaptation are considered and finally conclusions are provided.

THE LANDSCAPE METAPHOR IN ORGANISATION SCIENCE

The Schumpeterian view of creative destruction considers that the lure of profit drives a never-ending search for innovation in both technological (Schumpeter, 1934) and

organisational design (Schumpeter, 1934). Considering innovation as a search process suggests a number of relevant questions:

- ◆ What is being searched?
- ◆ How do organisations search ?
- ◆ What search strategies work best?

Although search and landscape metaphors closely parallel the concept of seeking sustainable competitive advantage (Porter, 1980, 1985), it is only in recent years that the importance of landscapes has been recognized in organisational science (Beinhocker, 1999; Kitts, Edvinsson and Beding, 2001; McCarthy and Tan, 2000; Rivkin, 2000; Gavetti and Levinthal, 2000). As suggested in the questions above, the landscape metaphor is implicit in a long stream of literature which conceptualises organisational learning as search (Levinthal and March, 1981; March ,1991), in that the concept of search is only meaningful in the context of :

- ◆ A defined search space
- ◆ A means of traversing that search space
- ◆ The ability to determine the ‘quality’ of a proposed solution

The landscape metaphor can be enhanced by coupling it with the biological concept of co-evolution (Eisenhardt and Galunic, 2000), whereby strategic landscapes are continually evolving or ‘deforming’ in response to the actions of other firms, suppliers, staff and customers. Dynamic, deforming landscapes imply that firms cannot simply continue to exploit their current capabilities, as there is no guarantee that current locations of high fitness will remain unchanged over time. Rather, it is suggested that the challenge for management is to strike an appropriate balance between exploitation and exploration (March, 1991) and to attempt to influence the evolutionary trajectory of their industry. A key argument of this paper is that

although an appeal to a landscape metaphor is often made in discussing organisational adaptation, a definition of the search process and the related landscape is rarely provided in a rigorous fashion. This paper both provides a definition of these interlinked terms and outlines a number of implications stemming from them. As a result, the paper links the concept of search in an organisational context to the well developed ideas of search in the domains of operations research and artificial intelligence. It is noted that a strong tradition of computational/simulation research exists in organisation science (Carley, 1995; Gavetti and Levinthal, 2000; Rivkin, 2000). The utility of these studies depends critically on their analytical foundations. A clear definition of landscapes and search processes can assist in establishing the likely robustness of the findings of these studies.

SEARCH AND LANDSCAPES

A search algorithm can be defined as '*a strategy or plan to efficiently locate a global extrema of a mapping*' (Wolpert and MacReady, 1995). Although some search algorithms such as simulated annealing (Kirkpatrick, Gelatt and Vecchi, 1983) and genetic algorithms (Holland, 1975) have global search potential, no algorithms exist which guarantee the determination of global extrema for the kind of non-linear problems facing management strategists. Indeed it is highly probable that no such algorithms exist. In an organisational setting, management can encourage various forms of search by varying resource allocations, determining the degree of tolerance of failure and through hiring/promotional policies. For example, by permitting employees to spend a portion of their time on personal research projects, an organisation is engaging in a populational (global) search process. Conversely, by

insisting that all projects pass stringent financial criteria, an exploitive or ‘greedy’ search strategy is perused. As will be discussed further below, the optimal form of search depends critically on the nature of the search space (Wolpert and MacReady, 1995). Hence, discussions as to the form of innovation/search that an organisation should undertake, maintain an implicit, often unrecognised, assumption regarding the nature of the landscape which is being searched.

In general terms, a search process can be characterised as ‘find an object with the following properties’. The search is considered to take place amongst a potentially infinite set of objects and the search heuristic determines how the search process proceeds around the set of these objects. The desired property of the object will be problem-specific.

The first step in a search process is to select the set of objects, this selection provides a basic definition of the problem. Let \mathcal{O} represent the object space. The second step in a search process is to select the representation of the objects in \mathcal{O} . This representation determines a set, \mathcal{R} , the representation space. For example, in organisational settings, the choice of \mathcal{R} will reflect firm culture, perhaps arising from a ‘founding imprint’ or the past experience of decision-makers (Gavetti and Levinthal, 2000). Thus, \mathcal{R} represents a specific mapping of \mathcal{O} . For $o \in \mathcal{O}$ and $r \in \mathcal{R}$. In particular, we can write $o \zeta r$ to indicate that o is represented by r . In a broad class of search problems, the aim is to find as good an object as possible, within a time constraint. As it is generally impossible in large problems, to enumerate for the purpose of search every $o \in \mathcal{O}$, a non-random search process will require a function $g: \mathcal{O} \rightarrow \mathcal{G}$ for some set \mathcal{G} and a partial order $>_{\mathcal{G}}$, such that for $o_1, o_2 \in \mathcal{O}$ if $g(o_1) >_{\mathcal{G}} g(o_2)$, then $g(o_1)$ is considered a better answer to the search problem than $g(o_2)$.

Define \mathbf{O} as the object space (for example) for organisational structure. The search process is limited to a subset of this space, \mathbf{R} . Taking a specific vector of variables which represent an organisational structure, $\mathbf{v} \in \mathbf{R}$, a ‘neighbourhood’ $N_\phi(\mathbf{v})$ can be defined such that it represents the set of elements of \mathbf{R} which can be reached through a single application of a ‘move’ operator, ϕ . A move operator is a function $\phi: \mathbf{R} \times \mathbf{R} \in [0..1]$. The value of $\phi(\mathbf{v}, \mathbf{w}) = p$ for $\mathbf{v}, \mathbf{w} \in \mathbf{R}$ indicates the probability (p) that a single application of the operator transforms $\mathbf{v} \rightarrow \mathbf{w}$. Therefore, in a search context, a move operator usually generates a variant of a parameter vector such as \mathbf{v} . Finally, a search algorithm requires the existence of a mapping $f: \mathbf{R} \rightarrow \mathbf{F}$ for a set \mathbf{F} (a fitness or cost function) and a partial order $>_F$ over \mathbf{F} such that if $\mathbf{v}, \mathbf{w} \in \mathbf{R}$ and $f(\mathbf{v}) >_F f(\mathbf{w})$, then \mathbf{v} is considered ‘better’ than the \mathbf{w} . The choice of this fitness function $f(\mathbf{v})$ lies in the hands of the modeller.

From the above, a landscape may be formally defined as: $L = (\mathbf{R}, \phi, f, \mathbf{F}, >_F)$ (Hordijk, 1995; Jones, 1995). The components are:

- ◆ The representation space \mathbf{R}
- ◆ A move operator ϕ
- ◆ The function $f: \mathbf{R} \rightarrow \mathbf{F}$
- ◆ A set of fitness values \mathbf{F}
- ◆ A partial order $>_F$ over \mathbf{F}

It is posited that this formal definition can be employed to understand the behaviour of search algorithms, and therefore organisational adaptation. It defines a graph, where the vertices correspond to individual configurations, labelled with the relevant fitness values, and the edges correspond to actions of a move operator (Jones, 1995). Traversing an edge of this graph is equivalent to taking a step on the related landscape. If the move operator is applied in a stochastic fashion, the traversal of an

edge can be recast as a Markov chain, whereby the probability of traversing an edge is equivalent to the transition probability of moving from state (or configuration) x to state y .

A simple example, drawn from Hordijk (1995) to demonstrate the above definition follows:

R The space of bit-strings of length 2 (or $\{0,1\}^2$)

φ Flip a randomly chosen bit in the string ($0 \leftrightarrow 1$)

f Number of 0s in the bit string

F $\{0,1,2\}$

$>_F$ $2 >_F 1 >_F 0$

The landscape can be represented as the following, undirected, graph:

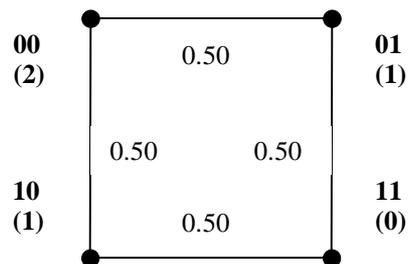


Figure 1: Fitness landscape, represented as a graph, for bit-strings of length two. The vertex labels show the bit-string and associated fitness of each vertex. All transition (edge) probabilities are 0.50.

This leads to a natural definition of the basis of attraction for a vertex v , as the set of vertices that are connected to v by a directed path. Equivalently, the basis of attraction is comprised of the set of vertices from which a random walk, governed by the move operator φ , could reach v .

The definition of a landscape in turn allows a concise definition of a search algorithm. A search algorithm consists of a navigation strategy, and a structure to be searched. The structure is the representation space and the navigation strategy determines from where the search process is started, how the search process moves between points on the landscape at each iteration and finally, the stopping rule for the search process.

Implications

The formal definition of a landscape highlights a number of issues. The choice of representation space is crucial to the discussion of a landscape. Typically although not necessarily, the representation space serves to compress the object space in order to simplify the search process. This accords with

- (a) The idea of bounded rationality in decision-making.
- (b) The necessity of restricting the domain of search.
- (c) It also highlights the importance of the specification of move operators.

(a) Bounded rationality

Decision makers must rely on imperfect cognitive representations in forming mental models of their environment (Gavetti and Levinthal, 2000). Simon (1996) suggests that ‘...*solving a problem simply means representing it so as to make the solution transparent*’ (p. 132). Examples of this in an organisational setting include, mental models of managers based on prior experience, organisational culture values and so on (Gavetti and Levinthal, 2000).

(b) The search domain

As Kauffman (1993) points out, the potential search domain is immense and representation that limit its size are naturally sought after by managers. The importance of choice of representation is noted by Levinthal and Warglien (1999), in that the choice of strategy can *'serve to create an illusionary surface - a social construction as opposed to a caricature of an actual fitness landscape'* (p. 353). Drawing a linkage with the idea of generic strategies as such a social construction, the authors suggest that *'The articulation of these [Porter's] generic strategies transforms what is otherwise an extraordinary rugged terrain of a vast number of local peaks, to a relatively smooth surface with perhaps three local peaks corresponding to the coherent articulation of each of the three generic strategies.'* (p. 353)

(c) *Move operators*

Choice of representation space implicitly determines the perceived range of move options available to the organisation. Move operator(s) specify a neighbourhood function which enable us to identify higher and lower fitness regions nearby: peaks or valleys can only be defined in relation surrounding landscape. For example, a local maximum $x \in X$ is a point for which a neighbourhood $N \in N(x)$ exists such that $f(x) \geq f(y)$ for all $y \in N$.

Generalising the definition of a neighbourhood, the distance between two vertices $a, b \in X$ can be defined as the number of applications of the move operator required to move from one vertex to the other. Critically, each move operator gives rise (a) to traversing different parts of a landscape and (b) to different landscapes (since different move operators may define a different representations or configuration sets, Jones, 1995) and hence may present varying difficulties of search.

In an organisational setting, varying move operators may result in outcomes as diverse as incremental change to existing operations or radical redesign of business processes.

Following from these implications, concepts of exploitation and exploration (March, 1991) implicitly rely on the definition of representation space, search domain and move operator. In the absence of clear definitions, discussion of exploitation or exploration are meaningless. As each move operator defines its own landscape, it may be more appropriate to consider that exploration and exploitation *occur on different landscapes*, or at least different parts of a given landscape, with differing transition probabilities governing each move operator and different expected (risk-adjusted) costs and benefits associated with each. These separate landscapes need not be directly mappable to one another.

RUGGEDNESS OF LANDSCAPES

The difficulty of search on a landscape is often considered as a function of the ‘ruggedness’ of the landscape. Intuitively, a rugged landscape is one with multiple local maxima (minima), the more rugged a landscape, the lower the average fitness correlation between neighbouring points. The notion of ‘ruggedness’ fits well with the discussion of local or ‘myopic’ search (Levinthal and March, 1993; Levinthal, 1997) and competency-traps (Levitt and March, 1988) in organisational settings. The more rugged the strategic landscape facing an organisation, the less viable the management choice perspective becomes.

On rugged landscapes, it is difficult to see or search for global peaks. Even if their location is known, attempts to move an organisation towards them may encounter deep valleys of low fitness on the landscape. The landscape metaphor can also

incorporate the polar views regarding strategic change, those of inertia and strategic choice. The ruggedness of the landscape facing the firm determines the success of managerial choice.

The definition of a landscape provides insight into the causes of ruggedness, highlighting depends both the importance of the geometry of the search space induced by the move operator (Stadler and Stadler, 2001), and the nature of the fitness function, as demonstrated by Kauffman's NK[C] model (Kauffman, 1993).

To date, the importance of the fitness function has been recognised (Levinthal, 1997) but the role of the move operator in determining the degree of landscape ruggedness has received little attention. The relative efficiency and expected payoff of alternative search processes is determined by the ruggedness of the landscape they create. On a highly rugged landscape, the fitness associated with a particular location contains little information regarding the likely fitness of neighbouring points. Rivkin (2000) employs the idea of a rugged landscape to explain why successful organisational strategies can defy imitation by competitors over long periods, even when the success of the strategy is widely known. If the strategic landscape faced by organisations is rugged, even small errors in copying a successful strategy can lead to large penalties in terms of organisational performance. Unfortunately, despite the intuitive appeal of the importance of the concept of landscape ruggedness, the nature of organisational landscapes is likely to defy any simple attempt at quantification of their ruggedness.

Measures of Correlation

A variety of formal measures of ruggedness do exist, including correlation functions, length of adaptive walks and the density of local optima (Reidys and Stadler, 2001). Most attention has been paid to the correlation metric, typically defined in the context

of landscapes as the degree of positive autocorrelation between fitness values produced by a random walk on the landscape (Weinberger, 1990). Intuitively, the higher the degree of autocorrelation, the smoother the fitness landscape (Kauffman, 2000). On a correlated, or smooth, landscape, close neighbours have similar fitness. As the landscape becomes increasingly rugged, the AR(1) correlation tends to zero and there is no distinction between local and distant search as both represent voyages into the unknown.

Following the definition of Hordijk (1995), the autocorrelation function ρ_i relates the fitness value of two points in the time-series of fitness values produced by a random walk which are i steps apart. Consequently, the autocorrelation for time lag i of a time series y_t , $t = 1, \dots, T$ is defined as:

$$\rho_i = \text{Corr}(y_t, y_{t+i}) = \frac{E[y_t y_{t+i}] - E[y_t] E[y_{t+i}]}{\text{Var}(y_t)} \quad (\text{of course, } E[y_t] E[y_{t+i}] = E[y_t]^2)$$

where $E[y_t]$ is the expected value of y_t and $\text{Var}(y_t)$ is the variance of y_t over the entire landscape. If $|\rho_i|$ is close to one, then there is strong correlation between fitness values i steps apart.

The definition of the autocorrelation function can be extended to encompass a definition of the *correlation length* of a landscape (Hordijk, 1995). Intuitively, this measures the distance from a random starting (or subsequent) point before the landscape fitness values obtained from a random walk are considered to be uncorrelated. Therefore, beyond the correlation length, the random walk has no memory. The step i auto-correlations can be condensed into a single metric of correlation length:

$$\sum_{i=0}^{\infty} p_i$$

The smaller this measure, the more rugged the landscape.

Anisotropic Landscapes

The discussion above assumes that the landscape is statistically *isotropic*. In this case a random walk is “representative” of the entire landscape and is sufficient to gather statistics on the shape of the landscape (Greenwood and Hu, 1998). A landscape is defined as being isotropic (Stadler and Gruner, 1993) if for any connected sub-graph A of a partition of the configuration space B :

$$[(f(x) - f(y))^2]^A_d \approx [(f(x) - f(y))^2]_d$$

where $[\cdot]^A_d$ represents an average taken over all pairs (x,y) , with $x,y \in A$ and $\text{distance}(x,y) = d$. Similarly, $[\cdot]_d$ represents an average taken over all pairs (x,y) , in the total configuration space with $\text{distance}(x,y) = d$.

Despite the widespread usage of correlation function as a measure of the ruggedness of a landscape, the metric is suspect in considering landscapes in an organisational setting. It may be plausibly argued that searches on organisational landscapes are typically constrained. Constraints may arise from organisation history (Levitt and March, 1988; Boeker, 1989), culture, the organisation’s existing network of suppliers and customers, its ability to absorb new technology (Cohen and Levinthal, 1990 and 1994) and the *stickiness* encountered when attempting to transfer innovations in-house between organisational sub-units (Szulanski, 1996). The affect of these constraints or ‘routines’ (Levitt and March, 1988) is to create infeasible regions (areas of very low fitness or *sink holes*) on the fitness landscape. The problems of the autocorrelation measure on such a landscape can be clearly seen by considering a landscape with large infeasible regions. If the random walk enters such a region and stays there, the

resulting the (low) fitness values may be highly correlated, giving the false impression that the landscape is smooth.

Greenwood and Hu (1998) point out that if the landscape represents a constrained combinatorial problem, the assumption of statistical isotropy is inappropriate, rather the landscape is considered to be statistically *anisotropic*. In such cases, the fitness values produced from a random walk depend both on the starting point and the random choices at each step. Stadler and Gruner (1993) speculate that on highly anisotropic landscapes *'the particular geometric details of the anisotropies are most important.'* (p. 383). In addition, the appropriateness of a correlation measure generated from a random walk is further suspect as organisational search is unlikely to be completely random, rather it will represent a 'biased' walk, whereby the transition probabilities between neighbouring states are not independent of expectations as to their fitness values. In terms of organisational science, the idea of searching on an anisotropic landscape fits well with existing research regarding the significance of a 'founding imprint' (Boeker, 1989), path-dependence (Levinthal and March, 1991) and also to Arthur's work on tipping industries (Arthur, 1989).

History matters, not just in terms of defining the current location of an organisation on a landscape, but in defining the representation (move-set and connectedness of firm activities) utilised by decision-makers, the very nature of the landscape itself. History also impacts on the 'discount factor' applied by decision-makers in evaluating the potential payoffs from movements on the landscape. Even if a move is possible, the likelihood of its occurrence depends on the discount factor applied to perceived future costs and benefits arising from the move. If organisations primarily engage in local search as suggested by Levinthal (1997), the founding imprint is a key determinant of organisational success.

No-Free-Lunch Theorem

The importance of understanding the nature of the landscape being searched is reinforced by the no-free lunch (NFL) theorem of Wolpert and MacReady (1995). This suggests that *all* algorithms that search for an extremum of a cost function perform *exactly the same, according to any performance measure*, when averaged over all possible cost/fitness functions, if none of the properties of the cost/fitness function are taken into account. This implies that the utility of a search process depends on the nature of the landscape being searched which in turn is partly defined by the form of the move operator. Unless the search algorithm embeds knowledge concerning the cost function, there is no reason to suppose *a priori* that it will perform any better than random guessing. This non-intuitive result implies that universal search procedures which will produce high-quality results on all possible landscapes, do not exist. In essence, the usefulness of any search procedure depends on the properties of the landscape it is searching.

The NFL theorem is restricted to the case where the cost/fitness function is constant, in many business settings the fitness function will be a relative measure. The degree of applicability of the NFL theorem to this latter case is as yet, unknown. However, it does suggest caution against assuming that a ‘good’ method of search in one industry or organisational setting, will be useful in another. Analytically, the NFL theorem explains why generic search methods such as genetic algorithms or simulated annealing are outperformed by tailored, domain-specific methods, for example, specialised heuristics for the NP-hard, TSP problem.

From an organisational science perspective, the NFL theorem raises several open questions:

- ◆ What form do fitness (cost) functions take in organisational science settings?
- ◆ Following from the above, how generalisable are the results of simulation-based studies in organisation science which purport to show the utility of search processes in finding an extremal peak on a rugged landscape?

CONCLUSIONS

This paper addresses a number of the implications arising from the conceptualisation of organisational adaptation as a search process. It is posited that these implications can only be understood by starting from a clear definition of both a search process and the related landscape which is being searched. Formal definitions of these were provided leading to a discussion of the implications of these definitions. The formal definitions permit the linking of several important concepts in organisational science which arise from the common root of bounded rationality, those of myopic search, founding imprint, exploration vs exploitation, competency trap and the choice of problem representation. All of these embed the idea of a search metaphor.

Three key points emerge from the definitions:

- ◆ Search and landscape definitions are intertwined. The choice of move (search) operator partly determines the topology of the related landscape
- ◆ A distinction exists between object space and representation space. Invention takes place in representation space rather than object space.
- ◆ Landscapes have different degrees of '*ruggedness*'

The search heuristic(s) adopted by an organisation are a matter of choice. Insistence that all project proposals pass a stringent financial test corresponds to a 'greedy' search strategy. In contrast, permitting employees to engage in 'personal projects' corresponds to a more permissive, exploration search heuristic.

Search takes place in representation, not object, space. Due to bounded rationality, not all possibilities can be considered, hence strategists restrict the search space by adopting a specific representation of it. This eases the search task by smoothing the landscape, but may limit the search potential. Many representations of the same object space exist, with varying potential. Landscape ruggedness is impacted by the choice of representation and the choice of search heuristics. The degree of ruggedness determines the ease with which strategists can assess ex-ante, the likely payoffs to alternative strategies.

It is acknowledged that organisational adaptation, at populational level, may result from a variety of processes, including evolutionary pressures (Hannan and Freeman, 1977). The paper focuses on adaptation of individual organisations and therefore does not consider the ecological argument.

REFERENCES

- Arthur, B. (1989). Competing Technologies, Increasing Returns and Lock-In by Historical Events, *The Economic Journal*, 99:116-131.
- Beinhocker, E. (1999). On the Origin of Strategies, *The McKinsey Quarterly*, 4:46-57.
- Boeker, W. (1989). Strategic Change: The Effects of Founding and History, *Academy of Management Journal*, 32(3):489-515.
- Carley, K. (1995). Computational and Mathematical Organization Theory: Perspective and Directions, *Computational and mathematical Organization Theory*, 1(1): 39-56.
- Carley, K. and Svoboda, D. (1996). Modeling Organizational Adaptation as a Simulated Annealing Process, *Sociological Methods and Research*, 25(1):138-168.
- Cohen, W. and Levinthal, D. (1990). Absorptive Capacity: A new Perspective on Learning and Innovation, *Administrative Science Quarterly*, 35:128-152.
- Cohen, W. and Levinthal, D. (1994). Fortune Favors the Prepared Firm, *Management Science*, 40(2):227-250.
- Darwin, C. (1859). *The Origin of Species by Means of Natural Selection*, reprinted (1985), London: Penguin Books.
- Eigen, M. (1971). Self-organisation of matter and the evolution of macromolecules, *Naturwissenschaften*, 10:465-523.
- Eisenhardt Galunic 2000
- Gavetti, G. and Levinthal, D. (2000). Looking Forward and Looking Backward: Cognitive and Experiential Search, *Administrative Science Quarterly*, 45:113-137.
- Gillespie, J. (1984). Molecular Evolution Over the Mutational Landscape, *Evolution*, 38(5): 1116-1129.
- Greenwood, G. and Hu, X. (1998). On the use of random walks to estimate correlation in fitness landscapes, *Computational Statistics & Data Analysis*, 28:131-137.
- Hannan, M. and Freeman, J. (1977). The Populational Ecology of Organizations, *American Journal of Sociology*, 82(5): 929-964.
- Holland, J. (1975). *Adaptation in Natural and Artificial Systems*, 6th reprinting (2001) Massachusetts, MIT Press.
- Hordijk, W. (1995). A Measure of Landscapes, *Santa Fe Institute Working Paper 95-05-049*.
- Jones, T. (1995). *Evolutionary Algorithms, Fitness landscapes and Search*, PhD Thesis, Department of Computer Science, University of New Mexico.
- Kauffman, S. (1993). *The Origin of Order*, Oxford, England: Oxford University Press.
- Kauffman, S. (2000). *Investigations*, Oxford, England: Oxford University Press.
- Kirkpatrick, S., Gelatt, C. and Vecchi, M. (1983). Optimization by Simulated Annealing, *Science*, 220(4598):671-680.
- Kitts, B., Edvinsson, L. and Beding, T. (2001). Intellectual Capital: from intangible assets to fitness landscapes, *Expert Systems with Applications*, 20:35-40.
- Levinthal, D. (1997). Adaptation on Rugged Landscapes, *Management Science*, 43(7):934-950.
- Levinthal, D. and March, J. (1981). A Model of Adaptive Organizational Search, *Journal of Economic Behavior and Organization*, 2:307-333.
- Levinthal, D. and March, J. (1993). The Myopia of Learning, *Strategic Management Journal*, 14:95-112.

- Levinthal, D. and Warglien, M. (1999). Landscape Design: Designing for Local Action in Complex Worlds, *Organization Science*, 10(3):342-357.
- Levitt, B. and March J. (1988). Organizational Learning, *Annual Review of Sociology*, 14:319-340.
- Macken, C. and Perelson, A. (1989). Protein evolution on rugged landscapes, *Proceedings of the National Academy of Science (USA)*, 86:6191-6195.
- March, J. (1991). Exploration and Exploitation in Organisational Learning, *Organization Science*, 2(1):71-87 .
- McCarthy, I. And Tan, Y.K. (2000). Manufacturing competitiveness and fitness landscape theory, *Journal of Materials Processing Technology*, 107:347-352.
- Matthews, Robin, (1999a), Firms, Strategies and Co-operative Games: Elements of a Stakeholding Approach, Arestis, P., Daniel, S.S. and Grahl, J. Ed. *Festschrift in Honor of Maurice Peston and Bernard Corry*, Edward Elgar, London.
- Matthews, Robin (1999b), A Spin Glass Model of Diffusion in Organizations, in Fredel Bolle et.al. ed. *Inquiries into the Nature and Causes of Behaviour, Proceedings the International Association for Research in Economic Psychology Colloquium XXIV Vol 2*, Belgirate, Italy.
- Porter, M. (1980). *Competitive strategy : techniques for analyzing industries and competitors*, New York, The Free Press.
- Porter, M. (1985). *Competitive Advantage: Creating and Sustaining Superior Performance*, New York, The Free Press.
- Reidys, C. and Stadler, P. (2001). Combinatorial Landscapes, *Santa Fe Institute Working Paper 01-03-014*.
- Rivkin, J. (2000). Imitation of Complex Strategies, *Management Science*, 46(6):824-844.
- Schumpeter, J. (1934). *The Theory of Economic Development*, eighth printing, 1968, Harvard Economic Studies, Volume XLVI, Harvard University Press.
- Simon, H. (1996). *Sciences of the Artificial* (3rd edition), Cambridge, Massachusetts: MIT Press.
- Stadler, P. and Gruner, W. (1993). Anisotropy in Fitness Landscapes, *Journal of Theoretical Biology*, 165:373-388.
- Stadler, B. and Stadler, P. (2001). Fitness Landscapes, to appear in Ldssig and Valleriani (eds), *Biological Evolution and Statistical Physics*, Springer-Verlag. (<http://www.tbi.univie.ac.at/papers/>)
- Szulanski, G. (1996). Exploring Internal Stickiness: Impediments to the Transfer of Best Practice Within the Firm, *Strategic Management Journal*, 17:27-43.
- Weinberger, E. (1990). Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference, *Biological Cybernetics*, 63:325-336.
- Weinberger, E. (1991). Local properties of Kauffman's Nk model: A tunably rugged energy landscape, *Physical Review A*, 44(10):6399-6413.
- Wolpert, D. and MacReady, W. (1995). No Free Lunch Theorems for Search, *Santa Fe Institute Working Paper 95-02-010*.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution, *Proceedings of the Sixth International Congress on Genetics*, Vol 1., pp.356-366.