
Theses and Dissertations

Spring 2009

Innovation as a complex adaptive system

Joseph John Engler
University of Iowa

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Industrial Engineering Commons](#)

Copyright 2009 Joseph John Engler

This thesis is available at Iowa Research Online: <https://ir.uiowa.edu/etd/233>

Recommended Citation

Engler, Joseph John. "Innovation as a complex adaptive system." MS (Master of Science) thesis, University of Iowa, 2009.

<https://doi.org/10.17077/etd.bnphvn4x>

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Industrial Engineering Commons](#)

INNOVATION AS A COMPLEX ADAPTIVE SYSTEM

by

Joseph John Engler

A thesis submitted in partial fulfillment of the
requirements for the Master of Science degree
in Industrial Engineering
in the Graduate College of
The University of Iowa

May 2009

Thesis Supervisor: Professor Andrew Kusiak

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

MASTER'S THESIS

This is to certify that the Master's thesis of

Joseph John Engler

has been approved by the Examining Committee
for the thesis requirement for the Master of Science
degree in Industrial Engineering at the May 2009 graduation.

Thesis Committee

Andrew Kusiak, Thesis Supervisor

Pavlo Krokhmal

Albert Ratner

TABLE OF CONTENTS

LIST OF TABLES	iv
LIST OF FIGURES	v
LIST OF EQUATIONS	ix
CHAPTER 1. INTRODUCTION	1
CHAPTER 2. HISTORIC PERSPECTIVES OF INNOVATION	5
TRIZ.....	5
The Linear Model of Innovation.....	7
The Pull Model of Innovation.....	9
The Feedback Model of Innovation.....	10
The Strategic Model of Innovation.....	11
The Networked Model of Innovation	12
The Innovation Process.....	14
Conclusions.....	15
CHAPTER 3. MINING THE REQUIREMENTS OF INNOVATION	16
Mining User and Expert Reviews from the Web.....	17
Population of Transactional Database	22
Patent Database Mining	25
Mining Frequent Requirements	27
Conclusions.....	32
CHAPTER 4. THE NK LANDSCAPE OF INNOVATION.....	33
NK Landscapes	34
The Landscape of the System of Innovation.....	37
Historic Landscapes of the System of Innovation	42
Wind Turbine Case Study.....	50
Turbine Blade Case Study	58
Epistatic Interactions in Simple Systems.....	61

Conclusions.....	65
CHAPTER 5. A CELLULAR AUTOMATA MODEL	67
Boolean Networks in the System of Innovation	68
Boolean Network Experimental Results	70
Boolean Networks with Random Neighborhoods	77
Cellular Automata Model of the System of Innovation.....	79
Evolution of Individual Innovations with Multi-Parent Crossover	86
Industrial Examples	90
Conclusions.....	92
CHAPTER 6. AN AGENT-BASED INNOVATION ECOSYSTEM MODEL	94
Agent-Based Modeling.....	94
Updating.....	97
Model Results	100
Discussion and Examples	110
Conclusions.....	113
CHAPTER 7. CONCLUSION.....	114
REFERENCES	116

LIST OF TABLES

Table 1 . The mean starting iteration and mean cycle length of the system of simple boolean agents..... 63

Table 2. Length of period for von Neumann neighborhoods..... 74

LIST OF FIGURES

Figure 1. The linear model of innovation	8
Figure 2. Dell Corporation's innovation model.....	9
Figure 3. The feedback model of innovation	10
Figure 4. Generic Innovation Process.....	14
Figure 5. Sample MP3 player review	18
Figure 6. Sample Walkman review (Authors construct based on multiple review sites). 18	
Figure 7. Heuristic versus classifier results	20
Figure 8. Comparison of heuristic to classifier.....	21
Figure 9. Christmas tree stand patent. U.S. Patent Number 5492301	26
Figure 10. Three frequent itemsets of cardinality 3	31
Figure 11. Partial AND-OR tree for requirements in Fig. 10. (only the branch for cardinality of 3 is shown. The full AND-OR tree would have branches for other cardinalities as well.....	31
Figure 12. The fitness landscape of innovation for the years 1905-1910.....	40
Figure 13. The fitness landscape for the years 1910-1915	41
Figure 14. Fitness landscape for the system of innovation for the years 1800-1900.....	43
Figure 15. Peak fitness versus mean fitness for the innovation landscape of the years 1800 to 1900.	44
Figure 16. Historic fitness landscape of innovation for the years 1900 to 1950.	45
Figure 17. Peak fitness versus mean fitness for the innovation landscape of the years 1900 to 1950.	46
Figure 18. Historic fitness landscape for the years 1950 to 2008.....	47
Figure 19. Peak fitness versus mean fitness for the innovation landscape of the years 1950 to the present.....	48
Figure 20. Relation of the fitness of the optima to the mean fitness of the population for the years 1800 to the present.....	49

Figure 21. The fitness landscape for wind turbines for the years 1900 to the present.....	51
Figure 22. Fitness of the optima versus fitness of the mean of the population for wind turbines for the years 1900 to the present.	52
Figure 23. The relation of the fitness of the optima to the mean fitness of the population for wind turbines for the years 1900 to the present.	53
Figure 24. The number of patents which cite a given patent as basis for their validity ...	55
Figure 25. The number of patents which cite a given wind power related patent as basis for their validity	56
Figure 26. The number of patents which cite a given non-wind power related patent as a basis of their validity.....	57
Figure 27. Historic fitness landscape of a turbine blade for the years 1940 to present	58
Figure 28. Peak fitness versus mean fitness of the population of patents considering turbine blades for the years 1900 to present	59
Figure 29. The number of turbine blade patents citing a given patent as a basis for their validity	60
Figure 30. Mean starting iteration of the period for each value of K and each population size	64
Figure 31. Mean cycle length for each value of k and each population size.	64
Figure 32. Initial populations of a simple Boolean network.....	71
Figure 33. Example of von Neumann neighborhood of size 2	72
Figure 34. Number of transients prior to the start of period with von Neumann neighborhoods.....	73
Figure 35. . Period length for von Neumann neighborhoods.....	74
Figure 36. Period length for epistatic interactions of single bit string experiment.....	75
Figure 37. Transient length with Moore Neighborhoods.....	76
Figure 38. Period length with Moore neighborhoods	76
Figure 39. Number of transients prior to the start of a period for Boolean networks with random neighborhoods.....	78

Figure 40. Period length for Boolean networks with random neighborhoods	78
Figure 41. Sample of initial populations of the CA model. Starting from left, the populations size is 95%, 30% and 5% of the system.....	81
Figure 42. Run of CA model with 5% initial population and only 30% of neighbor success required. The iteration count starting from the left is 0 (initial), 10 iterations, 25 iterations, 50 iterations and 100 iterations.	82
Figure 43. Run of CA model with 5% initial population and 50% of neighbor success required. The iteration count starting from the left is 0 (initial), 10 iterations, 25 iterations, 50 iterations and 100 iterations. Borders were added to the last two graphs due to	83
Figure 44. Run of CA with 95% initial population and 50% of neighbor success required. The iteration count starting from the left is 0 (initial), 25 iterations, 50 iterations, 100 iterations and 250 iterations.	84
Figure 45. Multi-parent crossover illustrated with $n = 3$ parents of a genetic algorithm whose genotype is represented as bit strings.	87
Figure 46. The average population fitness of single point crossover GA compared to a multi-parent GA with 5 parents per offspring.....	89
Figure 47. The fitness of the fittest individual in the population of a single point crossover GA compared to a multi-parent GA with 5 parents per offspring.	89
Figure 48. Iterative agent-based model.....	96
Figure 49. Mean fitness for an average run of the model with the fitness change rate = 0.01, the cost of interaction = 0.01, and the payoff = 0.01.	102
Figure 50. Mean Interaction rate for an average run of the model with the fitness change rate = 0.01, the cost of interaction = 0.01, and the payoff = 0.01.	102
Figure 51. Mean agent resource for an average run of the model with the fitness change rate = 0.01, the cost of interaction = 0.01, and the payoff = 0.01.	103
Figure 52. Mean interaction for an average run with cost of interaction = 0.01 and the payoff = 0.01.	104
Figure 53. Mean interaction for an average run with the interaction cost = 0.01 and payoff = 0.05	104

Figure 54. Mean interaction for an average run with the interaction cost = 0.01 and the payoff = 0.0.10.....	105
Figure 55. Mean interaction for an average run with the interaction cost = 0.01 and payoff = 0.20.	105
Figure 56. Mean interaction for an average run with the interaction cost = 0.01 and the payoff = 0.50.....	106
Figure 57. Mean resources for an average run with the interaction cost = 0.01 and the payoff = 0.01.....	107
Figure 58. Mean resources for an average run with the interaction cost = 0.01 and the payoff = 0.50.....	107
Figure 59. Mean interaction for an average run with the fitness rate of change = 0.01. 108	
Figure 60. Mean resources for an average run with the fitness of rate change = 0.01. ..	109
Figure 61. Mean interaction for an average run with the fitness rate of change = 0.50. 109	
Figure 62. Mean resources for an averaged run with the fitness rate of change = 0.50. 110	

LIST OF EQUATIONS

Equation 1. Ideality	7
Equation 2. Support	28
Equation 3. Fitness.....	35
Equation 4. Genotype Fitness	35
Equation 5. Optima to Mean Relationship Function	49
Equation 6. Updating	97
Equation 7. Agent Fitness.....	97
Equation 8. Market share	98
Equation 9. Payoff	98

CHAPTER 1. INTRODUCTION

It has long been understood that innovation is a driving factor in the economic success of an enterprise. As new and, more importantly successful, innovations come to market older innovations are replaced. The companies, or enterprises, that bring the innovation to market are gaining a temporary dominance in the market place and thus erode the previous dominance of other products or services. This can clearly be seen in the case of the automobile that replaced the horse drawn carriage, or the acceptance of air travel that eroded the once popular train travel industry.

In order to realize success in the global market place companies and enterprises must continuously introduce new inventions to the global marketplace in hopes of having them accepted as innovations. As such, great emphasis placed on innovation in the corporate environment. Understanding the system of innovation and its principles are imperative to the success and longevity of the innovating entity.

Unfortunately, the literature does not present a solid definition of innovation. Nor is there a well formatted hypothesis as to the principles and methodologies by which the system of innovation operates. Various researchers have posited methodologies for new invention but few have considered innovation as a system with principles that govern its operation. The literature fails to agree on even a simple definition of the word innovation.

Stokic *et al.* [1] suggested that innovation is a new product that is introduced to an environment. This statement is restricted to only those products or services that have been fully implemented. Alternatively, Kusiak [2] stated that innovation is an iterative process

aimed at the creation of new products, processes, knowledge or services through the use of new or existing scientific knowledge. Kusiak further stated that product innovation is concerned with the introduction of new goods and services which differ from those currently existing in the marketplace [3].

Incumbent in the later statements, concerning innovation, is the acceptance of the global marketplace while the former suggests that innovation may occur in any environment. While it is well known that invention may occur in any environment, innovation suggests the need for market acceptance. Echerhmann *et al* [4] stated that innovation is more than discovering new ways to create products or processes; innovation is often rooted in the culture itself. This statement clearly differentiates innovation and the mere concept of inventing a new product or process.

The literature further describes innovation as taking place in one of two manners. Incremental, or non-disruptive innovation, is seen as the manner in which innovation builds upon previous innovations and ideas. Radical, or disruptive innovation, is the less common manner in which innovation is sparked from completely new ideas. It has been stated that radical innovation constitutes only about 10% of all innovations [5]. Most often it is the incremental type of innovation that is considered in the context of the literature as it is easily followed.

This thesis considers innovation as a system rather than distinct products. While distinct products can be considered innovative themselves the focus here is on the system of innovation and its governing principles. Various theories of innovation methodologies have been published, (e.g., [6], [7], and [1]) and while each innovation methodology

contributes to better understanding of how innovation can be managed, studies of the origins and principles of the system of innovation are limited.

This thesis hypothesizes that the system of innovation is a complex adaptive system and as such its principles may be fully realized only through modeling the system in its entirety. Interest in complex systems has seen a marked increase over the past two decades. Various approaches have been developed to gain better understanding of these systems. Bandini *et al.* [8] and Chopard *et al.* [9] modeled complex systems as cellular automata. A cellular automaton, with its lattice of agent cells whose behavior is directly influenced by its neighbors, can be viewed as a hierarchical parent to the agent-based model. Cellular automata can thus be considered as a more generalized conception of the full agent-based modeling approach.

An agent-based model utilizes simple agents interacting with their environment through sensors and actions [10]. The environment that the agents interact with is defined by the model which defines the interaction with other agents in the environment as well as external constraints or boundaries. The use of agent-based modeling to understand complex systems is well documented in the literature. For example, Koritarov [11] applied agent-based modeling to simulate electricity market scenarios and Tesfatsion [12] described an agent-based model of complex economic conditions.

Complex adaptive systems (CAS) are also well accounted for in the literature. John Holland posited the concepts of CAS in his well known book [13]. Miller *et al.* [14] also contributed to the field with their treaty on the subject. The topic of technological innovation has also been given brief mention in the CAS literature [15].

The remainder of this thesis is organized as follows. Chapter 2 presents a historic overview of the methodologies that have been hypothesized concerning innovation. Chapter 3 illustrates a methodology for mining the requirements for individual innovations. Chapter 4 introduces the concept of the system of innovation through the view of biological systems and describes the interconnectedness of the system of innovation. Chapter 5 posits a set of rules for the interconnectivity of the system of innovation through a model of CAS called cellular automata. Chapter 6 illustrates an agent-based model for the understanding of driving factors involved in the market perspective of the system of innovation. Finally, conclusions are presented in Chapter 7.

CHAPTER 2. HISTORIC PERSPECTIVES OF INNOVATION

In this chapter, an overview of the historic models of innovation and their associated methodologies will be discussed. This section is divided by model type starting with the Theory of Inventive Problem Solving (TRIZ) and then addressing five generations of innovation models.

TRIZ

The first major model of innovation was put forth by a Soviet Navy patent expert, Genrich S. Altshuller. Altshuller had access to literally thousands of ideas in his trade as a patent expert and noticed some very specific patterns within the patents themselves. Through his observations and subsequent research, Altshuller developed a minimum of criteria that he felt any theory of innovation should satisfy [6]. He stated that any theory of innovation must:

1. be a systematic, step-by-step procedure
2. be a guide through a broad solution space to direct to the ideal solution
3. be repeatable and reliable and not dependent on psychological tools
4. be able to access the body of inventive knowledge
5. be able to add to the body of inventive knowledge
6. be familiar enough to inventors by following a general approach to problem solving.

By the 1960s Altshuller had combed over 200,000 patents looking for inventive problems and how they were solved. In the 1960s and 1970s Altshuller categorized the solutions to those problems into five levels [6].

- Level One. Routine design problems solved by methods well know within the specialty. No invention needed. Accounts for 32% of the solutions.
- Level Two. Minor improvements to an existing system, by methods known within the industry. Accounts for 45% of the solutions.
- Level Three. Fundamental improvements to an existing system, by methods know outside the industry. Contradictions resolved. Accounts for 18% of the solutions.
- Level Four. A new generation which uses a new principle to perform the primary functions of the system. Solutions found more in science than technology. Accounts for 4% for the solutions.
- Level Five. A rare scientific discovery or pioneering invention of an essentially new system. Accounts for 1% of the solutions.

As can be seen from the categorization of the solutions as represented by Altshuller, the majority of innovation takes place in an incremental fashion, building upon previous ideas and technologies. This process is formalized in one of the laws of the Theory of Inventive Problem Solving (TRIZ), as Altshuller's theory became known, called the Law of Increasing Ideality [6]. This law states that technical systems evolve

toward increasing degrees of ideality, where ideality is defined as the quotient of the sum of the system's useful effects, U_i , divided by the sum of its harmful effects, H_j as given in Equation 1.

$$Ideality = \frac{\sum U_i}{\sum H_j}$$

Equation 1. Ideality

Altshuller eventually devised a complete system to aid in the inventive process that included 39 engineering principles, 40 inventive principles, and a table of contradictions that aimed at resolving conflicts between the two sets of principles when applied to innovation problems. Over the years TRIZ has been applied to multiple application areas such as industrial, mechanical, electrical and biological engineering problems. TRIZ has also been expanded over the years in such areas as computer science to instruct critical thinking skills in IT departments [16].

The Linear Model of Innovation

The first generational theory model of innovation is considered to be the linear model. In this model innovation is seen as being pushed from the research phase to the commercial application phase in a unidirectional manner as shown in [17] and [18]. This model poses the designer as king, where the engineer or product designer dictates the product requirements. The designer pushes the product to market based upon technological advances more than market requirements. In the typical linear model, innovation is represented by a pipeline of sequential processes that starts at the pure

scientific research and ends with commercial applications [17]. Savoiz *et al.* [19] depicts this process, for a single product as in Fig. 1.




Figure 1. The linear model of innovation

Collins [20] likened the linear model to a funnel in which an organization pours a host of ideas into the top and, through a series of research efforts and formal assessment steps, a set of useful innovations emerges out the bottom of the funnel. Collins stated that this model has served some companies well such as Intel with its unforgiving delivery schedules but has caused other industries such as the pharmaceutical industry to lag behind and even lose large amounts of revenue. Zhang *et al.* [21] gave a similar description stating that the linear model is based upon the process where scientific and technological advances push a new product to market.

It should be noted that the linear model of innovation deals more with tangible products such as technology rather than the intangible objects of service innovation. Service innovation is more market dictated than product innovation but it may very well still follow the linear model in some instances.

The Pull Model of Innovation

In the second generation model of innovation the consumer or user is considered the guiding entity unlike the designer as king state in the first generation model. In this model, the consumer or end user drives the requirements of the technology. This model has traditionally been known as a pull model as stated by [21] and [20]. Zhang *et al.* [21] note that this model was dominant in the 1960s and was still a linear process where the market needs “pulled” a new product to market.

The Dell Corporation states [22] that the core of their innovation approach is to deliver solutions that directly address customer needs. In this reference, Dell depicts their innovation process model as a simple linear flow diagram shown in Fig. 2.




Figure 2. Dell Corporation's innovation model

It has been stated by Ulwick [23] that this traditional approach of asking customers for solutions tends to undermine the innovation process. This appears to be because most customers have a very limited frame of reference. As stated by Kusiak [24], many customers buy milkshakes based upon the drink's thickness and strong flavor. But the milkshake ends up competing in functionality with the complementary products of

sandwiches, soft drinks and salads. This type of competition in functionality must be known a priori to ensure that customer requirements do not undermine the innovation process.

The Feedback Model of Innovation

The third generation innovation model is known as the feedback model. In this model an initial product offering may have been pushed to market following the first generation linear model but now the end user or consumer has influence over the incremental innovation of this product through their feedback and requirements. Zhang *et al.* [21] stated that in this model the innovation process can be divided into a series of interdependent stages and feedbacks to the previous stage. Kline *et al.* [25] called this model the Chain-Linked Model due to its nature of continuous feedback and improvement.

This model combines the previous two generations of innovation models as shown in Fig. 3 below.

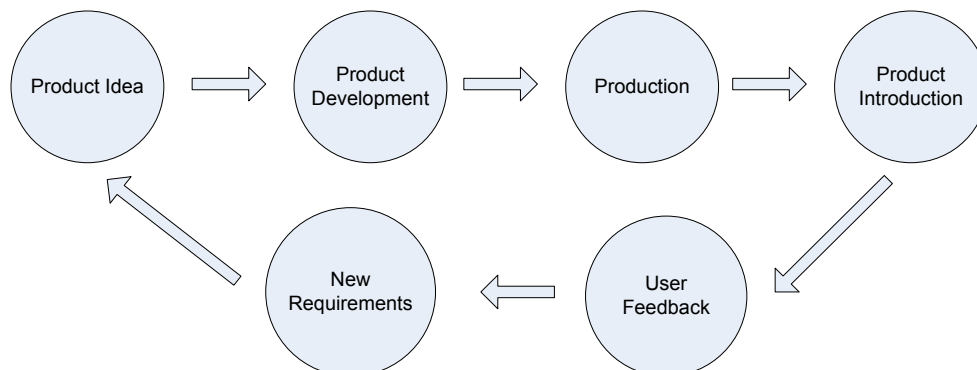


Figure 3. The feedback model of innovation

Since this model combines the previous two generations it does maintain some of the pitfalls of these previous models. Ulwick [23] stated that in this model there exists a culture of “me-too” such as when a automobile company delivers a new vehicle to market and consumers begin to request additional features to the vehicle that are simply duplicates of those found in vehicles of other manufactures. While this may have the tendency to improve the vehicle overall it is certainly not innovation.

The Strategic Model of Innovation

Although there is not agreement in the literature that this is a separate model, the fourth generation model consists of the model of innovation lining up with company strategy. Berkhout *et al.* [17] stated that this model and the third generation model are actually the same, due to the fact that the innovation chain is linked to the company strategy. Alternatively, Zhang *et al.* [21] stated that this model is the functional integration innovation process theory in which a parallel approach is functionally integrated around a core project. This hypothesis has been derived from the observation of the Japanese automobile and electronics industry. While end user requirements come into play as in the third generation models, this is only the case should the innovation line up with the strategic plan of the company as a whole, thus eliminating spawning products outside of the scope of the company plan.

The Networked Model of Innovation

The fifth generation, and according to the literature the current generation, of innovation models is based upon collaboration and networking. The term *Open Innovation* [20] is often used synonymously with this model. In this model, networks of innovation are created in which extra-enterprise and cross-discipline organizations join together to innovate. Collins [20] utilized the funnel analogy once again to discuss this model, but stated that the funnel is now porous. Companies still manage and assess their innovation processes in a series of phases, but ideas can come from external as well as internal sources and can enter the innovation process at any stage.

One important side effect of this model of innovation is that product innovation now often spurs service innovation as well. Products are increasingly accompanied by services and advanced service provision is increasingly facilitated by technological products [17]. Another exciting side effect of this model is that valuable ideas can leave the corporate innovation process at any stage, by being sold, spun off or licensed to another company [20]. Hence good ideas that do not necessarily line up with a company's strategic plan are not wasted.

Berkhout *et al.* [17] offered another view of this model with the Cyclic Innovation Model. The Cyclic Innovation Model (CIM) describes the generic innovation process by a '*cycle of change*' that links changes in scientific insights, changes in technological capabilities, changes in product design and manufacturing, and changes in market demand. In this model small incremental innovations may take place that do not involve

all the links in the cycle. For example an innovation in manufacturing technology may influence a change in market demand but not impact scientific insights.

Through the use of the fifth generation innovation model, corporations need to focus as much attention on outside influence on their current innovation work as they do on their own private research and development. Collaboration now forms a large sector of the mainline innovation process. This does present challenges when various innovating entities interact. Different company cultures and strategies may have an impact on the innovation process. As Collins [20] stated, Microsoft regards introductions as business experiments, while one of its corporate partners in the LUCIO project, KPN Mobile, requires introductions to be completely stable and reliable. This cultural difference may very well impact the innovation process in either a positive or negative manner.

Due to the networked nature of the fifth generation models, a number of data mining and computationally intelligent algorithms may be applied to the innovation process and the network selection in particular. Link analysis and social networking algorithms offer the ability for innovation minded companies to rapidly seek optimal partners in the innovation process. The use of computer algorithms is not limited to innovation partner selection as all models of innovation can make use of certain data mining and computationally intelligent algorithms. The social networking algorithms though are unique to the fifth generation models and should be exploited as deemed possible.

The Innovation Process

All of the innovations models discussed follow approximately the same generic innovation process. Kusiak [3] showed that innovation of all types perform five basic activities: search, evaluate, develop, refine and connect as shown in Fig. 4.




Figure 4. Generic Innovation Process

In the search phase of innovation, ideas are brought forth through various means: brainstorming, requirements gathering, market analysis, and so forth. Each idea is evaluated to determine its fitness to be further refined and developed. Should the idea be deemed fit, it is further refined to line up with requirements and company strategy. The refined idea is developed into a working product and is finally brought to market in the connect phase.

Stokic *et al.* [1] introduced the idea life cycle, stating that the idea will undergo a complete cycle in order to be collected, documented and used. This view of an idea life cycle is in line with the above stated generic innovation process. It should be noted that the idea life cycle in the fifth generation model may be more sporadic than in the other models as the idea itself may actually change company ownership a number of times

before undergoing the full innovation process. As Kusiak [3] stated, ideas can be generated locally within an organizational unit, or can be obtained from external sources.

As previously mentioned, data mining and computationally intelligent algorithms can play a large role in the innovation process. Kusiak [3] showed how evolutionary computation and data mining classification can be utilized to aid in the innovation solution. Even the use of TRIZ may benefit from searching the fitness landscape to determine the most appropriate engineering principle to apply and the optimal manner to abate contradictions with given inventive principles.

Many various frameworks for innovation have been proposed in the literature. Tung *et al.* [26] proposed an intelligent design framework for service innovation. Chen *et al.* [27] introduced a framework for managing service innovation as well. Xi-Song *et al.* [28] developed a framework for the detection of key elements in innovation as it relates to the defense science and technology industry. Felfernig *et al.* [29] developed toolkits for open innovation as it relates to e-government.

Conclusions

This chapter has presented a historic overview of various innovation methodologies and generations. From the inception of TRIZ to the Open Innovation of today, various morphologies have taken place in the perceived process of individual innovation creation. While none of these methodologies consider innovation as a system they have added to a rich corpus of information with which to understand innovative product generation.

CHAPTER 3. MINING THE REQUIREMENTS OF INNOVATION

One of the keys to understanding innovation as a system is to understand the requirements that drive that system. Individual innovative products and services are deemed as such only through market acceptance. Without the acceptance of the market a product that could have been an innovation is relegated to simply an invention which never gained acceptance. Therefore, a key to ensuring that an innovative product becomes deemed innovation is to fully understand the market requirements for the domain of that product.

In this chapter a methodology is illustrated for gathering the market requirements for a given product domain. This chapter is an aside from the rest of this thesis in that it centers on individual innovation instead of the system of innovation, however the requirements gathered utilizing this methodology are utilized in the construction of the principles of the system of innovation.

Determining market acceptability of a new concept during its early development stage is difficult and time consuming. There is no well established methodology or tool that would guarantee such success. Ulwick [23] stated that the traditional approach of asking customers for solutions tends to undermine the innovation process. This is due to the limited frame of reference of most customers. Yet, corporations spend millions on increasing the market acceptance of their products and services. While Ulwick [23] does have a justifiable case for limiting the utilization of customer input, there is an obvious advantage to ensuring that a product or service is more market worthy. In contrast to

Ulwick [23], von Hippel [7] argued that research in innovation suggests involving customers. Blazevic [30] suggested that electronic interaction with customers allows for a mutual understanding during the innovation process.

Focus groups and market research are often limited to a select group chosen either randomly or with segmentation in mind. This process is sub-optimal and may offer a false sense of security when releasing a product or service due to the limited size of the sampling. As in data mining, the larger the sample size for testing the better the model will be. This chapter extends the scope of the focus groups and marketing research tasks to the entire World Wide Web.

Mining User and Expert Reviews from the Web

It has become common practice on the World Wide Web for commerce sites to host user reviews. The practice is so engrained in the web that an entire field of study known as Collaborative Intelligence with sub-fields of study such as Collaborative Filtering [31] has been established. Collaborative Intelligence involves the collective reasoning of multiple users to achieve some goal. Netflix, an internet commerce site aimed at movie rentals, utilizes Collaborative Filtering to offer recommendations to its users. Netflix is so involved in this type of Collaborative Intelligence that they have even offered up a substantial prize for any team which develops an algorithm that is at least 10% more efficient than the current [32]. From such user reviews much can be gleaned to assist in the formulation of requirements for innovation.

User reviews often contain statements of opinion that indicate the strengths and weaknesses of a given product. When combined with user reviews for products of similar type a collective interpretation of the attributes that make a product a success or a failure can be obtained. Figures 5 and 6 illustrate this point for two different portable music players. Note that while the devices may be different, one being an MP3 player and the other a walkman style radio, they share similar review attributes including the need for good sound quality. It is these reviews that assist in forming the requirements for innovation.

Pros: Good sound quality, Can record normal conversation. The menu system is easy to use.

Cons: The case scratches easily, the ear buds are uncomfortable.

Pros: Stylish case and screen. It is nice to be able to customize the screen to my desires.

Cons: The battery life is a bit short.

Figure 5. Sample MP3 player review (Authors construct based on multiple review sites)

Pros: I liked the sound quality that this radio puts out.

Cons: Looses stations easily.

Pros: I dropped this while riding bike and it still plays like a charm.

Cons: None that I can see.

Figure 6. Sample Walkman review (Authors construct based on multiple review sites)

To extract expert and user reviews of existing products and services from the web it is necessary to discover the web pages which contain such reviews. There are two main methodologies for automated searching of the web. Standard, or unfocused, crawling of the web is intended to fetch a page, parse out-links, and repeat [30]. Standard crawling does not consider a specific topic of inquiry; rather its job is to index all pages available on the internet. Focused crawling seeks only those pages related to a specific query string.

Successful standard crawling requires massive hardware and bandwidth. This drawback prevents most corporations from performing this type of crawl internally. Focused crawling requires far less hardware and bandwidth but does require some sophistication of algorithms to weed out the undesirable links as they relate to the given query. Often data and text mining algorithms can be used to develop a classifier for organizing reviews presented on web pages. The classifiers are not always necessary especially when a specific domain of product types is to be utilized. In the case of a specific domain, a heuristic may be utilized as replacement for the classifier. Experimental results though have shown that the classifier system is more accurate at discovering correct pages while the heuristic discovers a greater number of correct pages faster.

An experiment was performed to test a classifier against the heuristic for the domain of MP3 players. The two algorithms were run for one hour and were set to discover the user reviews related to the domain. Figure 7 below shows the results of this

search. Figure 7 depicts the running of the heuristic based focused crawl and the classifier based one. Each chart shows the total pages that were discovered during the run versus the total number of pages containing actual reviews. The charts depict the page count on the y-axis and the run time of the crawl on the x-axis. The charts illuminate the fact that the classifier based crawl produces a larger percentage of correct pages while the heuristic based crawl discovers the correct pages faster but with lower accuracy.

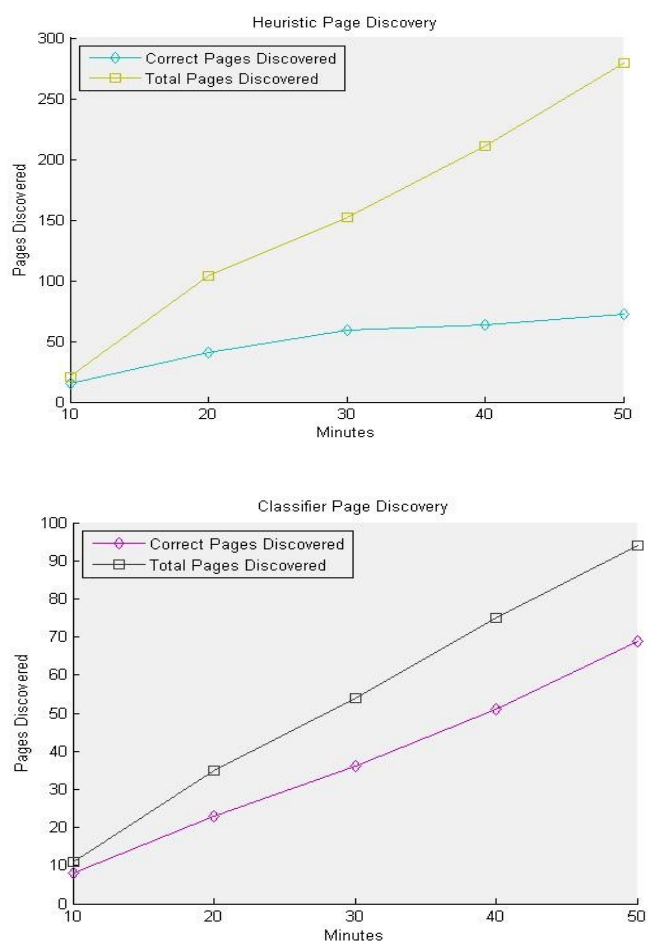


Figure 7. Heuristic versus classifier results

Fig. 8 directly compares the correct results of the two methodologies clearly showing that a domain specific heuristic is sufficient for rapid review detection. The two crawls were ran simultaneously and the results of this experiment are shown in Fig. 8.

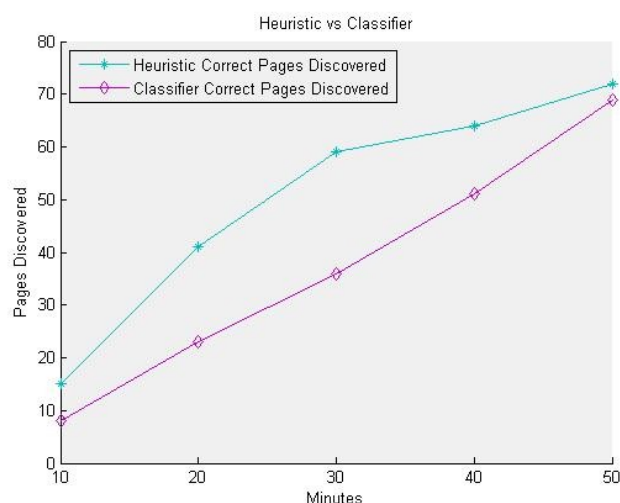


Figure 8. Comparison of heuristic to classifier

Once the reviews are discovered on the web they are parsed using standard text mining methods and stored in a database of reviews. The database is structured similar to a transactional database in a supermarket with one column for each review attribute (i.e., good sound quality, rugged case, and so on). This database, once populated with the other types of requirements sources discussed below, will be utilized to create a frequent itemset of requirements for the product type being considered.

Population of Transactional Database

Population of the transactional database from the reviews that were discovered in the web crawl is a twofold process. The first step is to filter out those reviews that offer little or no value to the requirements gathering and determination. A simple classification system, discussed below, is used to remove the unwanted reviews. The second step is to perform a semantic analysis of each review to segment out the individual requirements that form the attributes of the database. This section presents the methodology that was used to perform both of these operations.

To filter the unwanted reviews from the dataset a simple decision tree classifier was built utilizing a training set of data. The training set of data was cauled from the dataset and a class label of good review or MP (moaner/praiser) review was assigned through a manual inspection. It was discovered that this decision tree is actually extremely short with over 75% of the data being properly classified by the first node. The feature space of the decision tree consists of the words of the reviews with stop words such as “I”, “the”, “and”, and so forth removed. Each word was given an integer count associated with the number of times the words occurred in the corpus of the training set.

The set of nodes that made up the decision tree are, in order of importance, “none”, “nothing”, “awesome”, “awful”, “junk” and “kidding”. As previously stated, the first node correctly classified the moaner and praiser reviews 75% of the time. This should not be considered surprising when considering a praiser will most commonly have nothing negative to say about the product as a moaner will have nothing good to say about the product. The reviews are all formatted to the style of pros and cons through a

simple parsing manipulation. Thus, reviews of the type “Pros: I just love this MP3 player, Cons: none” were then easily filtered with the decision tree algorithm (a widely used data mining algorithm).

The task of breaking apart the reviews into individual requirements or attributes for the database is a matter of utilizing Part of Speech (POS) tagging. To accomplish this task the reviews are broken apart by pros and cons individually. Each segment is then further broken apart through natural stops such as commas or periods. Finally, those low level segments are tagged with their POS tags. For the task of POS tagging the WordNet API [33] is utilized. This Application Protocol Interface (API) takes a string as a parameter and returns the POS tag for that string. The WordNet API is a freely available software library that may be integrated into most higher-level programming languages.

Following the lead of Hu *et al.* [34], only those segments which contain nouns or noun phrases are utilized to generate attributes for the database. This is due to the fact that other components of the segment are unlikely to contain product features. Following this vein, a segment such as “Good Battery Life” would be captured as an attribute to store. Prior to the POS tagging, stop words are removed utilizing a standard stop word list. Con type reviews are changed to pro type reviews so that the attributes all reflect a desire on behalf of the consumer. Thus a review such as “The battery life is too short” under goes a transformation to “Long Battery Life” which becomes a requirement in the mining task.

Some amount of fuzzy matching is also involved in the attribute population. The system must have the ability to correlate, as effectively as possible, attributes that mean

the same thing but are expressed with different words. For this task the WordNet API [33] is again utilized. This time the API is used to determine the definition of words. This action is performed on the nouns rather than the adjectives. Through a list of definitions it is easy to determine the similarity of words. Words such as “Display” and “Screen” are found to contain the same root meanings and thus are correlated to the same word. The word that is chosen to be the attribute is the one first encountered. The fuzzy matching is also used to account for word similarities such as “auto focus” and “auto-focus”.

In some cases the noun may be separated by one or more words from the adjective that describes it. Such is the case of “The battery life overall is really good”. In this case a simple heuristic is used to mate the adjective with the noun. Simply, the noun closest to the adjective in the same sentence is used. This seems to account for most of the cases and little risk is induced when using this heuristic.

Once the reviews are segmented and the attributes discovered a transactional database is created. Each row of the database is formed from exactly one review. Each column of the database is formed from exactly one attribute that was discovered in the dataset. Each column is an integer data type and represents the number of times that the given review contained the given attribute. To allow for dynamic running of this system, columns may be added to the database as new attributes are encountered. All new columns are given a default value of 0 as the attribute was not found in the existing reviews that make up the data rows.

It is from the transactional database that the mining of the frequent requirements for innovation takes place. Each review represents the desires of individual consumers.

As such, it is important to ensure that any new innovation meet at least a good amount of these requirements.

Patent Database Mining

While user and expert reviews offer a wealth of information concerning the needs and desires of the market they are not the only source of requirements for innovation. Patent databases are also utilized to formulate requirements in a data-driven innovation approach. Patent databases offer complete and summary descriptions of inventions that have been previously envisioned. These text documents can offer a great deal of evidence for the trends in current innovation. Shih *et al.* [35] have devised a patent trend monitoring system that autonomously searches for patents, dissects them through semantic analysis and reports them in a comparison fashion to show recent trending in specific technological areas.

Patent databases are often more effective or requirements gathering than publications and thesis information. Wen [36] stated that patent gazettes reveal over 90% of research results for the patents, while more than 80% of that information is not enclosed in academic theses and publications. This wealth of information must be tapped in order to formulate true requirements for innovation.

In addition to the gathering of requirements, this research information assists in the ideation phase of innovation. To that end, Wen *et al.* [37] researched utilization of patents for the invigoration of the creativity processes. Fig. 9 illustrates a simple advancement that was discovered in a patent database. The use of extending arms to aid

in stability of the Christmas tree stand could be utilized in multiple areas during the ideation phase of an innovation that would require balance and stability.




Figure 9. Christmas tree stand patent. U.S. Patent Number 5492301

Mining web based patent databases is similar to the mining of user reviews given that both may be viewed as text documents with html markup added. A classifier is not necessarily needed to discover which documents are patents as a simple heuristic may suffice to mine sites such as Google Patents [38] due to the domain specificity of these sites. Similar to review mining, the utilization of patents requires semantic analysis to determine which text represents a possible requirement for innovation.

Determining which text represents a possible requirement can be more challenging in patent documents than in reviews due to the amount of text present in the documents. To disseminate the requirements in the patent the text is segmented naturally at the sentence breaks. Each sentence is then run through a trained classifier to determine

its relevance to requirements. This classifier, being previously trained by domain experts, is utilized to automatically sort the sentences, or other segmentations, into categories of relevance. Once sorted, each segment is then dissected to extract the requirements for innovation.

Requirements discovered in the patent documents are stored in the same transactional database as the review requirements. Storing the requirements in one transactional database simplifies the frequent itemset generation. Additionally, the user may choose to place these requirements in a creativity database for use during the ideation phase of an innovation project.

Mining Frequent Requirements

Once the requirements gathering process is complete for a given period of time, the task of determining requirements pertinent to making the innovation more market acceptable emerges. Within the transactional database there will be various requirements that were gathered that may not be good indicators of market success since they may have only been mentioned once or twice in the entire set of reviews or patents. To determine reviews that may increase market acceptance of an innovation, an apriori type market basket analysis is utilized.

The transactional database described above forms a sparse dataset from which the frequent requirements are mined. The dataset is sparse in the sense that each entry contains a very small number of requirements attributes. Scarcity is not a concern of the

frequent itemset generation algorithm utilized here. Due to this scarcity an apriori market basket analysis is performed to generate the frequent requirements.

To mine the frequent requirements a metric is defined for support. Support, similar to that defined by Tan *et al.* [39], determines how often an itemset is applicable to a given dataset. Given an itemset $X = \{x_1, x_2, \dots, x_n\}$ and a dataset D with number of transactions $|D|$ the support, σ , of the given itemset is defined to be the number of transactions k in D that contain the itemset divided by the cardinality of the dataset as in Eq. (2).

$$\sigma = \frac{1}{|D|} \sum_{k \in D} X_k$$

Equation 2. Support

For an itemset to be considered frequent the support of that itemset must be greater than some threshold τ which is a user defined parameter. The use of the minimum support threshold ensures the generated itemsets can truly be considered a requirement and not simply an anomalous entry.

Generation of the frequent itemsets is performed in accordance with the standard apriori algorithm [40], Candidate 1 item itemsets are generated for each of the attributes in the transactional dataset. Each candidate is checked to see if its support is above the minimum support threshold τ . Those candidates whose support is less than τ are discarded as they do not hold to the apriori property. The apriori property states that all nonempty subsets of a frequent itemset must also be frequent. Once the frequent 1 item itemsets are generated they are combined to form frequent 2 item itemsets. These

candidate itemsets are then checked for support. This process of candidate creation and checking is performed iteratively until itemsets available for combination no longer exist.

The generation of the frequent itemsets from the transaction database forms the requirements for innovation. Starting with the largest frequent itemset, one can determine the requirements that are described by the reviews simply by observing the items in the itemset. Thus, an itemset generated from the review of MP3 players may be represented as:

Long battery life, Good sound quality, tough case, and accepts standard accessories.

While frequent itemset generation takes a large step forward in determining the requirements for innovation they do not offer complete understanding of how the requirements interact. In example, there may exist five sets of frequent requirements of length 5 in the transactional database as determined by the apriori mining above. The requirements *Long battery life* and *Good sound quality* may be found in each of the itemsets. This indicates that these requirements are special in relation to the others in the itemsets. This quality of being special though is not described by the frequent itemset generation.

In order to understand the unique relationship each of the requirements has within the frequent itemsets the use of an AND-OR tree is considered. AND-OR trees have been utilized in areas from logic representation to web mining, see [41], [42] and [43]. The

construction of the AND-OR tree is taken a bit differently though. The purpose of this tree is to aid in determining the significance of each attribute within the frequent itemset.

The AND-OR tree is constructed iteratively on each set of itemsets in descending cardinality. A root node, with name requirements, is generated. Starting with the set of itemsets whose cardinality, number of items in the itemset, is the largest and whose member sets count is greater than 1, a node with the name of the cardinality of the itemsets is appended to the root node. The process of relationship determination is then performed for each attribute in the itemsets.

Relationship determination is performed by iterative search through the set of itemsets to find the frequent attributes in the set. A measure of support is utilized here such that, for each attribute to be considered frequent it must appear in at least ρ number of itemsets in the set of itemsets. Each frequent attribute forms a node that is appended to the cardinality node in the AND-OR tree. These nodes are then joined with arcs to represent AND conditions. The remaining attributes which are not frequent in the set of itemsets form child nodes to the last frequent node in the current branch of the tree. These nodes form the OR conditions of the relationship. This process is continued iteratively for each successive smaller set of itemsets until a stopping criterion (user determined) is met or the remaining itemsets are of cardinality 1.

For example, given a set of three frequent itemsets of cardinality 3 illustrated in Figure 10, the AND-OR tree representing the relationship of these itemsets, with $\rho = 2$ is presented in Figure 11.

Long Battery Life, Good Sound Quality, Clear Display

Tough Case, Long Battery Life, Good Accessories

Good Sound Quality, FM Tuner, Long Battery Life

Figure 10. Three frequent itemsets of cardinality 3

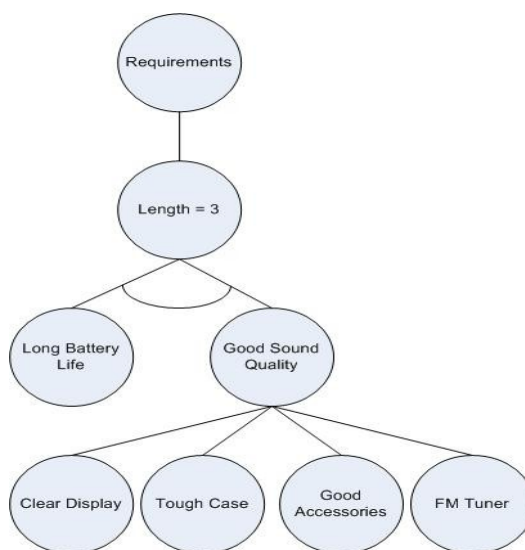


Figure 11. Partial AND-OR tree for requirements in Fig. 10. (only the branch for cardinality of 3 is shown. The full AND-OR tree would have branches for other cardinalities as well.)

Thus the AND-OR tree in Fig. 11 defines the relationship of *Long Batter Life* and *Good Sound Quality* as the most important requirements within the set of frequent itemsets of cardinality 3. Additionally, the children of *Good Sound Quality* determine less

frequent requirements of importance to the innovation process. As one can see, the AND-OR tree visualizes the frequent itemsets and makes it easy to determine the most important factors when considering an incremental innovation.

Conclusions

Incremental innovation is an evolutionary process. The requirements for an idea to be deemed innovation through market acceptance are changing in time. To keep up with these changes and to aid in the market acceptance of an idea, the mining of current requirements for innovation is crucial. Knowing the current requirements for innovation in a given domain is helpful in ensuring the idea will not fail as well as in the ideation process of innovation.

This chapter presented a methodology for mining the frequent requirements of innovation in an automated fashion. Once the requirements have been gathered the process of dissection and determination of the frequent requirements is achieved by data mining algorithms. The output of the data mining process, the AND-OR tree, can be utilized to evaluate and evolve current innovations. Through the continued use of the methodology presented in this paper, the changes in requirements may be monitored and a company utilizing the methodology may be better positioned for market success.

CHAPTER 4. THE NK LANDSCAPE OF INNOVATION

Examining innovation through the use of the principles of biological systems, namely evolutionary computation, is not an entirely new concept. Fleming [15] studied technology as a complex adaptive system which utilized concepts of evolvability. Brabazon *et al.* [44] modeled product innovation with the use of agent-based modeling techniques. Windrum *et al.* [45] discussed the product life cycle as utilizing the complex adaptive system components of emergence and niche creation. Rogers *et al.* [46] equated complex adaptive systems to the diffusion of innovation and its adoption by entities external to the supplying firm.

The studies above do not account for the nature of the system innovation as a complex adaptive system traversing rugged, ever-changing landscapes. Kauffman *et al.* [47] discussed the landscape of technology and offered recipes for product configuration that mimicked Kauffman's NK Landscape [48]; however, they did not include a discussion of the system of innovation in whole or part, but rather looked towards product manipulation.

The NK landscape proposed by Kauffman [48] considers the behavior of systems comprised of N elements, each of which is interconnected with K of these N elements [44]. The degree of interconnectedness of the N elements dictates the smoothness or roughness of the landscape that is traversed by the system under consideration. Other research papers such as [49], [50], [51] and [52] have added to the work of Kauffman and have solidified the NK landscape as a viable means of representing systems that utilize the biological principles of evolution and Darwinian selection.

This chapter presents a novel study of the landscapes that are traversed by the system of innovation as a complex adaptive system. The NK landscape of Kauffman [48] is used to evaluate complexity that is innate to the innovation system. Patents are utilized as surrogate representations of innovation to reconstruct historic landscapes that have been encountered by the system of innovation and to show that the innovation system is poised at the edge of chaos due to the interconnectedness inherent in the nature of the system of innovation. Patents clearly do not represent every innovative product that was ever created, nor do they imply that every patented product was a true innovation; however they can be utilized in the general purpose here to show the interconnectivity that does exist in the system of innovation.

NK Landscapes

The NK Landscape model presented in [48] considers a population of genotypes containing $N = \{1, \dots, n\}$ genes. Each genotype takes on a specific value (allele) at each of the N loci from a finite set of values A . The simplest genotype is the binary genotype in which $A = \{0, 1\}$. Each of the N components of Kauffman's model interacts with K other components of the genotype, where $K = \{0, 1, \dots, N - 1\}$. The fitness (strength) of each genotype is given as a function of N and K . For $K = 0$, all N loci are independent of each other and contribute to the fitness of the overall genotype independently, as shown in Eq. (3).

$$F = \frac{1}{N} \sum_{i=1}^N f(n_i)$$

Equation 3. Fitness

where $f(n_i)$ is the fitness contribution of the i^{th} gene, and F is the fitness of the overall genotype.

Kauffman [48] stated that in a system with N genes, the fitness contribution of an allele of a gene often depends on the alleles of some of the remaining $N - 1$ genes. Such dependencies are called epistatic interactions. The number of epistatic interactions is represented by K in Kauffman's model. Given that $K > 0$, the overall fitness of each genotype in the NK model is given in Eq. (4).

$$F = \frac{1}{N} \sum_{i=1}^N w_i$$

Equation 4. Genotype Fitness

where w_i is a function of the allele at the i^{th} locus and the alleles of the K other loci that interact with the i^{th} locus. It has been stated that the value of w_i can be considered as independently distributed variables approaching a standard Gaussian distribution $U(0, 1)$ as the number of loci increases [50]. This follows closely from the central limit theorem. The value of w_i is not critical to the understanding of fitness landscapes. The important fact to understand is what happens to the landscapes as K increases.

For $K = 0$, the landscape that is described by the population of all genotypes is smooth with, simplistically, a single global optimum defined by the fittest genotype. A

smooth landscape is one in which genotypes near one another have nearly the same fitness [48]. A landscape in which $K = 0$ is easily seen as a rare occurrence, especially in complex adaptive systems.

As K increases, the landscape becomes more rugged due to the epistatic interactions of the loci of the genotypes. A maximally rugged landscape is one in which the fitness values are entirely uncorrelated [48]. Thus, given a particular genotype in the fitness landscape, a neighboring genotype's fitness cannot be determined due to its proximity to the given genotype. Additionally, as K increases, the number of local optima increases as well. A special case of this instance occurs when $K = N - 1$, which corresponds to a fully random landscape [48]. Kauffman stated that for landscapes where $K = N - 1$, the system approaches complete chaos. It is the landscapes with values of $0 < K < N - 1$ that are most interesting and present ordered complexity.

For $K = 0$, the landscape is defined by a small number of regularities. When the number of epistatic interactions, or K , increases, the number of regularities defining the landscape (complexity) increases. Regularities are a means of measuring the amount of complexity in a data set [53]. Another method to measure the complexity of a data set is to utilize the Kolmogorov complexity metric, which measures the randomness of the description of the set [54]. Either measurement will suffice to make the point that as K increases so does the complexity of the landscape.

Kauffman [48] stated that as K approaches $N - 1$, the height peaks of the local optima tend to decrease towards the mean of the population. Conversely, as K approaches zero, the peaks of the optima increase over the mean of the system. While it may not

always be possible to predict or ascertain the exact value of K in a complex adaptive system, it is possible to determine if the peaks of the optima are approaching the mean of the system, thereby proving that the system is degrading into a chaotic state.

The Landscape of the System of Innovation

Complex adaptive systems remain a topic without a solid accepted definition. Instead, complex adaptive systems (*cas*) are expected to meet a number of criteria. Holland [55] enumerated seven basic elements in two categories that must be met for a system to be considered a *cas*. The properties' category consists of the elements of aggregation, non-linearity, flows and diversity. The mechanisms' category consists of the elements of tagging, internal models, and building blocks. The system of innovation meets the criteria proposed by Holland [55]. The innovation system utilizes aggregation by accepting groupings of specific types of innovations, such as product innovations, service innovations, technological innovations, and so forth. The innovation system as a whole considers a space of all inventions that have ever been deemed market acceptable at a given point in time. Incremental evolution within this enormous space becomes difficult, as the local optima are not clearly defined. Rather, through the use of aggregation, individual innovations can be grouped to reduce the search space for incremental evolution and thus heighten the optimal peaks within the landscape.

Strogatz [56] described linear systems as the weighted sum of the individual components of that system, while non-linear systems utilize the property that the system is more than the sum of its parts. It is easy to see that the system of innovation is far more

than the sum of its parts. The individual innovations upon which a particular innovation is built do not completely describe the previous individual innovations. As a closely related principle, individual innovations rely on passing information from one to the next (flows). For a particular innovation to be deemed successful, it must pass market acceptability [57]. As such, individual innovations that are market worthy increase the information flow and thus increase the chances of success of the new inventive idea.

The system of innovation is diverse in the common understanding of the term, and utilizes niche creations. Considering the interconnectedness of the system of innovation, current innovations building upon previous ones, it is easy to see that an individual niche innovation is replaced prior to further evolution of the current innovation.

Finally, the system of innovation meets the required criteria by making use of tagging, internal models, and building blocks. It utilizes the same tagging mechanisms as other *cas* in that it allows users to filter through large amounts of data to discover individual innovations that may be aggregated into groups of similar items. The innovation system constructs internal models through its use of the exchange of genetic material in previous individual innovations. The internal model of the system of innovation operates such that previous individual inventions that have failed to meet market acceptance at the time are often excluded from the creation of future generations. Building blocks in *cas* are smaller elements that are previously defined and are combined to formulate a larger item. Such building blocks are used to create individual innovations as well as to form the system of innovation as a whole. In the case of the motor vehicle, for example, it can be seen that the previous individual innovations of the wheel, the

internal combustion engine, and the battery are building blocks that the system of innovation has utilized to form the vehicle.

Having shown that the system of innovation meets the criteria to be defined as a *cas*, it is implicitly understood that the innovation system generates a population of individual innovations of varying fitness, as demanded in part by the market. These individuals interact to generate new innovations that are composed of genetic material from previous generations. A fitness landscape is constantly changing due to market influences.

Patents serve well as surrogate representations of inventions. Fleming *et al.* [15] utilized patents to equate the diffusion of innovations to complex adaptive systems. Patents record the individual innovation as well as previous innovations (patents) that the current innovation is built upon. Additionally, patent databases contain a list of other inventions and individual innovations as references. It is well known that not every patent in the patent database records a true individual innovation as deemed so by the market. It is also known that not every true innovation is recorded in patent databases. However, for the purposes of the models being built here, patents serve well to represent the general concepts under consideration. To generate a historic landscape of innovation patent databases have been mined. A clustering algorithm is used to cluster patents into groups. To facilitate the clustering, each patent's abstract and title were used. Stop words, words common to all texts, such as "the", "in", and so forth, were removed from the patent abstracts and titles. A list of word counts was generated from the corpus of the patent abstracts and titles for each patent in the given time frame. To cluster the patents, a K -

means algorithm was used. The *K*-means algorithm utilizes a distance metric to group similar objects (similar content) [39].

Once the clusters were formulated, the patent interactions were utilized as a representation of fitness, as suggested above, for market acceptability. Each patent in the cluster then defines a point in the innovation landscape given in three dimensions. The *x*-axis represents the cluster of which the patent is a member. The *y*-axis represents the year in which the patent was issued, and the *z*-axis represents the number of patents that cite the given patent, thus representing its fitness. Fig. 12 illustrates the fitness landscape, as mined, for the years 1905-1910, while Fig. 13 presents a similar landscape for the years 1910-1915.

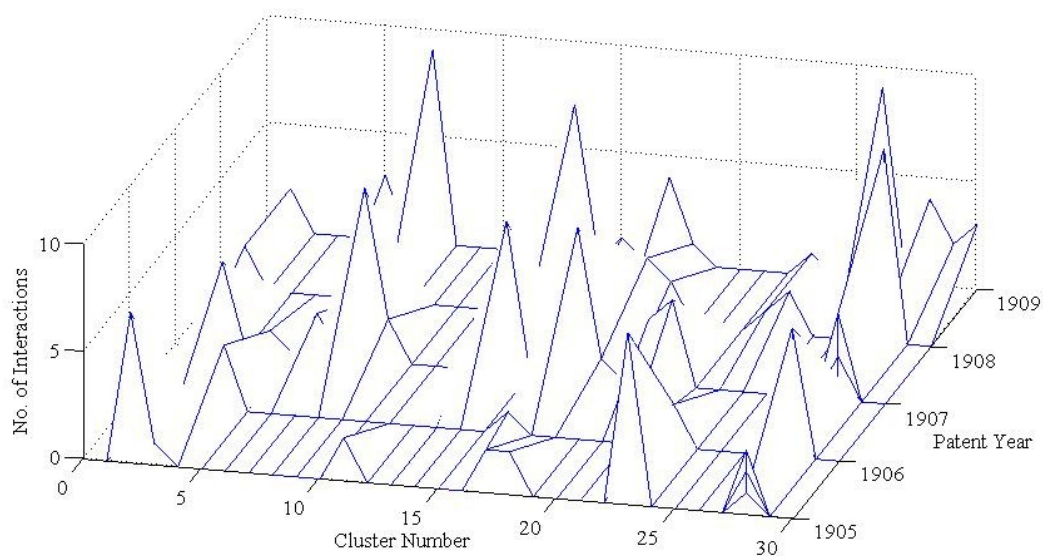


Figure 12. The fitness landscape of innovation for the years 1905-1910.

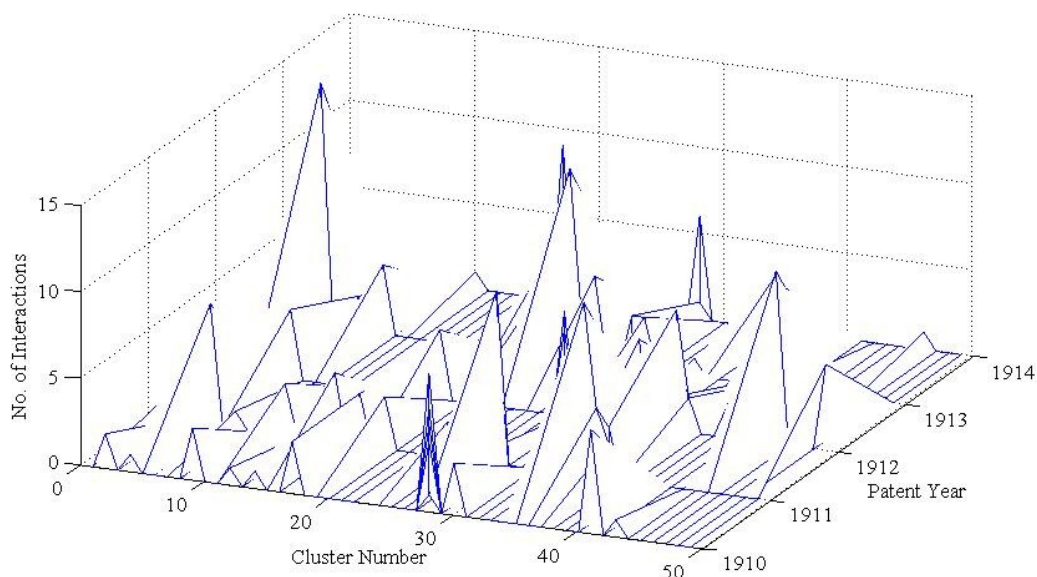


Figure 13. The fitness landscape for the years 1910-1915

As shown in Figs 12-13, the number of clusters increases greatly in the years 1910-1915 compared to the number of clusters in the years 1905-1910. In addition, the number of interactions has increased in the later years over the earlier years. It is from this historic recreation that a great deal can be learned about the evolution of the system of innovation and the landscapes which it traverses.

The innovation system landscapes presented illustrate a discovery of the historic landscapes that the system of innovation has traversed. From historic landscapes it is possible to prove that the system of innovation does not deteriorate to chaos, nor does it stay frozen with little or no epistatic interactions. Rather, as will be shown, the innovation

system forms an ordered complexity that can be utilized to predict the future landscapes of the system of innovation as well as extract specific properties.

Historic Landscapes of the System of Innovation

What information may be gleaned from the recreation of historic landscapes for the system of innovation? In the context of the NK model, historic landscapes can offer insights as to the number of epistatic interactions, the number and height of local optima, and the progress towards complexity found in the system of innovation. Though it may not be possible to calculate exactly the number of epistatic interactions, it is possible to establish bounds and to discover how these bounds change over time.

Fig. 14 illustrates the recreation of the fitness landscape for the system of innovation for the years 1800 to 1900. This landscape was created utilizing clustering interactions.

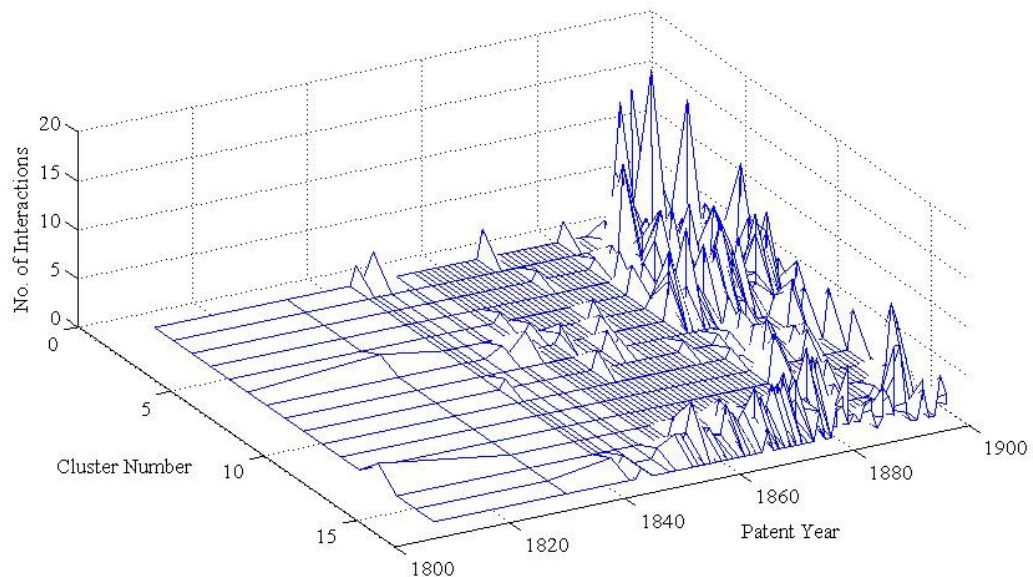


Figure 14. Fitness landscape for the system of innovation for the years 1800-1900

From Fig. 14 it can easily be seen that the system of innovation, in its surrogate form, was a sparse event during the early 1800s. Obviously, a good deal of this sparseness is attributable to the lack of patents during this time period, but for the purposes of this study, the recreation will suffice. Once patents began to be recorded in a meaningful manner, the epistatic interactions rapidly increased (see Fig. 14). It is this increase that intimates the nature of complexity in the system of innovation.

As previously stated, the epistatic interactions increase as the value K of the NK landscape increases. Thus, in comparison, K in 1900 is far larger than K in 1850. Does this increase in K indicate a breakdown of the system into a chaotic state? To determine the limit to which K has increased, it is beneficial to observe the trend of the height of the peaks within the landscape. If the height of the peaks trend towards the mean height of

the landscape, then the system is indicating a breakdown towards a chaotic state. Fig. 15 below describes the peaks and the means of the landscape shown in Fig. 14.




Figure 15. Peak fitness versus mean fitness for the innovation landscape of the years 1800 to 1900.

The solid line in Fig. 15 represents the peak fitness, normalized by the number of nonzero peaks, and the dashed line represents the mean fitness, normalized by the number of all points within the landscape, for the years 1800 to 1900. The heights of the optima in Fig. 15 do not trend towards the mean fitness of the landscape. This indicates that a breakdown of the system to a chaotic state does not take place. In fact, if the surrogate of patents is taken as the only indication, the exact opposite occurs within the system. From 1800 to approximately 1852, the value of K is exactly zero. After 1852, the value of K increases but does not increase drastically and is by no means near the limit of $K = N - 1$.

The historic landscape shown in Fig. 14 and the evaluation of the optima in Fig. 15 offer some insight into the system of innovation's complexity. It is necessary, though, to observe other time frames, particularly recent times, to ensure that the conclusions from the initial model will hold. Fig. 16 shows the recreated fitness landscape for the years 1900 to 1950. Obviously, there are many clusters and optima during this time, as would be expected given the surrogates of patents.




Figure 16. Historic fitness landscape of innovation for the years 1900 to 1950.

As shown in Fig. 16, a great deal more complexity has been introduced into the system of innovation. There are far more optima, indicating a greater amount of epistatic interactions. Additionally, the optima are well dispersed throughout the system and have much higher values than in the previous landscape of Fig. 14. Again, it is important to turn to the optima versus the landscape mean to determine if the system is breaking down

to a chaotic state, or if it is maintaining the ordered complexity that was hypothesized from Fig. 15.

Fig. 17 displays the mean fitness, normalized as before, and the fitness of the optima, also normalized as before. The mean fitness of the population in Fig. 17 has definitely increased. It can also be seen that the fitness of the optima has also increased. The important issue, though, is to determine if the fitness of the optima is trending towards the mean.




Figure 17. Peak fitness versus mean fitness for the innovation landscape of the years 1900 to 1950.

The fitness of the optima (see Fig. 17) is becoming closer to the mean fitness of the population, which is anticipated given the Central Limit Theorem and the large number of patents under consideration. This trend could indicate deterioration of the system to a chaotic state, or it could simply indicate that K , the number of epistatic

interactions, is increasing but not without a bound. To enhance the understanding of the complexity being imposed on the system of innovation, additional time periods are examined.

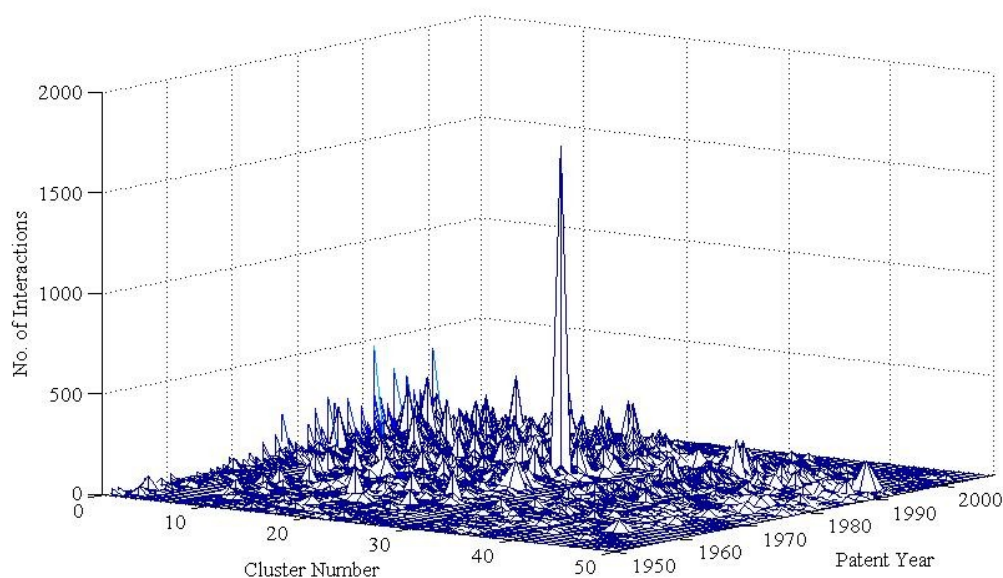


Figure 18. Historic fitness landscape for the years 1950 to 2008.

Fig. 18 reports the historic landscape for the years 1950 to the present. In this landscape, the complexity of the system innovation has increased, as well as the epistatic interactions. Previous interaction scales were in the order of 100, while the landscape in Fig. 18 shows interactions of many hundreds. Thus, if the system of innovation does indeed deteriorate into a chaotic system (as opposed to a complex system), the fitness of the optima of the population should see a marked trend towards the mean fitness of populations, as predicted by Kauffman [48].




Figure 19 . Peak fitness versus mean fitness for the innovation landscape of the years 1950 to the present

As can be seen from Fig. 19, the fitness of the optima does not approach the mean fitness of the population of all patents. In fact, on average, the distance between the fitness of the optima and that of the mean population varies only slightly from that of the previous time frame.

The complexity of the data set increases in relation to the number of generations. This increase in complexity is bounded as the optimum is approached. In fact, the epistatic interactions in the data set have decreased the distance between the optimum fitness and the mean fitness. However, this trend is tempered to a specific limit as the number of generations increase. As was shown in the graphs of the peak fitness versus the mean fitness of the population for the various time periods, there is a direct relation between the two fitness measures, as would be anticipated.

It has been found that as the number of generations increase, the function that represents the relation between the fitness of the optima and the mean fitness of the population approaches Eq. (5).

$$\lim_{t \rightarrow \infty} f(\mu_t) = \rho - \alpha(\rho_t)$$

Equation 5. Optima to Mean Relationship Function

where μ is the mean fitness of the population at time t and ρ is the fitness of the optima at time t . This function is shown in Fig. 20, which illustrates the fitness of the optima versus the mean fitness of the population of innovation for the years 1800 to 2008.




Figure 20. Relation of the fitness of the optima to the mean fitness of the population for the years 1800 to the present.

The limit of deterioration towards chaos that is approached by the fitness of the optima illustrates effectively that innovation is a complex adaptive system that lives on the edge of chaos, as noted in [48], [58] and [14]. Systems that reside at the edge of chaos exhibit the following property: Small fluctuations in the rules that define the system can perturb it so that it reaches stasis, or complete disorder, but it is normally resilient against such fluctuations. The system of innovation has proven to be just such a system, as can be seen from historic landscape recreations.

Wind Turbine Case Study

Wind turbines are an aggregation of individual innovations receiving widespread attention due to the demand for alternative energy. The concept of harnessing the power of wind is not unique to the 21st century or even the late 20th century. Wind power dates as far back as the ancient Egyptians. Therefore, harnessing the power of wind is a well-defined field which can serve as fertile ground for assessment of, and discovery within, its landscape. This section considers the historic landscape of the aggregated innovations of wind turbines through the use of the surrogates of patents. Due to the limited history of patents, the time frame considered in this section is that of the years 1900 to the present.

It is possible to create the landscape of a single type of innovation as opposed to the complete system of innovation that was discussed above. The landscape of a single type of innovation can be seen as two-dimensional due to the removal of other innovation aggregates. Above, the landscapes of innovation were depicted as three-dimensional. Here, landscapes are described on a two-dimensional graph, with the x axis representing

the patent year and the y axis representing the number of interactions. As such, Fig. 21 depicts the fitness landscape of wind turbines.




Figure 21. The fitness landscape for wind turbines for the years 1900 to the present.

The fitness landscape (Fig. 21) of wind turbines reached its maximum at about 1982. Of course, this is not to say that current research in the area of wind turbines is stagnating. Recall that the use of patents as surrogates invokes a delay, and there may be a large number of recent patent applications that have not made it into the patent databases. It is also important to note that there is a large gap in the historic landscape that has been removed, namely between the years 1932 and 1967.

Again, the historic recreation of the fitness landscape exposes the inherent complexity of even a single innovation aggregate. The number of epistatic interactions in a single innovation aggregate are maintained as in the whole of the innovation system.

This can be further seen in the relation of the fitness of the optima versus the mean fitness of the populations, as in Fig. 22. The fitness of the optima exhibits the same relation to the mean fitness of the population as that of the innovation system as a whole. Further, it can be seen in Fig.23 that the single innovation aggregate of wind energy does approach the same limit of epistatic interactions as the system of innovation as a whole.




Figure 22. Fitness of the optima versus fitness of the mean of the population for wind turbines for the years 1900 to the present.

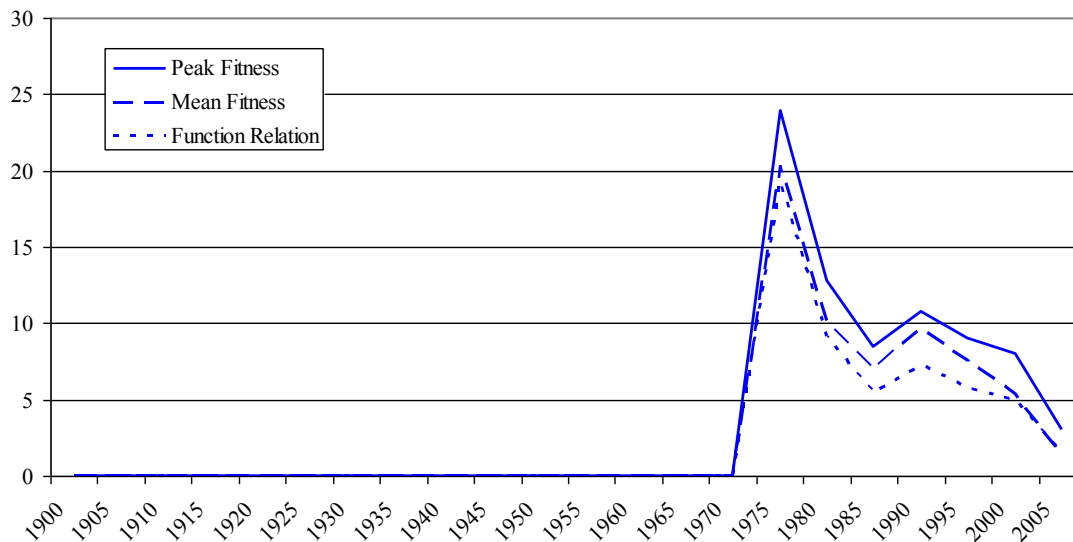


Figure 23. The relation of the fitness of the optima to the mean fitness of the population for wind turbines for the years 1900 to the present.

Some latency may be noticed in the graphs of the fitness values in Fig. 21 when compared to the graph of the fitness landscape in Fig. 22. This is due to the calculations used to formulate the graphs of the fitness values in Fig. 22, and their effect on the scale of the system. Regardless of this latency, it can be seen that this single aggregate of innovation maintains the properties inherent in the global system of innovation.

With the understanding that single aggregates of innovation maintain the same properties as the system as a whole, with perhaps differing K values, it is then proved that components of the innovation system are themselves adaptive in their creation, as one would expect. This then allows for a deeper study of the types of evolutionary operations at work in the innovation system.

It is initially conjectured that the evolutionary operator of mutation is extremely common, and that the operator of recombination is actually a rather rare occurrence. If this conjecture is proved to be true, then it may be seen that mutation breeds incremental innovations, while recombination breeds disruptive innovations [59]. Validation of this conjecture through empirical evidence can be accomplished.

The single innovation aggregate of wind turbines can prove a fruitful case study to verify that incremental innovation most often occurs through mutation rather than recombination. When using patents as surrogates, it is possible to view which other patents the current set cited as the foundation for its own validity. If mutation is the main evolutionary operator, each of the patents for a specific innovation aggregate should cite similar other patents. Otherwise, as is intuitively deduced, recombination would be the main evolutionary operator.

An experiment was conducted as described above to determine if mutation is the main operator, and that recombination is a rare event in the evolutionary system of innovation. For this experiment, 1905 patents on the single topic of wind turbines were considered. The patents which each of these 1905 patents cite are recorded, and a histogram of the counts for each cited patent is assembled. The results are presented in Fig. 24.




Figure 24. The number of patents which cite a given patent as basis for their validity

Unfortunately, the graph in Fig. 24 illustrates the exact opposite of the conjecture. It shows a large number of patents that are cited by the 1905 original wind turbine patents. Furthermore, it is easily seen that each of the cited patents contains few repetitive citations, eight at the maximum. The same experiment with the identical 1905 original wind turbine patents has been performed, with an added constraint of only recording those cited patents that related to wind power themselves. Fig. 25 presents the results of this second experiment.

As can be seen in the second experiment, even the constraint of only considering cited patents which relate to wind power does not diminish the randomness of the patent citations. From these experiments it is easily seen that, at least for the topic of wind

turbines, mutation is utilized in conjunction with recombination to evolve the innovation sub-system.




Figure 25. The number of patents which cite a given wind power related patent as basis for their validity

The final question to be answered in relation to the evolutionary processes is whether or not recombination utilizes a great deal of non-domain related individual innovations. In other words, are the cited patents generally from the domain of wind power in this case, or are they from other non-wind related domains? The answer to this question is that of the patents cited by the original 1905 wind turbine patents, 157 of them were in the domain of wind power while 210 of them were in other (non-wind power related) domains. It was found, however, that the number of repetitive citations was lower

for the non-wind power related domains than for the wind power related domains, as can be seen by comparing Fig. 25 and Fig. 26.




Figure 26. The number of patents which cite a given non-wind power related patent as a basis of their validity

Therefore, it is concluded that the system of innovation traverses the complete fitness landscape, utilizing both the mutation and recombination operators of evolution to generate the next generations of the population. It is noted, though, that the traversal of the landscape is made in small hops rather than large jumps, as the mutation operator heavily influences the system to achieve fitter individual innovations given the market need and acceptance.

Turbine Blade Case Study

Presented is a brief case study of the innovation of turbine blades. A number of patents, not limited to wind energy, in the patent databases utilize turbine blades. Thus, this section reviews a landscape similar to those found above. The data used here came from the same patent databases as the previous case study. In this data, a number of clusters were found which related to turbine blades. Fig. 27 illustrates the historic landscape for turbine blades. While the time period considered is from 1900 to the present, it was found that turbine blade patents were not significant in the patent database until the 1940s.




Figure 27. Historic fitness landscape of a turbine blade for the years 1940 to present

As shown in Fig. 27, even a small subset of individual innovations, turbine blades in this case, represent a fair amount of complexity and epistatic interaction. To validate that the epistatic interactions do not plunge the system into a chaotic state, the fitness of the optima is compared to the mean fitness of the population, as shown in Fig. 28.




Figure 28. Peak fitness versus mean fitness of the population of patents considering turbine blades for the years 1900 to present

The fitness of the optima in Fig. 28 does not trend towards the mean fitness of the populations. This indicates that subsets of the complex adaptive system of innovation approach the edge of chaos. It may be that the subsets of individual innovation are less stable due to the perturbations of initial conditions, and thus may be diverted to stasis or chaos, but the evidence for this is not conclusive.

An experiment was conducted to determine whether the mutation operator was singularly applied. Not surprisingly it was found that the subset of individual innovations related to turbine blades exhibited the same properties of evolution as the subset of innovations regarding wind energy. In this experiment, 1128 patents were used. Fig. 29 illustrates a histogram of the counts for each cited patent assembled.




Figure 29. The number of turbine blade patents citing a given patent as a basis for their validity

As illustrated in Fig. 29, a large number of patents are cited for the small subset of individual innovations related to turbine blades. While the number of patents cited by this subset is significantly less than the number cited by wind turbine related patents, the premise, that mutation is not the only evolutionary operator at work, holds.

The subset of individual innovations utilized in this section has sufficiently illustrated that there are common principles that hold true throughout all of the system of innovation. These principles state that the innovation system is a complex system that makes use of adaptation to ensure market acceptability. The system of innovation is a system that is poised at the edge of chaos. The latter principle has been proven to be correct on both a global innovation scale and on subsets of innovation. Finally, the system of innovation utilizes the evolutionary operators of mutation and recombination to traverse the fitness landscapes defined, in part, by the global marketplace.

Epistatic Interactions in Simple Systems

To further assist in understanding how epistatic interactions influence complex systems, an experiment was performed with simple networks where agents' values relied entirely on the values of their neighbors. This section describes this experiment and utilizes the results to offer guidance for the evolution of future innovations.

For this experiment, bit strings of length N (= the number of agents) were considered. Each bit in the string is considered an agent and is randomly selected from a uniform distribution to be an AND agent, OR agent, or XOR agent. Each agent in the string is initially assigned a value 0 or 1 based on a random selection from a uniform distribution. Once the initial values are assigned, iterative updating occurs based upon the agent's rule and the k nearest neighbors. Thus, if $k = 2$ and the agent's rule is AND, then the agent's value will be 1, if, and only if, the neighbors immediately to the left and right

of the agent both have a value of 1; otherwise the agent's value will be 0. The bit strings are considered to be connected at the ends such that the neighbor immediately to the left of the first agent is the last agent in the string. The value update has been accomplished in 1000 iterations for each run of the algorithm.

It was found that the bit strings settled into a periodic loop of some length l . The rate at which the bit strings settled into this period and the length of the cycle of the period were found to vary in direct relation to the number of epistatic interactions that were involved, as well as the population size (length of the bit string). To ensure accurate results, the same experiment was performed for different values of $k = \{2, 4, 6, 8, 10, 12\}$ and for population sizes of 50, 100, 150, 200, 500, and 1000 agents. Each experiment was performed 1000 times, and the mean starting iteration of the period, μ and the cycle length of the period, ρ , were calculated. Additional 1000 runs of 1000 iterations each were repeated four times, and the results were averaged.

Table 1 illustrates the results achieved through this experiment. They indicate that larger values of k decrease the mean starting iteration of the period μ (the iteration in which the cycle's first string appeared). Also decreased was the cycle length of the period ρ for smaller population sizes. In general, a trend towards this nature was detected regardless of population size, albeit some increase in μ and ρ were detected for the smaller values of k for larger population sizes.

Table 1 . The mean starting iteration and mean cycle length of the system of simple boolean agents.

Population Size	$k=2$		$k=4$		$k=6$		$k=8$		$k=10$		$k=12$	
	μ	ρ	μ	ρ	μ	ρ	μ	ρ	μ	ρ	μ	ρ
50	27.0695	15.416	20.422	11.646	10.8365	5.228	7.3915	3.627	5.637	2.822	4.7725	2.411
100	39.392	24.637	40.716	29.239	15.0585	7.879	9.9645	5.098	6.856	3.5105	5.4305	2.759
150	47.6955	31.432	65.7975	52.7985	19.42	11.2375	11.86	6.4255	8.2895	4.2935	6.0705	3.152
200	54.083	36.478	84.9215	70.764	22.796	13.988	14.188	8.0795	9.5505	5.2005	6.7565	3.565
500	74.998	53.098	174.267	153.356	38.401	26.652	23.869	16.1	15.198	9.438	10.142	5.722
1000	91.884	66.864	270.084	249.08	45.496	32.028	31.79	22.504	21.654	14.966	14.864	9.574

Figs. 30 and 31 illustrate the changes in the mean starting iteration of the periodic behavior and the mean cycle length, respectively, for each value of k and for each population size. As can be seen, for each population size, there is a single number of epistatic interactions k that describes the latest starting iteration of the period. This value of k also describes the greatest cycle length for the population size.




Figure 30. Mean starting iteration of the period for each value of K and each population size




Figure 31. Mean cycle length for each value of k and each population size.

Initially the results of this experiment seem to contradict the earlier discussion of epistatic interactions in the NK landscapes. This is due strictly to the nature of the single

simple rules governing the agents in the experiment. Had the agents been developed with more complex rules, the understanding of the previous sections would hold. Regardless of the simplicity of the rule, there is an important lesson that is learned from this experiment in relation to innovation generation.

The system of innovation often utilizes epistatic interactions from the same domain as the one being created to assist in the incremental creation of new innovations. This experiment has shown that if the domain of a particular individual innovation contains only a small number of other individual innovations, it would be beneficial to utilize only a small number of these (same domain) individuals in the generation of a new individual innovation so as to maximize the number of regularities in the system. Maximizing the number of regularities increases the complexity of the innovation system and, as such, aids in unique individual innovations being generated that, with market acceptance still influencing the success rate of the innovation, have a better chance of becoming successful due to this uniqueness. Thus, to increase the likelihood of success of the new individual innovation, domains other than the one of the innovation aggregate being generated should be used in conjunction with the given domain.

Conclusions

The system of innovation has been shown to be a complex adaptive system that lives at the edge of chaos with a large number of epistatic interactions. In a manner similar to the NK Landscapes of Kauffman [48], the innovation system traverses the fitness landscapes through the influence of the epistatic interactions but to a limit less

than $K = N - 1$. Given that this landscape traversal has been shown to be the manner in which individual innovations are created, future individual innovations may knowingly utilize the evolutionary operators of mutation and recombination to produce a new generation of individual innovations.

The landscape models have proven that the system of innovation is a system that is highly complex and functions at the edge of chaos but does not deteriorate into a chaotic state. It was shown that the mutation operator is not the sole evolutionary operation in incremental innovation. Rather all evolutionary operations are utilized by the innovation system to foster both incremental and disruptive generations.

Through the study of the historic landscapes of the system of innovation and the observation of their fitness measures, a formal representation of the relation between the mean fitness of the optima of each aggregate of individual innovations and the mean fitness of the population of individual innovations as a whole was developed. This function aids in the understanding of the system of innovation's placement at the edge of chaos. Such an understanding has not been presented in the innovation literature

CHAPTER 5. A CELLULAR AUTOMATA MODEL

In the final discussion of Chapter 4 a model was utilized to explain interactions in simple systems. This chapter extends this model through the consideration of larger neighborhoods about the agents. Additionally, this chapter generates a model of the system of innovation utilizing Cellular Automata (CA). From the CA model various principles are derived for the system of innovation.

The study of complex adaptive systems (*cas*) has gained considerable recognition in the research community over the past two decades. Institutions such as the Santa Fe Institute, the Krasnow Institute for Advanced Study at George Mason University and the Chaos Group at the University of Maryland have been influential in understanding *cas*. Simply stated, *cas* utilizes populations of agents, whose phenotypic traits are expressed with simple rules presenting behavior that is collectively non-linear [14].

As stated previously, the system of innovation has been considered as a *cas* in the literature. Fleming *et al.* [15] discussed the diffusion of innovation as a *cas* and utilized patent data as empirical evidence of such. Likewise, Goldenberg *et al.* [60] and Kun *et al.* [61] considered innovation diffusion as *cas*.

Some researchers explored the relationship between innovation and evolutionary concepts. Ma *et al.* [62] posited the use of agent-based modeling through evolutionary processes to model the influence of consumers on innovation. Andriani *et al.* [63] considered innovation in light of cellular neural networks and Gilbert *et al.* [64] simulated innovation policy with multi-agent networks in an evolutionary manner.

This chapter utilizes modeling techniques guided by Cellular Automata (CA) and Boolean Networks (BNs) to further illustrate the impact of epistatic interactions on the system of innovation. In addition to the BNs, simple cellular automata composed of agents whose phenotype are single rules are utilized to continue the enhancement of the understanding of the epistatic interactions involved in the innovation process. The modeling performed herein illustrates the emergent behavior of the system of innovation.

This chapter offers a methodology for individual innovation generation utilizing the existing population of individual innovations and evolutionary operations using multi-parent crossover. The application of the multi-parent crossover presents a way in which the genotype defines phenotypes that are in conflict and can be eliminated automatically, thereby leaving viable innovations to test for market acceptability as described in [57] and chapter 3 of this thesis.

Boolean Networks in the System of Innovation

Chapter 4 of this thesis highlighted an experiment for observing epistatic interactions at work in a simple system of agents whose updating is governed by the values of their neighbors. The section extends this experiment to discuss various neighborhood types.

Both Cellular Automata (CA) and Boolean Networks (BNs) have been applied to represent various aspects of *cas* as well as evolutionary systems. Kauffman [48] utilized Boolean Networks to assist in the formulation of his NK Landscape theory. Celada *et al.* [65] modeled interactions in the human immune system with CA. Clarke *et al.* [66] used

CA to predict urban growth. Perhaps the best known CA is that of Conway's Game of Life [67], a CA defined on an infinite two-dimensional grid of square cells, where each cell represents an entity that lives or dies based upon perceptions of loneliness and overcrowding.

The *cas* view of innovation allows a system to be modeled with very basic elements to elicit information about the emergent behavior of the system. Using CA to model innovation is not without precedence; Goldenberg *et al.* [68] utilized CA to model innovation exclusiveness in the marketplace. The innovation and CA literature does not include a study using CA or BNs to model the effects of epistatic interactions as they relate to the system of innovation.

The use of existing individual innovations to assist in generating new innovations [69] implies that there are possible implications of innovations being assembled from portions of existing innovations from various domain aggregates. Therefore, innovation spanning multiple domains which interact in an epistatic manner to generate new individual innovations is considered.

A CA consists of a regular lattice of sites. Each site takes on k possible values, and is updated in discrete time steps according to some rule that depends on the value of the sites in some neighborhood around it [70]. Each site in the CA's lattice represents a single entity, an individual innovation for the purposes of this section that is defined by some rule or rules, as the case necessitates. It is common to represent these sites with bits, where a 1 represents a positive value (e.g., alive, active) and a 0 represents the converse.

Boolean Networks are a special form of the generalized CA. In BNs each site is a binary variable with two possible states of activities (on or off). The sites are coupled to one another such that the activity of each site is governed by the prior activity of some other sites according to a Boolean function [48]. The number of sites that influence a given site is known as the neighborhood of the given site. There are two common types of neighborhoods that are considered in CA. In the von Neumann neighborhood the population is arranged in a rectangular matrix of sites, and each site is connected to sites above, below and to either side of it, not including diagonally connected sites [71]. Moore neighborhoods are then defined as von Neumann neighborhoods with the addition of the diagonal connections ($k = 8$). Due to their specialization of the generalized CA, BNs are initially considered in relation to the epistatic interactions of the system of innovation.

Boolean Network Experimental Results

Experiments utilizing Boolean Networks (BNs) to represent the epistatic interactions of the innovation system are now presented. The results presented here offer great insight as to the ontogeny of new individual innovations as well as the system of innovation as a whole. The individual innovations discussed in this section are not necessarily real world innovations but rather models. Later in this chapter, the ideas presented here are supported with industrial examples.

Consider a population of l bit strings such that each bit, $n \in \{1, \dots, l\}$, in the string represents a single innovation. Each bit string then represents an aggregate of individual

innovations to which each bit in the string belongs (e.g., a bit string of wind turbine innovations). Each bit in the string takes one of two states (on or off) as described by Kauffman [48]. The entire population of bit strings then represents a complete innovation system. Each n bit in the l bit strings is defined by a Boolean function of AND, OR, or XOR. The value of each of the n bits in the population is updated through the epistatic interactions of the neighborhood surrounding that bit. Thus, given a von Neumann neighborhood and an AND bit, the AND bit will assume a value of 1 if and only if all of the bits in its von Neumann neighborhood are 1. Likewise, an OR agent will assume a value of 1 if at least one of its von Neumann neighbors contains a value of 1, and the XOR agent requires exactly one of its von Neumann neighbors to contain a value of 1 for it to assume a value of 1.

The agents (bits) in this model are assigned their Boolean function randomly from a uniform distribution. Additionally, each of the n bits is initialized randomly from a uniform distribution of $\{0, 1\}$. Fig. 32 illustrates a sample initial population of bit strings where $l = 4$ and $n = 5$. Once the population is initialized, updates are performed synchronously to all of the n bit strings. Each bit string is updated according to its own rule and the values of the k members of its given neighborhood.

```

1 0 1 1 0
1 1 0 0 0
0 1 1 0 1
0 1 1 1 0

```

Figure 32. Initial populations of a simple Boolean network

The experiment was ran using populations sizes from {50, 100, 150, 200, 500, 1000} and neighborhood sizes from {1, 2, ..., 10} where each neighborhood size represents the number of individuals to be considered in each direction. Fig. 33 illustrates a neighborhood of size 2, where the agent under consideration is surrounded by its size 2 von Neumann neighborhood shaded in gray.

1	1	0	1	0
0	1	1	1	0
1	1	0	1	0
1	0	0	1	1
1	1	0	0	1

Figure 33. Example of von Neumann neighborhood of size 2

The experiment was run five times for each size of population for each size of neighborhood, and the results were averaged to account for the random nature of the initialization. The synchronous updating continued for each run until such time as a repeated bit string was found, which then formed a period. The lattice used for the experiment was considered to be connected from end to end of each bit string and from top to bottom of the population to form a torus shape. This was done to ensure that every agent could interact with the same number of other elements.

The first run of this experiment considered only von Neumann neighborhoods for the epistatic interactions. Fig. 34 illustrates the results of this trial by showing the average number of iterations (transients) that took place prior to the start of a period.




Figure 34. Number of transients prior to the start of period with von Neumann neighborhoods

Fig. 34 illustrates that as the number of epistatic interactions increases (*horizontal* axis), the number of transients prior to the start of a period decreases dramatically. It was found that the length of the period (number of iterations before the system began to repeat itself) also dramatically decreased as the number of epistatic interactions increased. Table 2 describes this result, and Fig. 35 illustrates this result graphically for population sizes of up to 200. Populations larger than 200 were removed from the chart due to skewing of the graphic.

Table 2. Length of period for von Neumann neighborhoods

No. of Interactions	Pop. of 50	Pop. of 100	Pop. Of 150	Pop. of 200	Pop. of 500	Pop. of 1000
1	26.4	56.8	324	108	7848	7848
2	4.8	12	20	28	73.6	73.6
3	2.4	2.8	2.4	3.2	5.2	5.2
4	2	2	2	2	2	2
5	1.4	1.8	2	2	2	2
6	1.2	1.2	1.6	1.2	1.8	1.8
7	1.2	1	1	1.4	1.6	1.6
8	1	1	1	1	1.2	1.2
9	1	1	1	1	1	1
10	1	1	1	1	1	1




Figure 35. . Period length for von Neumann neighborhoods

The use of the von Neumann neighborhood indicates a trend in the decrease of the transient and period length as the number of epistatic interactions increases. This

decrease then shows that the number of regularities (amount of complexity) in the system decreases. Given the nature of the Boolean Network, this is not surprising. In a previous experiment, it was discovered that given a single bit string (innovations of a single aggregate of innovations), there was an optimal number of epistatic interactions within that domain prior to the same decrease in complexity being observed, as illustrated in Fig. 36.




Figure 36. Period length for epistatic interactions of single bit string experiment

To determine if the decrease in period and transient length is truly a result of the increased number of epistatic interactions, the previous experiment was conducted, which used von Neumann neighborhoods on a given population of bit strings, utilizing Moore neighborhoods. The results of this experiment are shown in Figs. 37 and 38.




Figure 37. Transient length with Moore Neighborhoods




Figure 38. Period length with Moore neighborhoods

As indicated in Figs. 37 and 38, the decrease in complexity is exaggerated by the increased number of epistatic interactions indicative in the BN. Given market acceptance,

an increase in complexity in the innovation system as a whole enhances chances of individual innovation success. Decreases in the complexity of the innovation system indicate great similarity in innovation types, which result in a narrowed market.

Difficulty for new individual innovations to break into the market, due to increased similarity between innovations, is found. Is the trend towards stasis indicative of the neighborhood size and configuration, or is it a system-wide principle? To answer this question, an experiment utilizing random neighborhoods was performed.

Boolean Networks with Random Neighborhoods

The trend of the modeled innovation system above to attend to a fixed point attractor of stasis at or near 0 for the period and transient length could be an artifact of the neighborhood configuration utilized in those experiments. Therefore a third type of BN experiment was attempted.

In this experiment, each agent was initialized randomly as an AND, OR, or XOR agent with an initial value of 1 or 0, as described in Section 2.1. The neighborhood size remained the same as in the initial experiments. The factor that made this experiment unique from the previous ones is that the definition of neighborhood was changed dramatically. Neighborhoods in this experiment were considered to be randomly selected agents rather than topologically related ones. Thus, agents were randomly chosen from the population without replacement to become the neighbors of each agent. No record of the randomly chosen agents was maintained, and each agent chose randomly new

neighbors at each iteration. This definition of neighborhood was used to remove any configuration bias in the experiment.

Fig. 36 illustrates the findings of the number of transients prior to the start of a period in this experiment. Fig. 40 displays the period length for this experiment.




Figure 39. Number of transients prior to the start of a period for Boolean networks with random neighborhoods




Figure 40. Period length for Boolean networks with random neighborhoods

Figs. 39 and 40 illustrate the trend towards stasis of immediate period detection and period lengths of 1, as found in BNs with randomly selected neighbors, just as it was in the previous experiments.

The trend towards stasis at a period length of 1 and zero transient length offers the novel insight that new individual innovations should limit the number of existing innovations from which the new individual innovation is generated. There may be concern that BN rules are dissimilar to those in use by the system of innovation. The next section describes the use of a CA model of the system of innovation to address this concern.

Cellular Automata Model of the System of Innovation

This section describes the use of Cellular Automata to form a model of the system of innovation. Using CA models to understand aspects of innovation is not without precedence in the innovation literature. Sosa *et al.* [72] applied CA to model social influences in the adaptation and adoption of new innovations. Ratna *et al.* [73] discussed the use of agent-based modeling to capture the complexity of medical innovation diffusion. Similarly, Schwarz *et al.* [74] considered modeling the diffusion of environmental innovations.

The model presented here considers a regular lattice of agents, connected in a fashion similar to that of the lattice in the previous experiments. Each agent takes on a value from $\{0, 1\}$ and is assigned an initial value randomly from a uniform distribution

and a given initial population density variable. Each agent is defined by the same rules; if x percent of the agent's neighbors (von Neumann or Moore neighborhoods are determined at the start of the model run) are successful (alive), then it is possible for the given agent to be successful (alive). Additionally, each agent's domain is inspected to verify that the number of successful agents in the domain is not larger than y percent of the number of agents in the domain. If the number of agents in the given domain is less than the percent given, and the percent of agents in the neighborhood of the given agent is not less than the given x percent, then the agent is considered successful (alive).

While the rules of this model are simple, they do represent real rules of the system of innovation. If a new individual innovation is generated from existing successful innovation, then the greater the chance it will be successful. Likewise, if there are too many innovations in the current domain (aggregate of individual innovations), it is less likely that a new innovation will be able to enter the market in that domain and be successful.

Since individual innovations are deemed such by market acceptance [57], there is a time when an individual innovation is no longer viable in the market and therefore is no longer considered innovative. To account for this factor in the CA model being presented, a random variable that represented a death rate for individual innovations was added. Thus, after a prescribed number of iterations, each innovation agent can be deemed no longer successful based on whether a random number is less than the given death rate.

The results of the model are displayed graphically using a 100 pixel wide by 256 pixel tall bitmap. Each pixel represents a single innovation agent. Each line of the bitmap

represents a single domain of innovations. Population density, percent of neighbors required to be successful for new a generation, percent of domain members that bound a new generation, and the death rate are all variables that can be assigned by the user from the domain of (0,1), excluding the end points. Additionally, the user can determine whether to use von Neumann or Moore neighborhoods. Fig. 41 illustrates three examples of initial populations.




Figure 41. Sample of initial populations of the CA model. Starting from left, the populations size is 95%, 30% and 5% of the system

A sample run of this model is illustrated in Fig. 42. The model parameters were 5% initial population size, 30% of neighbors required for success, 65% of domain members as the success boundary, a death rate of 5% after 25 generations, and the use of Moore neighborhoods.




Figure 42. Run of CA model with 5% initial population and only 30% of neighbor success required. The iteration count starting from the left is 0 (initial), 10 iterations, 25 iterations, 50 iterations and 100 iterations.

Fig. 42 illustrates the complexity that becomes inherent in the innovation model after just 100 iterations. Prior to activating the death rate (at the 25th iteration), the system builds innovations bounded by only the neighborhood support and the parameters of the domain bounds. After the death rate is activated, the system continues to produce new innovations while allowing older innovations to die off. Fig. 43 demonstrates the system used in Fig. 42 with only the neighborhood success rate changed to 50%.




Figure 43. Run of CA model with 5% initial population and 50% of neighbor success required. The iteration count starting from the left is 0 (initial), 10 iterations, 25 iterations, 50 iterations and 100 iterations.

The results in Fig. 43 are remarkable. A 20% increase of the neighbor success parameter created a system in which no new innovations were generated. Additionally, as the death rate is activated, the current innovations die off at a startling rate, such that by 160 generations, zero innovations exist in the system. To observe if this trend is an artifact of the initial population size, the system was run with the same parameters as in Fig. 43, after changing the initial population percent to 95%. Fig. 44 illustrates the results of this run.




Figure 44. Run of CA with 95% initial population and 50% of neighbor success required. The iteration count starting from the left is 0 (initial), 25 iterations, 50 iterations, 100 iterations and 250 iterations.

Fig. 44 shows a trend towards the same results as seen in Fig. 43. In fact, the results were identical in the end. After 1421 iterations, the system described in Fig. 44 presents zero innovations. These results indicate that while the population size does influence the rate at which the system achieves the fixed point attractor of zero innovations, population size does not impact the end result of attending to the fixed point attractor.

After many runs with each parameter being changed systematically, it was discovered that the only factor governing whether the system attends to the fixed point attractor of zero innovations or remains in a complex state is the percentage of neighbors that are successful. It was found that if this parameter is set at 42% or higher, the system attends to the fixed point attractor. Lower values of this parameter allow the system to achieve the complex dynamic required for innovation.

In addition to the discovery of the fixed point attractor, it was determined that the innovation model in its complex state exhibited niche creation. It has been speculated that niche creation is common in *cas*. Niche creation is regarded as the systematic cascading of agents to fulfill a role that was fulfilled by a previous agent. In other words, as one innovation dies, another takes its place. This can be witnessed with the death of vinyl records and the emergence of compact discs for music representation.

The results of the CA model of innovation, as well as those of the BN models, permit the formulation of the following innovation hypothesis.

Innovation success relies on the limited use of epistatic interactions of existing innovations; an excessive number of interactions increase the chances of innovation failure. Incremental innovation proceeds through the use of these limited interactions to evolve new innovations.

Individual innovation success is truly determined through market acceptance, but limiting the number of epistatic interactions in innovation generation will increase the chances of market success. In other words, it is unwise to attempt to combine every invention under the sun into one innovation. In fact, it is wise to limit the number of interactions to three or less, according to the results of the CA model.

Additionally, individual innovations may be predicted based on the approach presented in this paper. They can be generated synthetically using a small number of epistatic interactions and a genetic programming algorithm. Evolving innovation candidates and their comparison against a fitness function that is in part defined using

market acceptance data mined from the internet offers a way to anticipate new innovations and to evaluate their possible success.

Evolution of Individual Innovations with Multi-Parent Crossover

This section illustrates the use of evolutionary operations to evolve new individual innovations from existing ones, based on the findings in the previous sections of this chapter. The evolutionary operator of multi-parent crossover to generate multiple children is used. Multi-parent crossover is not a new concept. Eiben *et al.* [75] discussed the use of multi-parent recombination in genetic algorithms to produce a single offspring. Similarly, Lis *et al.* [76] applied a genetic algorithm with multi-sexual parents to generate a single child for optimization.

Isaacs *et al.* [77] utilized a probabilistic form of multi-parent crossover to evolve sets for the representation of truss designs. This form of multi-parent crossover is dissimilar to the one assumed in this paper, as the children generated through this form are a combination of elements common to all parents and a probabilistic select of uncommon elements. Ting [78] offered a theoretical basis for the effectiveness of multi-parent crossover. This current literature fails to address the topic of evolutionary innovation with multi-parent crossover, which is a novel contribution of this chapter.

The evolutionary process of crossover traditionally selects two fit individuals from a given population to utilize as parents. A point in each individual is selected as the crossover point (often this point is in the same location for genetic algorithms), and the genetic material of the parents is exchanged at this point [79]. This technique is often

referred to as single point crossover. Multi-parent crossover utilizes n e fit individuals as parents and selects n loci as the point of exchange to produce $n!$ offspring. Fig. 45

illustrates multi-parent crossover for $n = 3$ parents.




Figure 45. Multi-parent crossover illustrated with $n = 3$ parents of a genetic algorithm whose genotype is represented as bit strings.

While the concept of multi-parent crossover is rudimentary, there remain hidden pitfalls to be considered. The largest of these pitfalls is phenotypic contradiction. Often the genotypic representation of an individual influences the phenotypic traits that are exhibited by that individual. In the case of innovation, the genotypic traits are the elements that define the innovation's makeup and determine to a great extent the phenotypic traits (behavior) of the innovation. Should contradicting phenotypes be discovered after the crossover operation, such as a loci whose allele states that the innovation should float in water being paired with a loci whose allele states that the innovation should have an outer casing, that is a holey material, that individual should be culled from the population.

While it is possible that the fitness of the malformed individual would be such that it would not be selected for evolutionary recombination, a more efficient methodology would be to remove all such individuals from the population. Determining individuals displaying contradicting phenotypes is not a trivial task to automate. Thus, while being more efficient to the overall evolutionary process to remove such individuals from the population, it would be a difficult task to perform. Therefore, a clearly defined fitness function is required to ensure such individuals are not propagated through evolutionary processes.

To illustrate the effectiveness of multi-parent crossover in evolutionary processes, a genetic algorithm utilizing single point crossover with two parents was compared to a genetic algorithm (GA) utilizing multi-parent crossover. Each GA had a population of 500 individuals whose representation was a bit string of length 100. A fitness function of a single randomly generated bit string of length 100 was used to assess the fitness of the individuals in the population of each GA. Both GAs used the exact same fitness string to insure proper comparison. For the purposes of this experiment, the mutation operator was not used in order to compare the effectiveness of each crossover operation properly. Additionally, both GAs started with identical populations and ran for 200 generations. To insure proper coverage, the experiment with the GAs was conducted 100 times, and the results were averaged. Multiple parent sizes were tested for the multi-parent GA for each 100 sample run of the GAs. The results of each 100 run test were similar. The multi-parent GA outperformed the single crossover GA with every parent size tested (from 3 to 10). Figs. 46-47 illustrate runs of the two GAs using a parent size of 5 for the multi-

parent crossover GA. As can be seen, the multi-parent GA is more efficient at approximating the fitness string.




Figure 46. The average population fitness of single point crossover GA compared to a multi-parent GA with 5 parents per offspring.




Figure 47. The fitness of the fittest individual in the population of a single point crossover GA compared to a multi-parent GA with 5 parents per offspring.

A genetic programming algorithm can be used to generate new individual innovations. The current innovations are represented in the genetic programming algorithm as trees whose nodes are properties of the innovation. Multi-parent crossover and mutation of the initial population of innovations result in new candidates for innovation. As shown above, the number of current individual innovations that undergo crossover should be limited to a small number. As seen in the GA experiment above, though, the number of current generation innovations should be larger than two to ensure the best performance.

Evolutionary generation of innovation requires the elimination of phenotypic conflicts. To assist in automating this elimination, it is suggested that the fitness of each new innovation be determined based upon market requirements and discounted if deep similarities to existing innovations occur. Market requirements are easily accessible with the use of data-mining techniques of Chapter 3 and [57]. Discounting the fitness due to close similarities to existing innovations assists the newly generated innovations in generating a niche market that is not currently filled with competitors.

Industrial Examples

This section presents examples of individual innovations that have been generated through the use of epistatic interactions. These cases support the hypothesis presented in the above sections.

It is not difficult to find examples in the current marketplace where individual innovations have been epistatically built. Consider, for example, the telephone and its evolution over the years. Initially, phones were single purpose objects that were wired to a global grid through telephone poles and companies. A new innovation aggregate, cell phones, was produced when the telephone was combined with wireless communication technology. Cell phones initially served the single purpose, like standard telephones, of voice communication, but unlike standard telephones this communication took place in a mobile environment. With the advent of email communications, a new innovation was produced by the epistatic interactions of cell phones and email. This innovation took various forms, such as text messaging and email on the cell phone (Blackberrys or Smart Phones).

Evolutionary progress has continued to combine individual innovations with the telephone to the point that cell phones are now also cameras, audio/video players, game machines, electronic filing cabinets, and more. While the initial epistatic interactions were extremely limited, phones and wireless communication, the number of interactions has grown. It can be seen, though, that each generation of the evolution of the phone has included only a limited number of these interactions, such as the addition of email and internet browsing to cell phones. Throughout the evolution of the telephone, the number of epistatic interactions was small.

The basic computer printer is another example of evolutionary processes at work in the system of innovation. Computer printers for the consumer market were originally dot matrix printers with a ribbon of ink and a rotating head to “type” the characters onto

the page. It did not take long for this innovation aggregate to evolve into laser and inkjet printers that replaced the ribbon and rotating head with technology that was incorporated from Xerox-type copying machines. With the expansion of facsimile machines into the residential market, another evolutionary step was taken, and the advent of multi-purpose printers was witnessed.

Today, it is possible to have a computer printer in the home that prints standard documents, but is also a photo printer and a fax machine, and can even send emails. As with the evolution of the telephone, computer printer evolution has taken place utilizing epistatic interactions from various domains, although these interactions were limited at each step. A record of a giant leap from dot matrix to multi-function printers is not found. It is not difficult to find numerous examples of epistatic interactions being utilized to generate new innovations. Nor is it difficult to see in the lineage of individual innovations that each evolutionary step occurred with a limited number of those interactions. Limiting the number of epistatic interactions is proven theoretically and empirically to be the most efficient manner for evolutionary innovation to proceed.

Conclusions

The literature contains numerous examples of research related to how one should innovate. Missing from that corpus is a systematic study of the innovation system itself. This chapter illustrates this novel study through the use of complex adaptive systems and modeling. Prior to this point the innovation process was thought of as a methodology rather than a system that could be studied and defined.

This chapter illustrates that the system of innovation is a complex adaptive system. It was shown that relatively simple innovation models can be constructed. A novel approach to modeling innovation through the use of Boolean Networks and Cellular Automata was described. From the models of the innovation system, it was shown that limiting the number of epistatic interactions in evolutionary innovation offers greater chances of innovation success in the marketplace.

Further, in the use of the cellular automata models, it was illustrated that individual innovations are most often generated through the use of previous innovations external to the domain currently being generated. The CA models, novel to this study, have shown that the limited use of cross-domain parents to generate innovation is the actual process that the system of innovation utilizes at each generation.

This chapter introduced the novel concept of the multi-parent crossover operator for evolutionary innovation. It was demonstrated that multi-parent crossover increases the efficiency of the genetic program used to generate new innovations from existing ones. The impact of multi-parent crossover on phenotypic contradictions was discussed, and methods for handling such contradictions were offered. The use of a fitness function of user requirements for innovation is also unique to this study of the innovation process.

The innovation hypothesis presented in this chapter demonstrated the evolutionary nature of the system of innovation, where a small number of cross domain epistatic interactions increased chances of success. Such a modeling approach can help predict future individual innovations by recombination and fitness functions of market acceptability which can be mined, for example, using cyberspace.

CHAPTER 6. AN AGENT-BASED INNOVATION ECOSYSTEM MODEL

In this chapter, an agent-based model of an innovation ecosystem is presented. The ecosystem is considered as an environment in which the individual agents (innovation entities) exist and interact. The interaction of the agents is with the other agents in the ecosystem as well as the dynamic environment itself. The model that is described in this chapter seeks to understand the factors of importance with which innovating entities (e.g., corporations, smaller businesses) realize an optimization of their resources. The literature lacks such a model of an innovation ecosystem.

An agent-based model utilizes simple agents interacting with their environment through sensors and actions [10]. The environment that the agents interact with is defined by the model which defines the interaction with other agents in the environment as well as external constraints or boundaries. The use of agent-based modeling to understand complex systems is well documented in the literature. For example, Koritarov [80] applied agent-based modeling to simulate electricity market scenarios and Tesfatsion [12] described an agent-based model of complex economic conditions.

Agent-Based Modeling

An agent-based model includes a set of simple agents that interact with each other and their environment. The agent interaction mode varies as so are the resultant models. Both, simple message passing interfaces with which the agents communicate [81] and

direct communication with various evolutionary processes to update the agents [82] have been researched.

Agents are often defined with an initial set of simple rules which are updated at discrete time intervals based upon a history of actions and sensed environment [14]. The rules that define the agent initially define the response of the agent to specific conditions in the environment and trigger that response to these conditions. The updating of the rules proceeds in discrete time intervals, often synchronously with the other agents and the environment. The updating process varies, though most often it relies on the agents making decision based upon past experience. Therefore, agents tend to maintain a history of their experiences and update their rules adapting to this history.

A key aspect in agent-based models is the lack of a centralized control [83]. Rather each agent controls its own updating and actions. This lack of centralized control makes agent-based modeling well suited for simulation of ad hoc networks of entities such as an ecosystem of innovating companies. The lack of centralized control over a system ensures that the only possible means with which the system can be understood is through simulation, which is the case with many social and economic scenarios.

An agent-based model iteratively produces its emergent behavior (see Figure 48). In an agent-based model, a population of agents is initialized with their defining rules, often randomly generated. The agents then sense their environment and interact with one another and record this information as a history. The agents update their rules based upon this history and begin the process of sensing and updating all over again.




Figure 48. Iterative agent-based model

During the iterative model development, it is often desirable to record the states of the environment and the agents to facilitate the observation of the emergent behavior of the system. Observing the complete system at discrete time intervals allows the user to understand more fully the behaviors emerging within the system and to possibly abstract rules based upon these behaviors. It is the emergent behaviors of the set of agents that makes agent-based modeling so appealing for understanding complex system.

Various methodologies are used for rule updating of agents. Some models utilize a simple look up tables of rules that the agents can choose from based upon their history. Other models utilize evolutionary concepts, such as genetic algorithms [84] or even a simple mutation to update the rules of each agent. Regardless of the updating

methodology, the rule updates are critical to the evolving of an optimal system and to agent survival within that system.

Updating

At each generation the resources of the agents are updated to reflect their adherence to the market requirements. This update is performed as a percentage of the agent's resources and is defined in (6).

$$\varphi_{t+1} = \varphi_t f(a_i) \rho$$

Equation 6. Updating

where $\rho \in (0,1]$ is a payoff parameter set by the user, φ is the resources of the agent, at time t and time $t + 1$, and f is the fitness function of each agent given by (7) where r_j is the allele at the j^{th} loci of the requirement string r with N properties.

$$f(a_i) = \frac{1}{N} \sum_{j=1}^N \delta_j \text{ where } \delta_j = \begin{cases} 1 & a_{ij} = r_j \\ 0 & \text{otherwise} \end{cases}$$

Equation 7. Agent Fitness

The resources of the agents are also used in determining the market share each agent holds in the domain of their given product. This market share is given by the agent's resources in relation to the total resources of all agents in the given domain as described by (8).

$$m_i = \frac{\varphi_{a_i}}{\sum_{j=1}^K \varphi_{a_k}}$$

Equation 8. Market share

where m is the market share of agent i , φ is the resources of agent i , and K is the number of agents whose products are in the given domain.

The market share of each agent is then utilized in each generation to determine if any agent will interact with a selected agent. Each agent maintains a minimum market share variable that indicates the required market share of each agent that the given agent will interact with. In the interaction, an agent chooses another agent to interact with. If the chosen agent's market share is below the threshold of the first agent then that agent will not be accepted for interaction. Upon interaction acceptance the agent with the larger market share is compelled to pay a portion of its resources to the other agent. The amount of resources paid, μ , is given as a product of the paying agent's resource and the cost of interaction parameter, θ , as shown in (9).

$$\mu = \varphi_i \theta$$

Equation 9. Payoff

The interaction of the two agents is performed using crossover. The properties of each agent are used as the parents for the crossover process. A random point within the bit strings of the properties is selected and the values are exchanged to form two new products. These new products replace the existing products of the agents. This form of

crossover interaction encourages diversity within the population of agents and allows for more fit agents to exchange information with less fit agents.

During the interaction the interacting agents have the opportunity to learn from one another. This learning is governed by each agent's learning rate. A random number is chosen and if it this number is less than the agent's learning rate, then that agent may learn from the other agent in the interaction. Learning takes place utilizing the evolutionary operation of cross-over similar to the product property exchange discussed above. Vectors of each agent's interaction rate, interaction minimum value, retool rate, amount of retool, learning rate, and the rates at which the agents adjust their own parameters are formed for each agent. A random loci within the parameter is selected and the alleles of the loci proceeding the swap point of one of the vectors is combined with the alleles of the loci following the swap point of the other vector. In this manner agents are able to learn from more successful agents in the population.

At each generation agents are given the opportunity to retool their products. A random number is selected and if that number is smaller than the agent's retooling rate, the agent may retool its product. Retooling proceeds with the use of the mutation operator from evolutionary computation. Each loci of the property string of the agent's product is selected for mutation if a randomly selected number is less than the agent's retooling amount parameter. Retooling in the real world incurs a cost (e.g. new machinery, new production methodologies). In the innovation ecosystem, retooling costs the agent a portion of its resources as a simulation of the real world costs. For the purposes of this

model the retooling cost was fixed at 15% of the agent's resources. Future models may set this as a parameter for user input.

With each generation the agents are given the opportunity to change their rate of interaction. If the history of the agent's interactions indicates a downward trend in resources the agent may decrease their rate of interaction. Should the agent observe an upward trend in its resources over the course of the given history, that agent may decide to increase its interaction rate. This evaluation of the history and subsequent alteration of the rate of interaction allows the agents to adapt to conditions present in the market place. Finally, agents whose resources have fallen below a specified threshold for more than 5 generations are allowed to perish. These agents are then replaced with new agents. All agents in the model are generated with their parameters randomly initialized. Thus, the agents replacing the extinct agents incur the same initialization procedure as the initial population, irrespectively of the generation number. In this model the population size is maintained at 1000 agents and therefore the number of extinct agents is replaced by the same number of randomly generated agents.

Model Results

For the purposes of exploration the model described in section 3 was ran with 125 different sets of parameters. The model maintained the constant parameters of:

- Population size = 1000,
- Number of product domains = 30 and

- Number of properties per product = 100.

The parameters of fitness change rate, cost of interaction and payoff were accounted for with values from (0.01, 0.05, 0.10, 0.20, 0.50). Each run of a set of parameters was repeated 10 times and the average of the ten runs was recorded. This allowed for the randomness of the model to be discounted in the results of the runs. During each run results were recorded for:

- The mean fitness of the population,
- The fitness value of the fittest individual,
- The mean interaction rate of the population,
- The mean market share of the population,
- The maximum market share of an agent in the population and
- The mean resources of the population.

The results discussed above were recorded at each generation to preserve a historic representation of the system for each parameter combination. These results were averaged over the ten repeated runs of each parameter combination as previously discussed. The results of these averaged runs are now presented.

All the runs of this model were limited to 1000 generations to ensure accurate comparison ability. Fig. 49 - Fig 51. illustrate the results of a single averaged run of the model, showing mean fitness, mean interaction rate and mean resources respectfully.

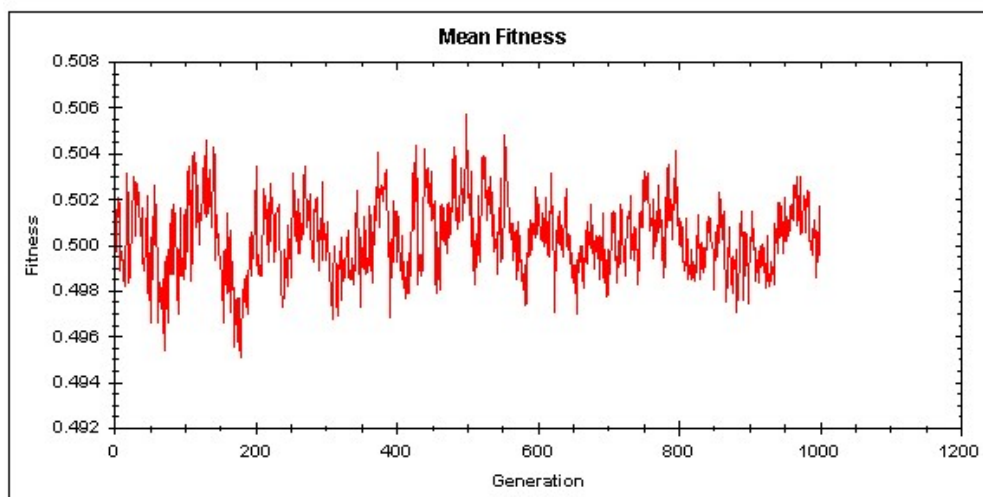


Figure 49. Mean fitness for an average run of the model with the fitness change rate = 0.01, the cost of interaction = 0.01, and the payoff = 0.01.




Figure 50. Mean Interaction rate for an average run of the model with the fitness change rate = 0.01, the cost of interaction = 0.01, and the payoff = 0.01.

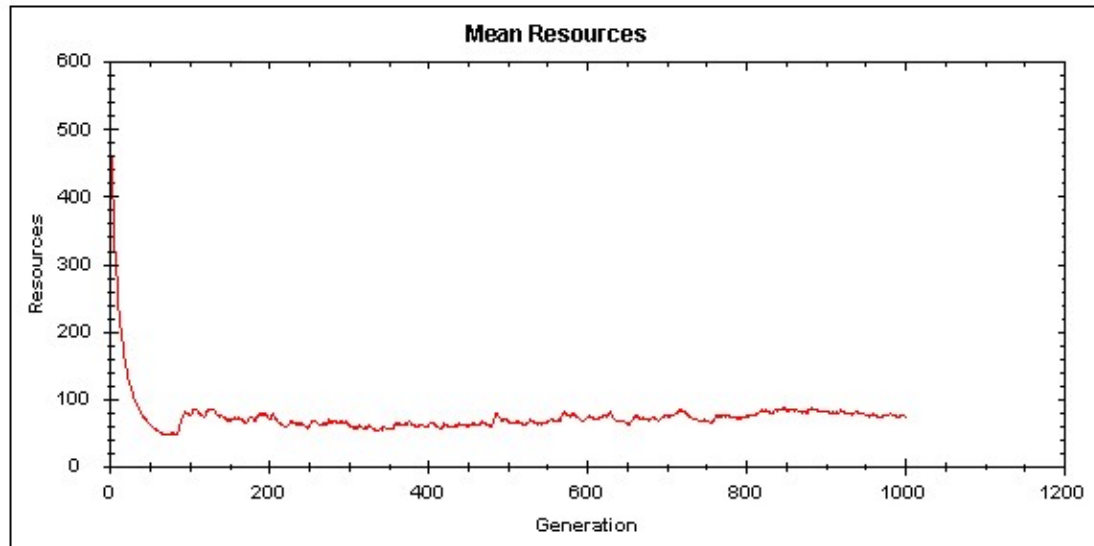


Figure 51. Mean agent resource for an average run of the model with the fitness change rate = 0.01, the cost of interaction = 0.01, and the payoff = 0.01.

As can be seen from the Fig. 49-Fig. 51 the mean interaction rate trends downward throughout most of the run with sharp declines in the first 100 generations. Similarly, the mean resources of the population take a sharp decline in the first 100 generation but then they appear to stabilize for the remainder of the run. The mean fitness of the population fluctuates constantly throughout the run but remains within a bounded region.

The initial hypothesis of this model was that the mean interaction rate would increase as the cost of interaction remained low and the payoff for each generation increased. Unfortunately, this hypothesis was proven to be faulty. Fig. 52 – Fig. 56 show that the mean interaction rate remains within a bounded constant region regardless of the cost of interaction and the amount of payoff. The results for runs where the fitness change rate remained constant at 0.05 are presented in Fig. 52 – Fig. 56.




Figure 52. Mean interaction for an average run with cost of interaction = 0.01 and the payoff = 0.01.

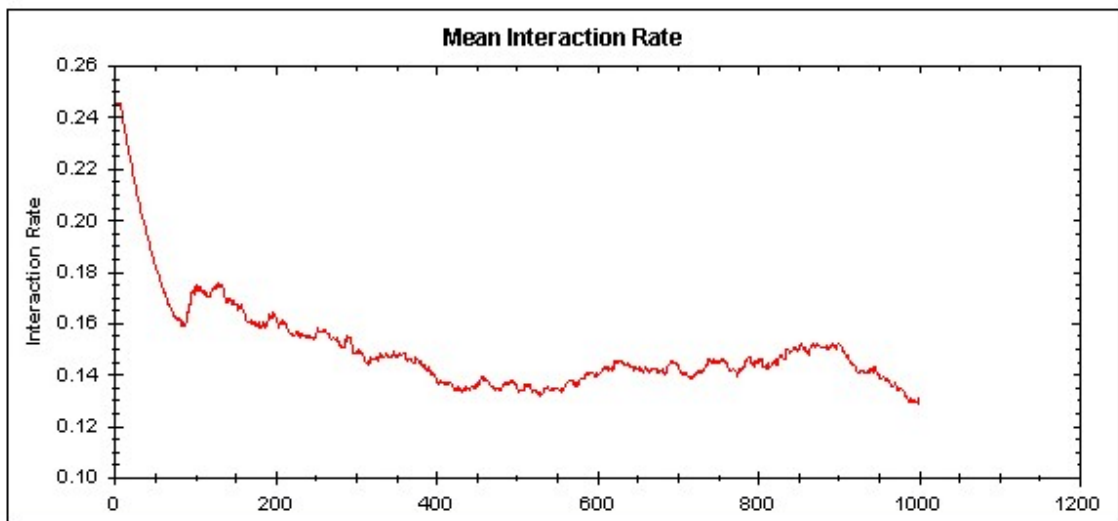


Figure 53. Mean interaction for an average run with the interaction cost = 0.01 and payoff = 0.05




Figure 54. Mean interaction for an average run with the interaction cost = 0.01 and the payoff = 0.010.




Figure 55. Mean interaction for an average run with the interaction cost = 0.01 and payoff = 0.20.




Figure 56. Mean interaction for an average run with the interaction cost = 0.01 and the payoff = 0.50.

As can be seen in Fig. 52 – Fig. 56, the mean interaction rate continued to trend low for all of the runs which is contrary to the initial hypothesis. In fact, the trend appears to be just the opposite, the higher the payoff the less likely the agents were to interact. Fig. 57 and Fig. 58 illustrate that in the same runs the mean resources of the population changed little from a payoff of 0.01 to a payoff of 0.50, ruling out the thought that resources increase as payoff increases.




Figure 57. Mean resources for an average run with the interaction cost = 0.01 and the payoff = 0.01.




Figure 58. Mean resources for an average run with the interaction cost = 0.01 and the payoff = 0.50.

The simulation has shown that the cost of interaction and the payoff rate had little, if any, impact on the change in the system behavior. The only factor that weighed heavily

on the behavior of the system was the rate at which the market requirements (fitness) changed. It was discovered that if the rate of change of the market requirements increased that the mean interaction rate and the mean resources of the population increased as well. In the case of the resources they began to increase in an exponential fashion. This fact is illustrated in Fig. 59 - Fig. 62 where the cost of interaction and payoff parameters were set to 0.10.




Figure 59. Mean interaction for an average run with the fitness rate of change = 0.01.




Figure 60. Mean resources for an average run with the fitness of rate change = 0.01.




Figure 61. Mean interaction for an average run with the fitness rate of change = 0.50.




Figure 62. Mean resources for an averaged run with the fitness rate of change = 0.50.

As shown in Fig. 59-Fig. 62, the mean interaction rate increases dramatically as the rate at which the market requirements change increases. In addition, the mean resources of the population experience exponential growth with the frequent changes of the market requirements. This type of emergent behavior cannot be accounted for through considering the agent definitions alone. Rather, this behavior is truly complex and can only be understood from the creation of a model such as the innovation ecosystem given here.

Discussion and Examples

The innovation ecosystem model built in this chapter has yielded some startling insights unto the emergent behavior of a complex system. The results presented above

where unexpected and point to some useful principles related to innovation and the global market place.

The amount of interaction between innovating entities is directly related to the rate at which the market changes. The model shown above illustrates that a rapidly changing marketplace calls for greater interaction between companies to ensure success. This is not counter-intuitive but such dramatic evidence for this statement has not been established in the literature. The model illustrates that a rapidly changing market requirement offers a great opportunity for investment and resource gain. When considering the rapidly changing technological markets the increase in resources shown in the model is justified.

It has been shown that in the real world market with rapidly changing technology, partnership innovation increases the chances of success. Companies such as Microsoft which initially concentrated on software production have teamed with other entities to produce entertainment devices and communications ability as seen in the Sync product in which Microsoft teamed with Ford Motors [85]. The technology of communications and entertainment is rapidly changing therefore innovation entities must seek partners to work with in order to keep up with the rapidly changing market requirements. With technologies that change rapidly it is doubtful that a single innovation entity can retool fast enough to maintain a market lead.

A simple search of the internet shows that there is prolific use of the concepts presented in this paper in the semiconductor industry. Siemens and Infineon teamed together in 2007 to produce power semiconductors for use in the power generation fields

[86]. Quimonda, a German semiconductor supplier of memory products, in 2008 teamed with Centrosolar Group, to produce crystalline solar cells [87]. In 2006 Ball Semiconductor joined forces with K. K. Dai-Nippon Kaken, a displays company, to produce maskless exposure equipment [88]. These are but a few of the many examples of semiconducting firms teaming with extra-domain entities to produce new innovations. Fields with rapid technological growth such as the semiconductor industry experience some of the highest change in innovation requirements. Therefore, it is essential for the survival of these companies to partner with other companies to keep up with the changing requirements.

An interesting behavior that was discovered in the innovation ecosystem was that of niche creation and filling. It was discovered that product domains with lower rate of change in requirements experienced agents whose interaction rates remained extremely low but whose resources would trend upwards. These markets can be viewed in similarity to the market for grandfather clocks. This market has changed little over the last 100 years in terms of its requirements, yet companies perform relatively well in this market if they are established. The model showed that it took a long amount of time for an agent to gain a large market share in these types of markets but once it had that larger share it was difficult for the agent to lose it. The niche creation was discussed in [13].

Therefore, it can be concluded that if an innovation entity is working within a rapidly changing domain, such as high end technology, it is beneficial for that entity to work in partnership with another innovation entity in order to meet market demands. Further, if an innovation entity is in a market domain in which requirements change very

little, that entity should seek to maximize their longevity in the market and minimize their innovation interactions with external entities.

Conclusions

An agent-based model of an innovation ecosystem was presented. The model presented illustrates the complex adaptive system that is innovation. While the model reduced the number of active variables to a minimum it fully represents the emergent behavior of innovation entities working within an innovation environment. This form of innovation ecosystem is novel to this chapter and is not found in the existing literature.

This chapter has shown that innovation entities within rapidly changing market domains benefit from the partnership with external entities. Furthermore, the chapter showed that innovation entities stand to see an increase in resource acquisition if the market domain they operate within is rapidly changing. Niche creation within slower changing markets was also discussed.

CHAPTER 7. CONCLUSION

This thesis has posited the existence of a system of innovation. This system of innovation has been shown to be a complex, non-linear, adaptive system. Due to its nature as a complex adaptive system, the system of innovation cannot be described by its components (individual innovations) alone. Rather, as has been shown in this thesis, the system of innovation produces emergent behavior that is understood through modeling.

This thesis presented a historic overview of the methodologies and theories of innovation that have been put forth in works of other researchers. From this overview it was understood that theories of innovation have encountered at least five different generational morphologies. Unfortunately, none of these generations considered innovation as a system with governing properties of its own, aside from the governance of the innovating entities.

Presented in this thesis is a methodology for gathering the requirements for individual incremental innovation. These requirements were gathered from the internet through the use of user and expert reviews as well as patent databases. An apriori style mining algorithm was presented for the discovery of the frequent current requirements for individual incremental innovations. These requirements were then placed into an AND-OR tree that may be utilized by evolutionary processes to predict new individual innovations that meet market acceptance.

This thesis presented several chapters which modeled the system of innovation as a complex adaptive system. Chapter 4 illustrated the historic landscapes that the innovation system has traversed. From these landscapes an important conclusion was

drawn. The system of innovation does not tend toward a chaotic state but rather resides at the “edge of chaos” where ordered complexity is observed. The discovery of this principle of the innovation system allowed for further work in modeling the system.

Chapter 5 described the system of innovation through simple models of cellular automata. These models, with their simple agents defined by simple rules, allowed for the discovery of another important principle of the system of innovation. The innovation system makes use of epistatic interactions which are small in number to produce new individual innovations. These interactions often occur outside of the domain aggregate of the individual innovation being generated. Further, these interactions illustrated the sensitivity of the system of innovation to initial external conditions.

Chapter 6 presented an agent-based model of an innovation ecosystem. This ecosystem considered innovating entities as a part of the system of innovation. The interactions that took place between the entities showed that domain aggregates whose phenotypes are defined in a rapidly changing marketplace are best served through partnered innovation with other, possibly aggregate external, entities. This model also illustrated niche generation in the complex adaptive system of innovation.

A system of innovation does exist and that system is truly a complex adaptive system. Through adherence to the principles posited here and model of the system one can expect to better understand the innovation system. Clearly, the system of innovation is far from being completely defined, although this thesis has contributed greatly to the formation and understanding of a complete system of innovation.

REFERENCES

- [1] Stokic, D., Campos, A., Sorli, M., and Gorostiza, A., "KM Systems to Support Incremental Innovation in Manufacturing Industry," in *10th ISPE Intl Conf on Concurrent Engineering*, Maderia, 2003.
- [2] Kusiak, A., "Innovation Science," in *Proc. of the 11th IEEE Intl Conf on Emerging Technologies and Factory Automation*, Prague, Czech Republic, 2006, pp. 507-514.
- [3] Kusiak, A., "Innovation of Products and Services: Bridging World's Economics," in *Proc. of the 19th Intl Conf on Production Research*, Valparaiso, Chile, 2007, pp. 1-6.
- [4] Eckermann, B. and Lin, G., "Setting the Focus for Innovation," in *Proc. of the 2000 IEEE Intl Conf on Management of Innovation and Technology*, 2000, pp. 283-288.
- [5] Kusiak, A., "Innovation: The Living Laboratory Perspective," *Computer-Aided Design and Applications*, Vol. 4, No. 6, pp. 863-876, 2007.
- [6] Chengde, Z., "Theory and Design Process of TRIZ," *Journal of Industrial Engineering and Engineering Management*, Vol. 17, No. 1, pp. 61-72, 2003.
- [7] von Hippel, E. and Katz, R., "Shifting Innovation to Users via Toolkits," *Management Science*, Vol. 48, No. 7, pp. 821-833, 2002.
- [8] Bandini, S., Mauri, G., and Serria, R., "Cellular Automata: From a Theoretical Paralledl Computational Model to is Application to Complex Systems," *Parallel Computing*, Vol. 27, No. 5, pp. 539-553, 2001.
- [9] Chopard, B., "Cellular Automata and Lattice Boltzmann Techniques: An Approach to Model and Simulate Complex Systems," *Advances in Complex Systems*, Vol. 5, No. 2/3, pp. 103-246, 2002.
- [10] Schleiffer, R., "An Itelligent Agent Model," *European Journal of Operations Research*, Vol. 166, No. 3, pp. 666-693.
- [11] Koritarov, V., "Real-World Market Representation with Agents," *IEEE Power and Energy Magazine*, Vol. 2, No. 4, pp. 39-46, 2004.

- [12] Tesfatsion, L., "Agent-based Computational Economics: Modeling Economies as Complex Adaptive Systems," *Information Sciences*, Vol. 149, No. 4, pp. 262-268, 2003.
- [13] Holland, J., *Hidden Order: How Adaptation Builds Complexity*. New York: Helix Books, 1995.
- [14] Miller, J. and Page, S., *Complex Adaptive Systems, An Introduction to Computational Models of Social Science*. Princeton, NJ: Princeton University Press, 2007.
- [15] Fleming, L. and Sorenson, S., "Technology as a Complex Adaptive SystemL Evidence from Patent Data," *Research Policy*, Vol. 30, No. 7, pp. 1019-1039, 2001.
- [16] Fullbright, R., "Teaching Critical Thinking Skills in IT using PINE-TRIZ," in *Proc. of the 5th Conf. on Information Technology Education*, Salt Lake City, UT, 2004, pp. 38-42.
- [17] Berkhout, G. and van der Duin, P., "Modile Data Innovation: Lucio and the Cyclic Innovation Model," in *Proc. of the 6th Intl. Conf. on Electronic Commerce*, 2004.
- [18] Sawatani, Y., Nakamura, F., Skskibara, A., Hoshi, M., and Masuda, A., "Innovation Paterns," in *Proc. of the 2007 IEEE Intl. Conf. on Services and Computing*, 2007.
- [19] Savioz, P. and Snnerman, E., "The Concept of the Integrated Innovation Process," in *Proc. of Technology and Innovation Management PICMET'99*, Vol. 2, Portland, OR, 1999, pp. 137-143.
- [20] Collins, L., "Opening up the Innovation Process," *Engineering Management*, Vol. 16, No. 1, pp. 14-17, 2006.
- [21] Zhang, J. and Tao, Y., "The Interaction Based Innovation Process of Architectural Design Service," in *Proc. of the 2007 IEEE Intl. Conf. on Idustrial Engineering and Engineering Management*, 2007, pp. 1719-1723.
- [22] Dell Corporation, "Cutomer-Driven Innovation," website:
http://www/dell.com.content/topics/global.aspx/about_dell/company/innovation.
- [23] Ulwick, A., "Turn Customer Input Into Innovation," *Harvard Business Review*, Vol. 80, No. 1, pp. 91-97, 2002.

- [24] Kusiak, A., "Innovation: From Data to Knowledge," BoneZOnline 2007.
- [25] Kline, S. and Rosenberg, N., "An Overview of Innovation, The Positive Sum Strategy: Harnessing Technology for Economic Growth," *National Academy Press*, pp. 275-306, 1986.
- [26] Tung, W. and Yaun, S., "iDesign: An Intellieght Design Framework for Service Innovation," , 2007.
- [27] Cjen, Y., Hsieh, P., and Lee, C., "An Intefrated Framework for Managing Service Innovation," in *Proc. of the 2007 Portland Intl. Conf. in Management Engineering and Technology*, Portland, OR, 2007.
- [28] Xi-song, L., Zhi, Y., and Xing-jian, X., "Discernment of Key Evaluating Elements of Independent Innovation Ability of Defense Science and Technology Industry," in *Proc. of the 2006 Intl. Conf. on Management Sciences and Engineering*, 2006, pp. 1965-1969.
- [29] Felfernig, A., Russ, C., and Wundara, M., "Toolkits Supporting Open Innovation in E-Government," in *Proc. of the 2004 Intl Conf. on Enterprise Information Systems*, 2004, pp. 296-302.
- [30] Blazevic, V. and Lievens, A., "Managing Innovation Through Customer Coproduced Knowledge in Electronic Services," *J. of the Acad. of Marketing Science*, Vol. 36, pp. 138-151, 2008.
- [31] Segaran, T., *Programming Collective Intelligence*. Sebastopo, CA: O'Reilly, 2007.
- [32] Netflix,., *The Netflix Challenge*.: Website: <http://www.netflixprize>.
- [33] WordNet,., *WordNet API*.: Website: <http://wordnet.princeton.edu/links>.
- [34] Hu, M. and Liu, B., "Mining Opinion Features in Customer Reviews," in *Proc. of the 19th National Conf. on Artificial Intelligence*, San Jose, CA, 2004.
- [35] Shih, M., Liu, D., and Hsu, M., "Mining Changes in Patent Trends for Competitive Intelligence," in *Proc. of PAKDD 2008*, 2008, pp. 681-690.
- [36] Wen, G., "Rotating Dynamics for Computational Creativity," *National Defense Industry Press*, 2008.

- [37] Wen, G., "Generating Creative Ideas Through Patents," , 2006, pp. 681-690.
- [38] Google, *Google Patent Search*.: Website: <http://www.google.com/patents>.
- [39] Tan, P., Stienbach, M., and Kumar, V., *Introduction to Data Mining*.: Addison Wesley, 2006.
- [40] Han, J. and Kamber, M., *Data Mining Concepts and Techniques*, 2nd ed.: Morgan Kauffman, 2006.
- [41] Bruynooghe, M., "A Framework for the Abstract Interpretation of Logic Problems," Katholieke Universiteit Leuven, 1987.
- [42] Muthukumar, K. and Hermenegildo, M., "Compile-time Derivation of Variable Dependency Using Abstract Interpretation," *Journal of Logic Programming*, Vol. 7, July 1992.
- [43] Luby, M., Mitzenmacher, M., and Shokrallahi, M., "Analysis of Random Processes via And-Or Tree Evaluation," in *Proc. of the Ninth Annual ACM-CIAM Symp. On Discrete Algorithms*, San Francisco, CA, 1998.
- [44] Brabazon, A. et al., "Agent-based Modeling of Product Invention," in *Proc. of the ACM Genetic and Evolutionary Computation Conf.*, New York, 2005, pp. 129-136.
- [45] Windrum, P. and Birchenhall, C., "Is Product Life Cycle a Special Case? Dominant Designs and the Emergence of Market Niches Through Coevolutionary-learning," pp. 109-134, August 1998.
- [46] Rogers, E., Medina, U., Rivera, M., and Wiley, C., "Complex Adaptive Systems and the Diffusion of Innovations," *The Innovation Journal: The Public Sector Innovation Journal*, Vol. 10, No. 3, pp. 579-584, June 2005.
- [47] Kauffman, S., Lobo, J., and Macready, W., "Optimal Search on a Technology Landscape," *Journal of Economic Behavior and Organization*, Vol. 43, No. 2, pp. 141-166, October 2000.
- [48] Kauffman, S., *The Origins of Order*. New York: Oxford University Press, 1993.

- [49] Stephens, C., "Effective Fitness Landscapes for Evolutionary Systems," in *Proc. of the IEEE Congress on Evolutionary Computation*, Washington D.C, 1999, pp. 703-714.
- [50] Skellette, B., Cairns, B., Geard, N., Tonkes, B., and Wiles, J., "Maximally Rugged NK Landscapes Contain the Highest Peaks," in *Proc. of the 2005 Conf. on Genetic and Evolutionary Computation*, 2005, pp. 579-584.
- [51] Wiese, T., Niemczyk, S., Skuch, H., Reichle, R., and Geihs, K., "A Tunable Model for Multi-Objective, Epistatic, Rugged, and Neutral Fitness Landscapes," in *Proc. of the ACM Genetic and Evolutionary Computation Conf.*, Atlanta, GA, 2008, pp. 795-802.
- [52] Kauffman, S. and Johnsen, S., "Coevolution to the Edge of Chaos: Coupled Fitness Landscapes, Poised States and Coevolutionary Avalanches," in *Artificial Life II*. Redwood City, CA: Addison Wesley, 1992, pp. 325-368.
- [53] Gell-Mann, M., "What is Complexity," *Complexity*, Vol. 1, No. 1, pp. 16-19, 1995.
- [54] Standish, R., "On Complexity and Emergence," *Complexity International*, Vol. 9, 2001.
- [55] Holland, J., *Hidden Order: How Adaptation Builds Complexity*. New York: Helix Books, 1995.
- [56] Stogatz, S., *Nonlinear Dynamics and Chaos*. New York: Westview Press, 1994.
- [57] Engler, J. and Kusiak, A., "Mingling the Requirements for Innovation," *Mechanical Engineering*, Vol. 130, No. 11, pp. 38-40, 2008.
- [58] Langton, C., "Computation at the Edge of Chaos: Phase Transitions and Emergent Computation," *PhysicaD*, Vol. 42, pp. 12-37, 1990.
- [59] Paap, J. and Katz, R., "Anticipating Disruptive Innovation," *Research Technology Management*, Vol. 47, No. 5, pp. 13-22.
- [60] Goldenberg, J. and Efroni, S., "Using Cellular Automata Modeling of the Emergence of Innovations," *Technological Forecasting and Social Change*, Vol. 68, No. 3, pp. 293-308, 2001.

- [61] Kun, F., Kocis, G., and Farkas, J., "Cellular Automata for the Spreading of Technologies in Socio-economic Systems," *PhysicaA*, Vol. 383, No. 2, pp. 660-670, 2007.
- [62] Ma, T. and Nakamori, Y., "Agent-based Modeling on Technological Innovation as an Evolutionary Process," *European Journal of Operational Research*, Vol. 166, No. 3, pp. 741-755, 2005.
- [63] Andriani, P., Conti, F., Fortuna, L., Frasca, M., Passiante, G. and Rizzo, A. "Innovation Systems by Nonlinear Networks," *Nonlinear Dynamics*, Vol. 44, No. 1-4, pp. 263-268, 2006.
- [64] Gilber, N., Pyka, A., and Ahrwieler, P., "Innovation Networks: A Simulation Approach," *Journal of Artificial Societies and Social Simulation*, Vol. 4, No. 3, pp. 8-22, 2001.
- [65] Celada, F. and Seiden, P., "A Computer Model of Cellular Interactions in the Immune System," *Immunology Today*, Vol. 13, No. 2, pp. 56-62, 1992.
- [66] Clarke, K., Hoppen, S., and Gaydos, L., "A Self-modifying Cellular Automaton Model of Historical Urbanization in the San Francisco Bay Area," *Environment and Planning B*, Vol. 24, No. 2, pp. 247-261, 1992.
- [67] Conway, J., *What is Life?* San Diego, CA: Academic Press, Vol. 2.
- [68] Goldenberg, J., Libia, B., and Muller, E., "Using Complex Systems Analysis to Advance Marketing Theory Development: Modeling Heterogeneity Effects on New Product Growth through Stochastic Cellular Automata," *Acad. of Marketing Science Review*, Vol. 20, No. 9, pp. 211-223, 2001.
- [69] Kusiak, A., "Innovation: A Data Driven Approach," *Drug Development 2007, Touch Briefings*, pp. 55-58, 2007.
- [70] Packard, N. and Wolfram, S., "Two-Dimensional Cellular Automata," *Journal of Statistical Physics*, Vol. 38, No. 5-6, pp. 901-946, 1985.
- [71] Kenedy, J. and Mendes, R., "Neighborhood Topologies in Fully-Informed and Best-of-Neighborhood Partical Swarms," in *Proc. of the IEEE SMC Workshop on Soft Computing in Industrial Applications*, 2003.

- [72] Sosa, R. and Gero, J., "Social Structures that Promote Change in a Complex World: The Complementary Roles of Strangers and Acquaintances in Innovation," *Futures*, Vol. 40, No. 6, pp. 577-585, 2008.
- [73] Ratna, N., Dray, A., Perez, P., Grafton, R., and Kompas, T., "Innovation Diffusion Amongst Heterogeneous Agents: Exploring Complexity with Agent-Based Modeling," in *Applications of Complex Adaptive Systems*. Hershey, NY: IGI Publishing, 2008, p. Ch. 5.
- [74] Schwarz, N. and Ernst, A., "Agent-Based Modeling of the Diffusion of Environmental Innovations - An Empirical Approach," *Technological Forecasting and Social Change*, Vol. To Appear (available online June 2008), 2008.
- [75] Eiben, A., Raue, P., and Ruttkay, Z., "Genetic Algorithms with Multi-Parent Recombination," in *Proc. of the 3rd Annual Conf. on Parallel Problem Solving from Nature*, 1994.
- [76] Lis, J. and Eiben, A., "A Multi-Parent Genetic Algorithm for Multiobjective Optimization," in *Proc. of the 4th IEEE Conf. on Evolutionary Computation*, Indianapolis, IN, 1997, pp. 59-64.
- [77] Isaacs, A., Ray, T., and Smith, W., "Set Representation and Multi-parent Learning with an Evolutionary Algorithm for Optimal Design of Trusses," *Evolutionary Computation*, pp. 419-439, 2008.
- [78] Ting, C., "An Analysis of the Effectiveness of Multi-parent Crossover," in *Parallel Problem Solving from Nature VIII*. Berlin: Springer-Verlag, 2004, Vol. 3242 of LNCS, pp. 131-140.
- [79] Eberhart, R. and Shi, Y., *Computational Intelligence: Concepts to Implementations*. Burlington, MA: Morgan Kaufman, 2007.
- [80] Koritarov, V., "Real-World Market Representation with Agents," *IEEE Power and Energy Magazine*, Vol. 2, No. 4, pp. 39-46, 2004.
- [81] Harp, S., Brignone, S., Wollenberg, B., and Samad, T., "SEPIA: A Simulator for Electric Power Industry Agents," *IEEE Control Systems Magazine*, Vol. 20, No. 4, pp. 53-69, 2000.

- [82] Babaoglu, O., Meling, H., and Montresor, A., "Anthill: A Framework for the Development of Agent-based Peer-to-Peer Systems," in *Proc. of the 22nd Intl Conf on Distributed Computing Systems*, Vienna, Austria, 2002.
- [83] Wooldridge, M., "Agent-based Software Engineering," *IEE Proceedings on Software Engineering*, Vol. 144, No. 2, pp. 26-37, 1997.
- [84] Axelrod, R., *The Complexity of Cooperation: Agent-based Models of Competition and Collaborations*. Princeton, NJ: Princeton University Press, 1997.
- [85] Fleming, B., "Endangered Technologies: Automotive Electronics," *IEEE Vehicular Technology Magazine*, Vol. 2, No. 3, pp. 58-60, 2007.
- [86] Marketwatch.com,., *news article.*: website:
<http://www.marketwatch.com/news/story/siemens-infineon-establish-semiconductor-joint/story.aspx?guid=>
- [87] sustainablebusiness.com,., *news article.*: website:
<http://www.sustainablebusiness.com/index.cfm/go/news.display/id/15950>.
- [88] ballsemi.com,., *press release.*: website:
http://www.ballsemi.com/corp/p_20060719.html.