

Imperfect Evolutionary Systems

Graham Kendall, *Member, IEEE*, and Yan Su

Abstract—In this paper, we propose a change from a perfect paradigm to an imperfect paradigm in evolving intelligent systems. An imperfect evolutionary system (IES) is introduced as a new approach in an attempt to solve the problem of an intelligent system adapting to new challenges from its imperfect environment, with an emphasis on the incompleteness and continuity of intelligence. We define an IES as a system where intelligent individuals optimize their own utility, with the available resources, while adapting themselves to the new challenges from an evolving and imperfect environment. An individual and social learning paradigm (ISP) is presented as a general framework for developing IESs. A practical implementation of the ISP framework, an imperfect evolutionary market, is described. Through experimentation, we demonstrate the absorption of new information from an imperfect environment by artificial stock traders and the dissemination of new knowledge within an imperfect evolutionary market. Parameter sensitivity of the ISP framework is also studied by employing different levels of individual and social learning.

Index Terms—Artificial intelligence, environmental variables, evolutionary computation (EC), imperfect evolutionary systems (IESs), individual learning, social learning.

I. INTRODUCTION

IT HAS BEEN a long-term aim for humans to build intelligent systems or humanoids just like ourselves. One of the important features of being a human is the ability to learn new things, i.e., the ability to accept new information from the environment and use it to update our existing knowledge. There have been advances in creating such intelligence in recent years. One of the notable works is *BLONDIE24* developed by Chellapilla and Fogel [1] and Fogel [2]. *BLONDIE24* is an evolutionary program that taught itself to play checkers from scratch, without the input of human expert knowledge. *BLONDIE24* developed its own game playing strategies to the level of a human expert. The key advance in *BLONDIE24* is that the system did not use any preprogrammed human expertise, which was the significant difference between *BLONDIE24*, *Chinook* [3], [4], and *Deep Blue* [5], [6]. *BLONDIE24* was a major step forward in answering Samuel's challenge about machine learning without pre-injected knowledge [7]. However, *BLONDIE24* is still a perfect end product, rather than adaptive intelligence. Fogel [2, p. 205] commented on *BLONDIE24* as follows:

“... Remember that the graph doesn't show the neural network learning during the one hundred games. Everything that it learned, it had learned during its evolution.

Manuscript received December 1, 2005; revised May 3, 2006.

The authors are with the Automated Scheduling, Optimization and Planning (ASAP) Research Group, School of Computer Science and Information Technology, University of Nottingham, Nottingham NG8 1BB, U.K. (e-mail: gsk@cs.nott.ac.uk; yxs@cs.nott.ac.uk).

Digital Object Identifier 10.1109/TEVC.2006.887348

The neural network was the same for game one hundred as it was for game one ...”

Therefore, *BLONDIE24* has not yet fully answered the challenge of learning with new information when the environment evolves. As commented by Harley in his book review [8]:

“... An interesting point is that the end product which looks intelligent is Blondie, yet she is not in fact the intelligence. Like the individual wasp, Blondie is fixed in her responses. If she played a million games, she would not be one iota smarter. In this sense, she is like Deep Blue. ... Perhaps a better example of intelligence would be ... a human, who can adapt her behavior to any number of new challenges ...”

Stanley *et al.* [9] developed another type of real-time computer game, called NERO, where artificial agents are able to improve and change themselves during the game. In NERO, each player trains a team of virtual robots which engage with other robots designed by other players. As the game is being played, a real-time neuroevolution method is employed to constantly evolve increasingly complex artificial neural networks (ANNs) which the virtual robots use to construct fighting strategies. NERO demonstrates the possibility of artificial intelligence being adaptive in real time, i.e., being able to learn and improve itself, while interacting with the environment. However, note that the training of robots in NERO is controlled by human players. The designers of NERO decide what inputs should be used to construct the ANNs that the virtual robots use to fight. In other words, the artificial robots in NERO have the option to improve themselves based on what they are given, but they do not have the ability to decide what they should perceive from their environments. If one of the robot teams developed a new kind of weapon that could see through solid obstacles, other teams that still use the old set of inputs, are likely to be defeated.

We argue that the fundamental problem lies in the ignorance of an intelligent entity's relationship to its environment. More precisely, it is a relationship to an *imperfect environment*. By saying an environment is imperfect, we mean that the environment is constantly changing and evolving as new information and knowledge emerge over time. For example, in the game-playing world, the problem facing a checkers player is not only different opponents, with different skills and characteristics, but also different types of games. A checkers player may play a good game and ask himself what he can learn from the game. The same player may also watch someone playing poker and may want to learn to play poker. The same thing happens to scientific researchers. The environment that a researcher is facing is also evolving where new algorithms, techniques, software, and hardware are continuously introduced into the

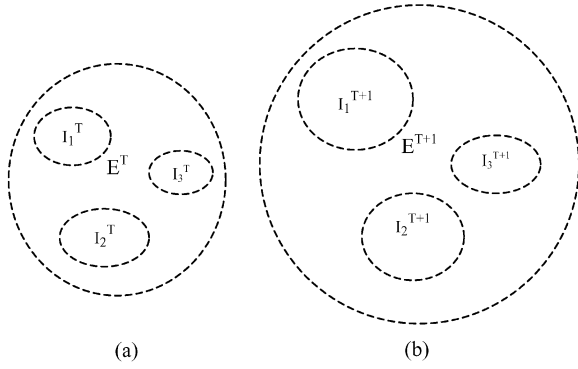


Fig. 1. Imperfect evolutionary system.

environment. In order to answer the new challenges from their research field, researchers need to constantly improve themselves. Therefore, an imperfect environment is an environment that constantly presents new challenges to its participants.

Inside an imperfect environment, there exist *imperfect individuals*. By saying individuals are imperfect, we mean individuals do not completely understand their environment due to the complexity of the environment and the fact that every individual has a (possibly) different view of the world. For example, a game player may only know certain strategies when playing checkers and does not know how to play poker at all. A researcher who investigates genetic algorithms may have no idea how to design an artificial neural network (ANN). A stock investor may know how to use price moving-averages to predict future movement of a market but may not understand the significance of trading volumes. It is not necessary for individuals to be of perfect intelligence, so long as what they know is enough for his/her survival. However, when new challenges are placed upon an individual by its imperfect environment, individuals may choose to adapt to the new challenge, e.g., by learning a new skill, and hence have a better chance of competing in the future.

Together, imperfect individuals existing in an imperfect environment form an imperfect evolutionary system (IES). Imperfect individuals constantly learn new things in order to adapt to the new challenges posed by the imperfect environment. The imperfect environment constantly evolves because things that were previously unknown become known and the unobservable become observable. We define an IES as a system where intelligent individuals optimize their own utility, with available resources, while adapting themselves to the new challenges from an evolving and imperfect environment.

Diagrammatically, we depict an IES in Fig. 1. An imperfect environment at time T , E^T , represented as the large circle (dashed), is defined as a set of information i_n^T , as shown in (1)

$$E^T = \{i_1^T, i_2^T, i_3^T, \dots, i_n^T\}. \quad (1)$$

We define each i_n^T as an *environmental variable* v_n^T ; thus, an imperfect environment E^T can also be represented as a set of environmental variables, as shown in (2), where n is the number of environmental variables in E^T

$$E^T = \{v_1^T, v_2^T, v_3^T, \dots, v_n^T\}. \quad (2)$$

Due to their imperfectness, imperfect individuals only partially understand their environment. Every individual (I_i^T) in E^T , represented as small circles (dashed) in Fig. 1, perceives a different set of information from its environment. Therefore, every individual has a different observational space of the environment, as defined in (3). We use dashed lines for both the environment and the individuals in Fig. 1 to indicate that none of them are finite and static

$$I_i^T = \{v_j^T, \dots, v_k^T \mid j, k \in n\}. \quad (3)$$

We use two examples to further illustrate E^T and I_i^T in Fig. 1. For the first example, we use the game playing problem discussed above, where *BLONDIE24* is unable to learn from its environment. Assume Fig. 1 represents a particular part of a human brain that is responsible for game playing. We call it a *game-playing mind*. The imperfect environment, E^T in this case, is then a set of games. At time T , all people that surround this person/brain only play checkers. E^T in Fig. 1(a) is shown as in (4), where v_i^{checkers} is a particular feature used for playing checkers, such as number of pieces on the game board

$$E^T = \{v_1^{\text{checkers}}, v_2^{\text{checkers}}, v_3^{\text{checkers}}\}. \quad (4)$$

Individual I_1^T in Fig. 1(a) is a game strategy that makes use of v_1^{checkers} and v_2^{checkers} . Individual I_2^T is another checkers game strategy that uses v_1^{checkers} and v_3^{checkers} . I_3^T does nothing, as shown in (5)

$$\begin{aligned} I_1^T &= \{v_1^{\text{checkers}}, v_2^{\text{checkers}}\} \\ I_2^T &= \{v_1^{\text{checkers}}, v_3^{\text{checkers}}\} \\ I_3^T &= \{\}. \end{aligned} \quad (5)$$

From time T to time $T+1$, this person/brain plays checkers with people using different features and strategies, and even has the opportunity to meet people who play poker. The imperfect environment of the game-playing mind evolves to E^{T+1} in Fig. 1(b), as shown by (6)

$$E^{T+1} = \{v_1^{\text{checkers}}, v_2^{\text{checkers}}, v_3^{\text{checkers}}, v_4^{\text{checkers}}, v_1^{\text{poker}}, v_2^{\text{poker}}\}. \quad (6)$$

During time T to $T+1$, individual I_1^T and individual I_2^T compete so that better strategies will survive and bad strategies will be discarded (or, in some sense, forgotten by the brain). In the case of Fig. 1(b), both strategies (I_1^T and I_2^T) survive. The individual I_1^T survives by learning to use a new feature v_4^{checkers} , as shown in (7), while individual I_2^T survives by refining itself (if the game strategy is implemented in the form of neural networks, refining means the tuning of weights or changing the network architecture). Individual I_3^T also adapts to the evolving environment by learning to play poker

$$\begin{aligned} I_1^{T+1} &= \{v_1^{\text{checkers}}, v_2^{\text{checkers}}, v_4^{\text{checkers}}\} \\ I_2^{T+1} &= \{v_1^{\text{checkers}}, v_3^{\text{checkers}}\} \\ I_3^{T+1} &= \{v_1^{\text{poker}}, v_2^{\text{poker}}\}. \end{aligned} \quad (7)$$

For the second example, we use the stock market where traders use different market indicators for constructing their trading strategies and evolve while new trading techniques constantly emerge from the market. At time T , the imperfect evolutionary market only provides two market indicators in E^T , as shown in (8), where MA is the moving average of stock prices and *volume* are the trading volumes (*price* is the primitive stock prices)

$$\begin{aligned} E^T &= \{\text{price, MA, volume}\} \\ I_1^T &= \{\text{price, MA}\} \\ I_2^T &= \{\text{price, volume}\} \\ I_3^T &= \{\text{price}\}. \end{aligned} \quad (8)$$

At time $T + 1$, a new market indicator, rate of change (ROC), is introduced to the market. The imperfect market evolves, as shown in (9)

$$\begin{aligned} E^{T+1} &= \{\text{price, MA, volume, ROC}\} \\ I_1^{T+1} &= \{\text{price, volume, MA, ROC}\} \\ I_2^{T+1} &= \{\text{price, MA, ROC}\} \\ I_3^{T+1} &= \{\text{price, ROC}\}. \end{aligned} \quad (9)$$

Obviously, there exist a number of features for playing checkers or poker, and there are more complicated market indicators for trading stocks in markets. We only use a few of them in the above example for the purposes of illustration. It is also the case that the second example of an imperfect evolutionary market is less complicated compared with the imperfect game-playing mind. This is because in the game playing system, imperfect individuals will have multiple objectives, i.e., playing checkers or playing poker, while in the imperfect evolutionary market stock traders only have one objective, i.e., making a profit. Therefore, in this paper, as a first attempt in developing IESs, we choose the stock market as a testbed to study the evolutionary learning of imperfect individuals in an imperfect environment.

Another reason we choose the stock market as the first testbed for developing IESs is that the financial markets have been intensively studied as imperfect information systems. The study of information economics stresses a change from the conventional perfect information paradigm to an imperfect information paradigm and has demonstrated the profound effect it has had on people's understanding of many economic problems (please refer to Stiglitz [10] and Arnott *et al.* [11] for comprehensive discussions on information in economic studies). It is, in fact, this change, i.e., from a perfect paradigm to an imperfect paradigm, in economic studies that has inspired us to rethink artificial intelligence from an imperfect perspective.

In the next section, we will discuss IESs in the context of artificial intelligence research and evolutionary learning. In particular, we discuss cultural algorithms that are similar to the individual and social learning paradigms (ISPs) proposed in this paper for the implementation of IESs. In Section III, we describe the integrated ISP. Through the comparison between cultural algorithms and the ISP learning paradigm, we further explain our idea of IESs. Section VI describes the implementation of an imperfect evolutionary market based on the ISP learning

paradigm. Experimental studies and discussions are presented in Section V. Through experiments, we demonstrate how imperfect artificial traders are able to learn new things and adapt to new challenges from an imperfect environment. Conclusions and future work are provided in Section VI.

II. RELATED WORK

There are a number of studies that have investigated uncertain and dynamic environments in artificial intelligence research, such as commonsense reasoning [12], [13], qualitative reasoning [14], probabilistic reasoning using Bayesian networks [15], and fuzzy logic [16], [17], etc. Each of these studies deals with the problem of modeling intelligence from a different perspective. Commonsense reasoning generally covers problems that are related to a human's ability to use common sense knowledge in daily life. Qualitative reasoning studies problems that can only be described qualitatively rather than quantitatively. Bayesian networks provide the means to reason with probabilities and fuzzy sets make it possible to describe the vagueness of a problem. In this paper, we propose the concept of an IES with an emphasis on how to sustain the continuous innovation of an intelligent system. For example, in the context of Bayesian reasoning, the problem for a checkers player would be, "If my opponent has a 30% probability to move to A and 70% probability to move to B, what should I do?" On the other hand, if we view the checkers player as an IES, the problem would be "What can I learn from this opponent, or how do I learn to play this new game?" If Bayesian networks are more about imprecision and probabilities, and fuzzy logic is more about vagueness and possibilities, we would say IESs are more concerned with incompleteness and continuity, i.e., the imperfectness of intelligence.

Evolutionary computation (EC) [18]–[20] provides an alternative way to study adaptive intelligence. In [19, p. 245], the author argues that evolution serves as a unifying description of all intelligent processes. Every intelligent learning system in nature adopts a functionally equivalent process of reproduction, variation, competition, and selection. Intelligent systems evolve and adapt their behaviors to achieve goals in a range of environments. Evolutionary learning is also unique in the sense that it not only recapitulates human behaviors, but also generates new behaviors and new solutions that are not currently known to humans. Fogel *et al.* [21] have recently guest edited a Special Issue of the IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION which presents advances in developing adaptive intelligence in computer games using EC techniques.

Cultural algorithms [22], [23] are evolutionary learning models that emulate cultural evolutionary processes. A cultural algorithm consists of two levels of evolution: the microevolution in a population space and the macroevolution in a belief space. Through an acceptance function, the experiences of individuals in the population space are used to generate problem solving knowledge that is to be stored in the belief space. The belief space manipulates the knowledge, which in turn guides the evolution of the population space by means of an influence function. Sternberg and Reynolds [24] use an evolutionary learning approach based on cultural algorithms to learn about the behavior of a commercial rule-based system for fraud detection. The learned knowledge in the belief space




Fig. 2. An integrated ISP.

of the cultural algorithm is then used to re-engineer the fraud detection system. Cultural algorithms have also been used for modeling the evolution of complex social systems [25], [26].

Reynolds and Saleem [27], Reynolds and Peng [28], and Reynolds and Saleem [29] describe the application of cultural algorithms for function optimization problems in dynamic environments. In their experiments, the dynamic environment is modeled as a two-dimensional (2-D) plane with four cones of different heights and slopes randomly scattered on it. The four cones shift their locations on the plane at certain generations, so that the location of the optimum solution constantly changes. Reynolds and Saleem [27] show that the cultural algorithm is more effective than an evolutionary algorithm with only one single-level evolution when they are applied to the problem of finding the new optima in dynamic environments. Reynolds and Peng [28] discuss how the learning of knowledge in the belief space ensures the adaptability of cultural algorithms. Reynolds and Saleem [29] further investigate the contributions of different types of knowledge from the belief space in guiding the search toward the best solutions in both deceptive and nondeceptive environments.

Note that the dynamic environmental change described in the above experiments is only caused by the change of values of existing environmental variables, i.e., the heights, slopes, and locations of cones. The dynamic environment in these experiments is not a result of the emergence or the introduction of new environmental variables over time. The dimension of the problem observational space remains unchanged, i.e., it remains a 2-D optimization problem. Other applications of cultural algorithms in dynamic environments, with multi-agent-based population space, are carried out in a similar way. Reynolds and Ostrowski [30] combine the cultural algorithm and agent-based computational economics for the evolution of successful pricing strategies in an original equipment manufacturer-consumer market. The market is a dynamic environment in the sense that a theta parameter, which indicates an economic recession, is applied for various lengths of time.

The *new challenge* we discussed in Section I not only indicates that the status, e.g., the values of current existing environment variables could change, but there is also the possibility of the emergence of new environment variables that presents additional problems for an intelligent entity. In other words, an

imperfect individual's observational space, and therefore its related problem space and problem domain, could both change in an IES (see our examples of the game-playing mind and the imperfect evolutionary market in Section I). In this paper, for the implementation of an imperfect evolutionary market, we introduce an integrated ISP that is described in next section.

III. AN INTEGRATED INDIVIDUAL AND SOCIAL LEARNING PARADIGM (ISP)

Fig. 2 schematically depicts an integrated ISP, which consists of four building blocks. Under the definition of IESs given in Section I, individuals of such a system absorb information from their environment and transfer the information into their own knowledge. The knowledge will then be refined through individual learning and disseminated within the society of imperfect individuals through social learning. The four building blocks of an IES are the following.

The Imperfect Environment (E): E is as defined by (2). The imperfect environment is the core of the IES. It has two functions: the supplier of resources for the evolution and the medium for the evolution. The environment provides its participants with the information and knowledge that they need for their survival. As an imperfect environment, it too also evolves when new information and knowledge become available, as shown in Fig. 1. The imperfect environment also acts as the medium where evolution occurs through the storage and dissemination of information and knowledge by means of social learning.

The Imperfect Individuals (I_i): I_i is defined by (3). An imperfect individual has two tasks. One is to optimize its own utility, with the available resources, through an individual learning process. The other task is to adapt to the changing environment by the absorption of new emerging information and knowledge from the imperfect environment, through a social learning process.

Individual Learning Mechanism (IL): The individual learning mechanism defines the evolutionary process of an individual's optimization of its own utility. *Social Learning Mechanism (SL):* Social learning defines the evolutionary process where participants of an imperfect environment learn from each other, while the information and knowledge is disseminated within the IES.

In Fig. 2, imperfect individuals (I_i) are represented as dashed rectangles of different sizes. This indicates that the observational spaces of individuals are of different dimensions and nonstatic, depending on the information sets used by individuals. The social learning mechanism (SL) works as the channel through which individuals will learn from each other so that information and knowledge, including new information from the evolved environment, is disseminated among the evolutionary intelligent entities. There are similarities between the integrated ISP and cultural algorithms. For example, they both evolve at the individual level and they both maintain a knowledge repository. However, there are some fundamental differences between these two learning paradigms.

- **Imperfect environment**—The ISP learning paradigm stresses the importance of the awareness of an imperfect environment by an intelligent agent. An imperfect environment is always changing and evolving. Things that were previously unobservable become observable. Things that were previously unknown become known. An imperfect environment does not only mean possible changes in quantity, but also the potential changes in quality. The concept of an imperfect environment is not reflected in the design of cultural algorithms.
- **Imperfect individuals**—More importantly, the design of an IES is dependent on the development of imperfect individuals. It is only when an intelligent entity considers itself to be imperfect, rather than a perfect end product, that it has the opportunity and motivation to learn and adapt to the new challenges from the environment. Imperfect individuals also imply that each individual can only understand its own world imperfectly and each one has a (possible) different view of the world. In terms of problem solving, this implies the problem observational spaces of individuals in an IES could have different dimensions because they use different sets of information, and even different problem domains, because they are looking at the world in different ways. In cultural algorithms, there is no concept of imperfect individuals.
- **Imperfect world**—Compared with other evolutionary learning paradigms in AI, IESs deal with learning problems from a different point of view. For example, imagine we are building a simulated brain (we do not care if it works as the way the human brain does, since we do not yet know, exactly, how the brain works). If we employ a cultural learning algorithm approach, the first thing we need to do is to create a population space e.g., using multipopulation genetic algorithms/genetic programming (MGA/MGP). Each subpopulation/individual in the MGA/MGP does a specialized task. The cultural algorithm then evolves the population space and updates the belief space. The question we now need to ask is how many subpopulations we need create in order to cover all the functions performed by the brain? There is no easy answer to this question. Also, how does the cultural learning system learn new things from its environment; by creating a new subpopulation with a different problem domain, or by modifying an existing subpopulation? Unfortunately, neither of these two mechanisms exists in cultural algorithms. On the other hand, if we employ an imperfect

evolutionary methodology, we can start with just one individual that basically does nothing (a group of dummy neural networks). When a new environmental variable emerges in the environment, the imperfect evolutionary “mind” perceives it and decides either to incorporate it into current networks in the ISP, or create a new individual for this piece of new information, or simply delete the new information. We stress the change from a perfect perspective where intelligence is viewed as a perfect end product, to an imperfect perspective where adaptability is the nature of intelligence in an imperfect world.

In this paper, as a first attempt in developing IESs, we implement a simple imperfect evolutionary market, where each artificial stock trader represents (possibly) different problem domains depending on the information set the trader is using. We will demonstrate how the ISP learning paradigm enables artificial traders to adapt to new environmental variables emerging from the imperfect market by experimental studies.

IV. AN IMPERFECT EVOLUTIONARY MARKET

The stock market has been extensively studied as a complex evolutionary system [31]–[35]. In [31], Palmer *et al.* point out that evolutionary stock market models can be used as a fertile testbed for exploring, not only financial markets, but also adaptive agents and a class of artificial lives. The imperfect evolutionary market model described in this paper extends our previous studies in [36]–[38], where we investigated the adaptive behaviors of imperfect traders in a simulated stock market. In this paper, we implement an imperfect market environment by means of gradually injecting new market indicators into the simulated stock market.

The implementation of the imperfect evolutionary market consists of four parts, corresponding to the four building blocks of an IES described in the previous section: imperfect stock market, imperfect stock traders, an individual learning mechanism, and a social learning mechanism.

A. Imperfect Market

The fundamental market scenario, such as stock prices, market indexes, and trading volumes comprise 20 technical indicators that will be used by artificial traders in the simulated stock market for trading stocks. These 20 technical indicators are the environmental variables of the imperfect evolutionary market. Initially, the simulated market starts with ten indicators, whereas the other ten indicators are gradually injected into the simulated market approximately every two years. These gradually introduced technical indicators are the new challenges, in the form of new information, for traders. Obviously, it is not always true that every two years some new market indicators are invented in real-world markets. However, we have designed the experiments in this way for the purpose of simulating the emergence of new challenges in an imperfect environment.

The imperfect market also acts as the medium for the evolution of information (indicators) and knowledge (trading strategies) in the market. As well as trading, artificial traders will also publish their successful trading strategies to the market so that good trading strategies can be disseminated to other traders

within the market. The trader will also assign credit to the indicators and strategies they have used for successful trading so that useful information and knowledge will be accessible by the other artificial traders.

B. Imperfect Stock Traders

We use 50 artificial stock traders in the simulated stock market. Each trader selects a set of indicators from the market for creating their trading models, which are in the form of multilayer feedforward perceptrons [36]–[38]. The imperfect traders evolve (hopefully) better trading models through individual learning. The imperfect traders also search for better information and learn from the other traders in the market through social learning. Initially, we give each trader £10 000 cash to start trading.

C. Individual Learning Mechanism

The individual learning mechanism in our experiments is implemented using an evolutionary programming algorithm [18] applied to evolve ANN models used by individual traders (see [39] and [40] for comprehensive surveys on evolutionary ANNs). The feedforward ANNs, with zero or one hidden layers, take technical indicators as inputs. Each trader will create ten random ANNs using the set of indicators the trader has chosen. The aim of individual learning is to search for the optimal relationship among the set of indicators that the trader has chosen by finding a suitable set of weights utilizing the evolutionary programming algorithm as a search mechanism. During the trading period, before every five days of trading, one model from the ten ANNs is selected to be used for making trading decisions. After every five days trading, the trader will calculate the selected model's rate of profit (ROP) by using (10), where W is the value of a trader's current assets (cash + shares). W' is the value of the trader's assets in the previous trading session

$$\text{ROP} = \frac{W - W'}{W'} \times 10. \quad (10)$$

The ROP describes a trading model's profitability. The selection process of the EP algorithm selects four ANNs with the lowest scores to be replaced by offspring produced by a mutation operation on other ANNs. The pseudocode of the individual learning algorithm is presented as follows; where m is a random integer between 0 and the total number of connections in the selected neural network; Δw is a random Gaussian number with a mean of zero and standard deviation of 0.1.

1. Select a model to be eliminated;
2. Select a model to be mutated using roulette wheel selection;
3. Decide number of connections to be mutated, m ;
4. $i = 0$;
5. While ($i < m$) {
 - Randomly select a connection;
 - Weight = weight + Δw ;
 - $i = i + 1$;
6. With 1/3 probability add a hidden node;
7. With 1/3 probability delete a hidden node;
8. Replace the model to be eliminated with the mutated model;

D. Social Learning Mechanism

The social learning mechanism in the imperfect market is implemented by a memory approach with a credit assignment scheme. The market works as a memory of all indicators, including newly introduced technical indicators, and the trading strategies developed by traders during the evolution. During the trading period, at the end of every 125 days trading, social learning occurs. The social learning mechanism accomplishes three main tasks.

- Traders self-assess their performance over the past six months (125-day) trading and assign credit to indicators and strategies that have been used in the past six months trading.
- Preserve successful trading strategies and disseminate them to the market.
- Traders detect new information that has recently been introduced to the market and disseminate the recently acquired information to others in the market.

The trader's self-assessment, at the end of 125 days trading, reflects how well the trader has performed in terms of his own profitability and relative performance when compared with others. The self-assessment process uses (11)–(13)

$$S_{\text{peer}}^i = \frac{R_i}{49}. \quad (11)$$

First, the trader's ROP for the past six months is calculated by using (10), and the 50 traders are ranked from 0 to 49 (R_i) according to their ROP. Equation (11) gives each trader a score in terms of peer pressure from other traders. In other words, this score shows trader i 's performance compared with other traders

$$S_{\text{self}}^i = \frac{\text{ROP} - \text{ROP}'}{100}. \quad (12)$$

In (12), ROP is a trader's rate of profit for the current six months trading. ROP' is a trader's rate of profit for the previous six months. Equation (12) gives the trader a score in terms of his own performance in the past six months compared with the previous six months. These two types of performance are then composed into (13), which gives the overall assessment for trader i

$$\text{assessment}_i = S_{\text{peer}}^i + \frac{1}{1 + e^{(1 - S_{\text{self}}^i)}}. \quad (13)$$

The final assessments (assessment_i) for the 50 traders are normalized into the range of [0, 1]. Depending on the value of each trader's assessment_i , there will be traders who publish their successful trading strategies to the market, traders who need to learn from other traders, and traders who search for new information from the environment. The artificial traders' social learning behaviors are presented below.

- 1) If a trader's assessment is 1, and the trader is not using a strategy drawn from a central pool, then publish the strategy into the pool. Go into the next six months trading using the same strategy.
- 2) If a trader's assessment is 1, and the trader is using a strategy copied from the pool, do not publish it again, but update this strategy's score in the pool using their six-month ROP. Go into the next six months trading using the same strategy.

TABLE I
FIVE SELECTED STOCKS TRADED IN THE SIMULATED STOCK MARKET

Company	Symbol	Market	Sector	Trading Period	Return (BnH)	Return (Bank)
CHENG KONG (CK)	0001.hk	HKEx	Real Estate	25/April/88 to 03/July/03 (3750 trading days)	410.46%	109.76%
CATHAY PAC AIR (CATHAY)	0293.hk	HKEx	Transports	25/April/88 to 03/July/03 (3750 trading days)	46.21%	109.76%
WHARF HOLDINGS (WHARF)	004.hk	HKEx	Conglomerate	25/April/88 to 03/July/03 (3750 trading days)	134.11%	109.76%
TOYOTA INDUS CORP (TOYOTA)	6201.jp	TSE	Automobiles	15/May/92 to 04/July/03 (2750 trading days)	122.6%	72.16%
SONY CORP (SONY)	6758.jp	TSE	Electronics and Electrical	15/May/92 to 04/July/03 (2750 trading days)	-16.93%	72.16%

HKEx refers to Hong Kong Stock Exchange. TSE refers to Tokyo Stock Exchange. Return refers to accumulated total return over the whole trading period; Return (BnH) = Cumulative total return from employing the Buy and Hold strategy (BnH) on that stock, Return (Bank) = Cumulative total return from investing all assets in bank savings over the same trading period (with an interest rate of 5%, paid annually).

TABLE II
ENVIRONMENTAL VARIABLES—20 TECHNICAL INDICATORS

Technical Indicator	Description	Technical Indicator	Description
IND1	10 days Moving Average (price)	IND11	14 days Relative Strength Index
IND2	20 days Moving Average (price)	IND12	21 days Relative Strength Index
IND3	50 days Moving Average (price)	IND13	Stochastic Oscillators (K%)
IND4	200 days Moving Average (price)	IND14	Fast Stochastics (D%)
IND5	Closing Price (normalised)	IND15	Slow Stochastics (slow D%)
IND6	Rate of Change (price)	IND16	Primary Market Index Rate of Change
IND7	Oscillator (price)	IND17	Relative performance to Primary Market Index
IND8	10 days bias (price)	IND18	Secondary Market Index Rate of Change
IND9	20 days volume Rate of Change	IND19	Relative performance to Secondary Market Index
IND10	10 days Relative Strength Index	IND20	Third Market Index Rate of Change

All values are normalised in the range of [0, 1] before being used as inputs to the neural networks. For all stocks, the Secondary Market Index refers to DJ INDU AVERAGE. For stocks from HKEX, the Primary Market Index refers to Hang Seng Index; the Third Market Index refers to NIKKEI 225 index. For stocks from TSE, the Primary Market Index refers to NIKKEI 225 index; the Third Market Index refers to Hang Seng Index.

- 3) If a trader's assessment is less than 0.9, the trader has a 0.5 probability of copying a strategy from pool, which means the trader will discard whatever model he is using, select a better trading strategy from the pool using roulette wheel selection, and go into the next six months trading with this copied strategy. Or, with a 0.5 probability, the trader will decide to discard whatever strategy he is using, select another set of indicators as inputs, build ten new models, and go into next six months of trading with these ten new models.
- 4) If the assessment is between 0.9 and 1.0, the trader is reasonably satisfied with his performance over the past six months. With 0.7 probability the trader will continue using the same set of indicators. With 0.3 probability the trader will decide to try a new set of indicators.

Traders will also update the scores of indicators they have used based on their performance over the past six months using (14) below

$$I_i^n = I_i^n + \text{ROP} \quad (14)$$

where i is the i th trader, n is the n th indicator used by trader i in the current six-month trading, and ROP is the rate of profit of trader i in the current six months trading. This mimics the public opinion of the technical indicators in the market.

V. EXPERIMENTAL RESULTS AND DISCUSSION

A. Imperfect Environment With New Challenges

1) *Absorption and Dissemination of New Information:* Table I lists the five stocks traded in the imperfect evolutionary market. Details of the 20 technical indicators are given in Table II. As discussed in Sections I and II, an imperfect environment constantly evolves and changes. Environmental variables that form an imperfect environment change when new information and knowledge are introduced or discovered. When new technical indicators are introduced into the simulated stock market, in the form of new environmental variables, the problem observational space of the artificial stock traders expands. The absorption and dissemination of new information

and knowledge among imperfect individuals becomes crucial to the evolutionary process. An evolutionary system becomes meaningless if it cannot perceive and digest new information from its environment.

The 20 indicators from Table II are divided into two groups: IND1–IND10 (the *static variables*) and IND11–IND20 (the *new challenges*). The group of ten indicators from IND1 to IND10 are given to the simulated stock market as initial settings, i.e., from the start of the evolutionary system. The remaining ten indicators are injected into the imperfect market as a new environmental variable every two generations, i.e., indicator IND11 is introduced to the market at generation 2, indicator IND12 is introduced to the market at generation 4, and so on. Thus, the market is gradually injected with new techniques such as stochastic oscillators and new knowledge such as Hang Seng Index, Nikkei 225 Index, and the Dow Jones Industry Average. This conforms to the real-life market, as technical indicators are not all invented at the same time. The environmental space of artificial traders expands as the number of environmental variables increases. Generations here refer to the generations of social learning, i.e., each generation consists of 125 days trading.

In the ISP, social learning acts as a mechanism through which the absorption and dissemination of new information into the trading society occur. During the social learning stage, besides being able to publish their successful trading strategies, traders with poor performances have the opportunity to discard the old set of indicators they used, and select a new set of indicators from the environment, which possibly includes newly introduced indicators. In addition, for traders who are satisfied with their performance, besides allowing them to keep their old trading models, we also give these traders the opportunity to try out other new indicators from the market, as shown in Section IV-D. When good trading strategies are developed through the use of new indicators, these good trading strategies will also be published to the market and learned by other traders. Thus, the new indicators are disseminated among the trading society. In order to put pressure on the use of new information, we assign the newly introduced indicators with the highest score among the current environmental variables, which mimics the phenomenon that new techniques appearing in a market usually attract public attention.

We define $U(e)$ as the utility of an environmental variable. In the imperfect evolutionary market, an indicator's $U(e)$ simply equals the number of times that it has been selected for use by traders. By examining the changes in an indicator's utility $U(e)$, we will be able to see if a piece of new information in the environment has been perceived and absorbed by the imperfect evolutionary market. Ten simulations are run on each of the five selected stocks with the dynamically changing environmental variables. Figs. 3–5 depict 20 indicators' $U(e)$ from the best run on each stock. The $U(e)$ of indicators IND1–IND10 are recorded starting from generation 1. For indicators IND11 and IND20, $U(e)$ is recorded starting from the generation before they are introduced into the market. For example, IND20 is introduced in generation 20, and its $U(e)$ is depicted starting from generation 19, when it has the value of 0. Therefore, the utility of indicators IND11–IND20 start with the value of 0, which, in

fact, indicates that they represent new variables emerging in the environment.

Fig. 3(a) shows the utility of the *static variables* from the simulation run on the CK stock with five randomly selected indicators' $U(e)$ from IND1 to IND10. Fig. 3(b) shows the traders' response to the *new challenges* in the environment from the simulation on the CK stock with five randomly selected indicators' $U(e)$ from IND11 to IND20. In Fig. 3(a), the indicators are static environmental variables given to the market at the beginning of the evolution. Each line in Fig. 3(a) starts at generation 1 with a nonzero value, which indicates the frequency at which the indicator is selected in the very first round of social learning. Since a trader is allowed to select an indicator more than once, we see some indicators' $U(e)$ start with a value that is greater than 50. As the evolutionary process continues, indicators' utility increases.

In Fig. 3(b), each indicator is introduced to the market at a different time. Before it is introduced, the indicator does not exist in the market. Each $U(e)$ line in Fig. 3(b) starts at a different point on the x axis, which corresponds to the generation before it is introduced to the market, and has a utility of 0. When it comes to the generation in which the indicator is injected, the social learning mechanism will put pressure on the traders to try this new indicator for building new trading models. This results in the increase in $U(e)$ values after the starting point of each line. As shown in Fig. 3(b), new indicators are gradually selected and used by traders. If an indicator is useful for trading, it will be accepted by more traders through the social learning process, and subsequently has a high utility, such as IND16 in Fig. 3(b). If the new information is not as useful, it gradually loses favor with traders, such as IND12 in Fig. 3(b). Therefore, it is clearly demonstrated in Fig. 3(b), the absorption and dissemination of new environmental variables by the trading society during the simulation on CK stock. Similar phenomena can also be observed on the CATHAY stock [Fig. 3(c) and (d)], the WHARF stock [Fig. 4(a) and (b)], the TOYOTA stock [Fig. 4(c) and (d)], and the SONY stock [Fig. 5(a) and (b)].

Some indicators' utility stayed at zero for a few generations after they have been injected to the market, such as IND12 in Fig. 3(b) and IND18 in Fig. 4(b). This is because they were not selected for use by traders after they were introduced to the market. There are two factors that could be responsible for a new environmental variable remaining undetected by the imperfect individuals: 1) the randomness in the selection procedure of social learning and 2) at the time when a new variable emerges, the traders are generally doing well and there is little pressure on social learning. Interestingly, by looking at Figs. 3(b), (d), and 4(b), we find that IND16, IND18, and IND20 are generally well accepted by the artificial stock traders. Referring to Table II, these three indicators are rates of change of the three major market indexes. It is generally accepted by investors in the stock market that market indexes have a significant impact on the stock prices. Often stock prices simply rise and fall along with the rise and fall in major market indexes regardless of the company's actual performance. We want to stress that the *new challenges* (IND11–IND20) are gradually introduced to the market as black boxes, i.e., the traders do not know what these indicators represent. The traders are only aware that some new infor-

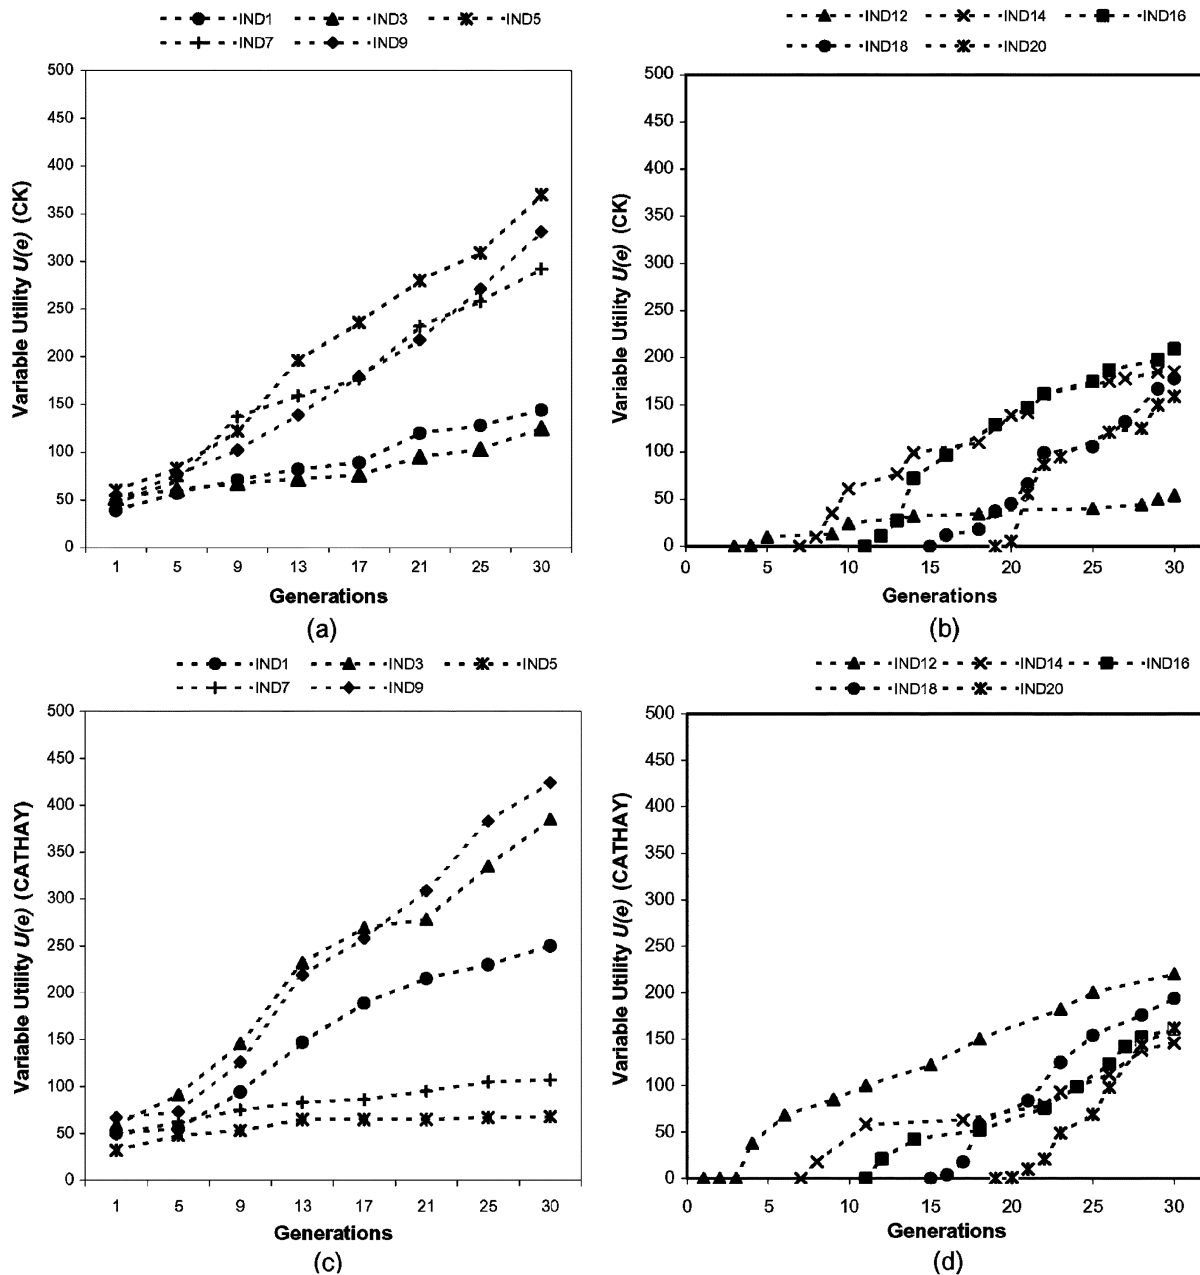


Fig. 3. The usage of environmental variables and absorption of new information [CK(a/b) and CATHAY(c/d)].

mation has appeared. Through evolution, artificial agents find the best way to make use of the new information and transfer it into skills and knowledge. As a summary, the experiments clearly demonstrate the absorption and dissemination of new information and knowledge within the imperfect evolutionary market using the ISP learning paradigm, and the adaptation of artificial stock traders to the new challenges from their imperfect environment.

2) *Static Environment Versus Imperfect Environment:* In Table III, we compare the performance of imperfect traders who evolve under the static 20 environmental variables, against traders who evolve with dynamically changing environmental variables. We use four criteria for the comparison of artificial stock traders' performance against two benchmark investment strategies: *Buy and Hold Strategy* (BnH) and *Bank Savings*.

The BnH strategy requires stock investors to buy and hold a stock over a long period in the hope that time will eventually ride out the volatility in the stock price. The four criteria used are: number of traders who outperformed the BnH strategy; number of traders who outperformed the risk-free bank savings; the best trader's cumulative total return; the average cumulative total return of 50 traders from the best run.

Clearly, from Table III, traders using dynamic environmental variables performed poorly compared with traders using static environmental variables. This is understandable as the traders with the static environmental variables have more opportunities to make use of indicators IND11–IND20 from the beginning of the evolution, whereas, the traders with dynamic environmental variables are only able to make use of indicators IND11–IND20 when they are introduced to the market. Obvi-

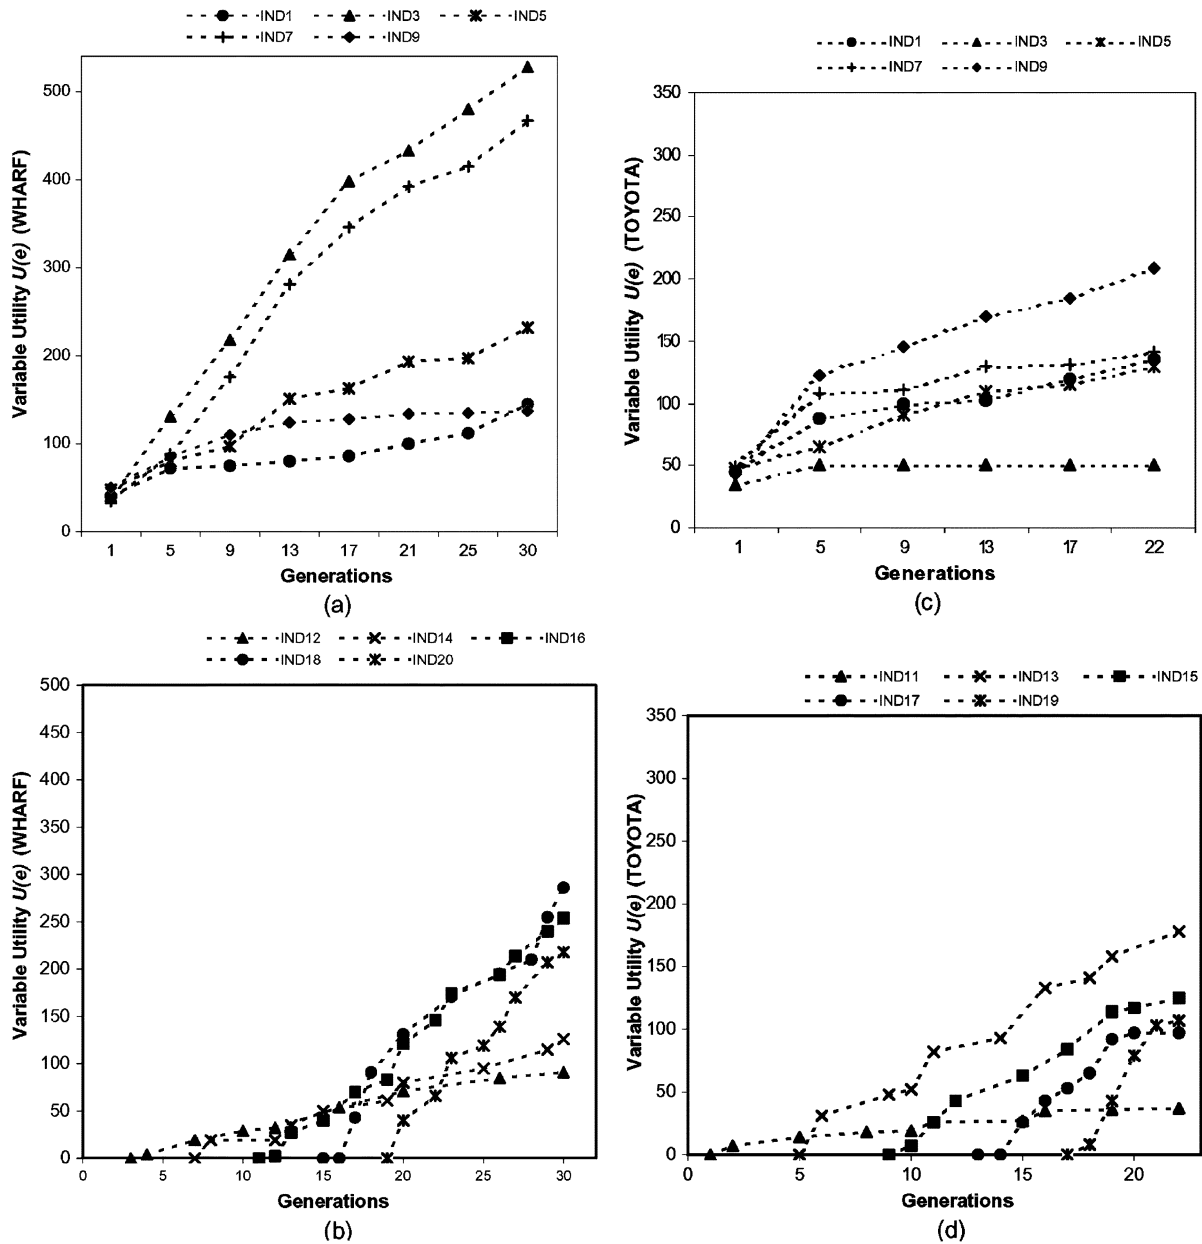


Fig. 4. The usage of environmental variables and absorption of new information [WHARF(a/b) and TOYOTA(c/d)].

ously, traders in an imperfect environment are disadvantaged in terms of knowledge and time. The comparison shows that for many scientific problems, people know what the exact observational space of a problem is. However, for real-life problems, observational spaces, and therefore their related problem spaces, are generally not of a static shape and a static domain. Dealing with an imperfect environment with new challenges is a problem that should not be ignored.

B. Individual Learning and Social Learning With Different Frequencies

The social learning mechanism in the IES plays an important role in ensuring artificial agents' adaptation to a changing environment. In the next set of experiments, we study the sensitivity of the parameters in the ISP learning paradigm by letting the individual and social learning occur at different frequencies.

As shown in Table IV, we let individual learning occur every 5 trading days or every 25 trading days. We let social learning occur every 125 trading days or every 250 trading days. A five-day trading period is approximately one week trading (sometimes, it is shorter, e.g., when there is a public holiday). A 125-day trading period is approximately a six-month trading period. The simulation of trading is run on each of the five stocks from Table I, with the four different frequencies shown in Table IV. Ten runs are carried out on each frequency setting. The results from the best runs of each stock are depicted in Figs. 6 and 7.

We use the same four criteria for the comparison of the traders' performance with difference frequencies in individual and social learning: the number of traders who outperformed the benchmark BnH strategy shown in Fig. 6(a); the number of traders who outperformed the risk-free bank savings shown in

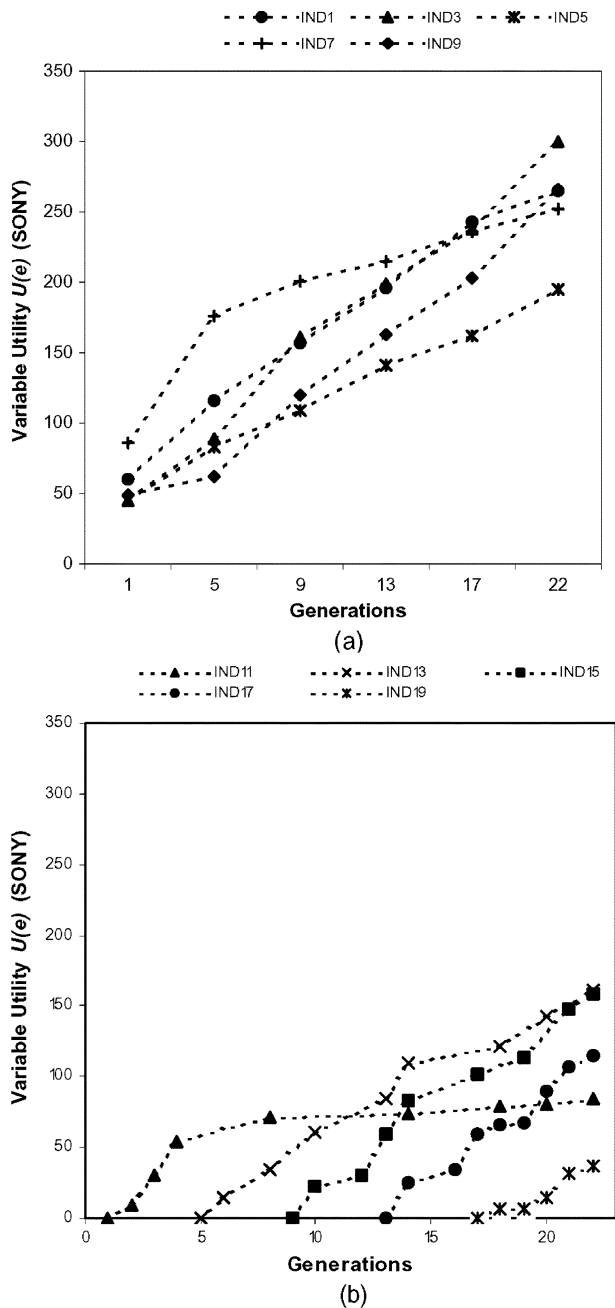


Fig. 5. The usage of environmental variables and absorption of new information (SONY).

Fig. 6(b); the top trader's cumulative total return, as shown in Fig. 7(a); and the average cumulative total return of 50 traders from the best run, as shown in Fig. 7(b). As shown in Table I, the classical BnH strategy does not suit every stock. CATHAY and SONY are two stocks which have suffered serious losses in the past few years due to the economic climate or a failed corporate development strategy. The apparent choice for investors in this case is to leave their money in the bank. Nevertheless, the artificial active stock traders in the simulated stock market, who use ANNs as trading models to detect buy and sell signals from the market, have developed successful trading strategies that beat investments in bank savings in the case of CATHAY and SONY, as shown in Fig. 6(b), and beat the BnH strategy in

the case of CK, WHARF, and TOYOTA, as shown in Fig. 6(a). With five different types of stocks, the results demonstrate the robustness of the ISP learning paradigm in the IES.

With regard to the effect of changing frequencies in individual and social learning, Figs. 6 and 7 clearly show that the learning ability of artificial traders declines when the frequency of social learning is reduced. Simulations with more frequent social learning (5-125 and 25-125) generally improve the agents' performance significantly in all four criteria. This phenomenon can be explained in two ways.

- 1) The reduction in the frequency of social learning results in less opportunities for agents to learn from other successful agents, and hence it takes agents much longer to find a good trading strategy in which time they may suffer from serious losses from which they may never recover.
- 2) The reduction in social learning also means less successful trading strategies are developed in the society as agents have less opportunities to publish their good strategies to the society and have less opportunities to discard inferior information and select new sets of indicators from the market for building new models.

However, the effects of individual learning occurring at different frequencies are not so clear. For example, considering CK in Fig. 6(a), more frequent individual learning improves the performance of the best trader when social learning also occurs frequently, whereas more frequent individual learning did not actually improve the performance when the social learning process occurs less frequently. However, for WHARF in Fig. 6(a), when individual learning occurs at higher frequencies, better results are produced under both social learning settings. We also see mixed results on other stocks with other criteria regarding the effects of different frequencies in individual learning. Apparently, it is not reasonable to simply assert that individual learning should occur more frequently or less frequently. The mixed results on individual learning are due to several reasons. When individual learning occurs more frequently, it accelerates the discovery of good trading models, and also enables traders to adapt to changes in the market more quickly. However, as the performance of artificial traders is evaluated using a trader's profit from the past week, or the past six months, there are problems that the fitness of a trading model may not be completely reflected within a short time. As an example, during the correction of a stock market, the stock prices will drop from a previous high, but are usually followed by another surge in the stock price. A good trading strategy in a period like this is to buy stocks when the prices have dropped to a certain resistance level, and wait for a bull market. The correction of a market may take a few weeks, or a few months. An individual learning over a short timescale, of (say) five trading days, will not fully reflect the trading strategy's profitability. This causes the possible loss of good trading strategies, and the evolution may move in a wrong direction.

In fact, what the results show us is that a dynamic individual and social learning scheme might be more appropriate in the ISP paradigm. Instead of designing the agents to learn at a certain time, the learning entities should have the right to decide when they need individual learning or when they need to turn to the society for help. Thus, every individual within the dynamic

TABLE III
COMPARISON OF RESULTS ON STATIC AND IMPERFECT ENVIRONMENT (ARTIFICIAL TRADERS' TRADING PERFORMANCE)

Stock	Static Environments				Dynamically Changing Environmental Variables			
	Outperform BnH	Outperform BS	Best Return (%)	Average Return (%)	Outperform BnH	Outperform BS	Best Return (%)	Average Return (%)
C.K.	44	50	2953.47	947.21	28	49	1862.19	672.33
CATHAY	50	49	1155.03	413.88	34	19	440.9	98.3
WHARF	50	50	1639.89	732.79	34	40	1312.99	280.17
TOYOTA	42	49	417.14	228.23	16	33	359.94	102.93
SONY	49	38	718.07	175.75	49	27	591.99	109.79

Results are taken from the best run out of 10 runs on each stock under different conditions, i.e., static environmental parameters and dynamically changing environmental parameters. .

Return refers to accumulated total return over the whole trading period; Outperform BnH = number of traders outperform the Buy and Hold strategy (BnH) out of the 50 traders, Outperform BS = number of traders outperform bank savings (BS) out of the 50 traders, Best Return = maximum return from the 50 traders, Average Return = average return from the 50 traders.

TABLE IV
INDIVIDUAL AND SOCIAL LEARNING OCCURRING AT DIFFERENT FREQUENCIES

Timescale	Description
5-125	Individual learning occurs every 5 trading days Social learning occurs every 125 trading days
25-125	Individual learning occurs every 25 trading days Social learning occurs every 125 trading days
5-250	Individual learning occurs every 5 trading days Social learning occurs every 250 trading days
25-250	Individual learning occurs every 25 trading days Social learning occurs every 250 trading days

imperfect environment will adjust their learning frequencies depending on their own evolutionary process. In addition, forward learning methods, i.e., evaluate a trading model's fitness by calculating its future profitability, should be used in coordination with backwards learning in order to offset the problems caused by fitness evaluation timescales.

Note that Fig. 7(a) shows the *cumulative return* of the best performer from ten simulations on each stock over 15 years trading. The figures may look abnormally high, for example, the CK stock achieved a nearly 3000% cumulative return. However, if we calculate an approximate annual return on the CK stock over the 15-year trading, we have the value as $3000\%/15 = 200\%$ per year. Assuming an initial investment of £1000 at the beginning of a year, a trader then needs to make £3000 at the end of the year to achieve an annual return of 200%. For a good market timer in the real-world stock markets, it is not difficult to achieve such a performance. For our artificial traders, we give each trader £10 000 cash to start with, where the effect of compound interest is more significant. For more discussions on active trading and market timing, please refer to Bodie *et al.* [41, pp. 919–923].

VI. CONCLUSIONS AND FUTURE WORK

We have described an imperfect environment as an environment that constantly presents new challenges to its participants. Inside an imperfect environment there exist imperfect individuals who do not completely understand their environment due




Fig. 6. Performance of artificial stock traders with different frequencies of individual and social learning.

to the complexity of the world surrounding them and every imperfect individual has a (possibly) different view of their environment. An IES consists of imperfect individuals living in an imperfect environment where intelligent individuals optimize their own utility, with the resources available, while adapting




Fig. 7. Performance of artificial stock traders with different frequencies of individual and social learning.

themselves to the new challenges from the evolving and imperfect environment. We have used two examples to illustrate our concept of IESs. One example demonstrated a simple imperfect evolutionary market and we discussed how imperfect evolution enables stock traders to learn new investment techniques. We also described a simulated game-playing “brain,” and how it would be possible for a brain to learn from the games it played and also to learn different kinds of games. We propose IESs as a change from a perfect perspective to an imperfect perspective in designing intelligent systems. The comparison between the cultural learning algorithm and the ISP further stresses that an IES concerns the incompleteness and continuity of intelligence, i.e., how to sustain the continuous innovation of an intelligent system.

In this paper, as a first attempt in developing IESs, we implemented an imperfect evolutionary market based on an integrated individual and social learning algorithm. Through experiments, we clearly demonstrate the absorption and dissemination of new information and knowledge within the imperfect evolutionary market under the ISP learning paradigm, and the adaptation of

imperfect artificial traders to the new challenges from their environment. The ISP learning paradigm demonstrates its robustness under various market scenarios. However, there are still a number of aspects where the ISP learning paradigm can be improved, e.g., the knowledge evolution in the social pool and the credit assignment problem in both individual and social learning processes.

Blondie24 is successful because it solved the problem of machine learning without preinjected human knowledge. However, does it mean we have created intelligence when we develop machines that can play games at expert levels or even beat human world champions? Not quite. Many people play checkers, but not all of them to an expert level. Nevertheless, they are intelligent. Under the IES methodology, we believe it is possible to develop a computer game-playing program that learns from the games it plays, and also learns to play different kinds of games. However, there are still many practical questions that need to be answered. For example, in the imperfect evolutionary market, the intelligent system only has one objective, i.e., to make a profit. If we are going to learn different kinds of games, we will have multiple objectives in one intelligent system, e.g., playing checkers or playing poker. How do we create a new individual, i.e., a new functional section, in the simulated game-playing mind for a new game? How can the ISP learning paradigm be improved to accomplish this? How do we differentiate information? Should the functional sections in the simulated “brain” interact with each other? And how? The development of the concept on IESs, presented in this paper, provides a new perspective in designing adaptive intelligent systems. To implement more complex IESs still requires many practical problems to be solved.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their valuable comments, which has led to a much improved paper.


REFERENCES

- [1] K. Chellapilla and D. B. Fogel, “Evolving an expert checkers playing program without using human expertise,” *IEEE Trans. Evol. Comput.*, vol. 5, no. 4, pp. 422–428, Aug. 2001.
- [2] D. B. Fogel, *Blondie24, Playing at the Edge of AI*. San Mateo, CA: Morgan Kaufmann, 2002.
- [3] J. Schaeffer, R. Lake, and P. Lu, “CHINOOK the world man-machine checkers champion,” *AI Magazine*, vol. 17, no. 1, pp. 21–30, 1996.
- [4] J. Schaeffer, *One Jump Ahead: Challenging Human Supremacy in Checkers*. New York: Springer-Verlag, 1997.
- [5] D. Goodman and R. Keene, *Man versus Machine: Kasparov versus Deep Blue*. Cambridge, MA: H3 Publications, 1997.
- [6] M. Campbell, A. J. Hoane, and F.-H. Hsu, “Deep blue,” *Artif. Intell.*, vol. 134, no. 1–2, pp. 57–83, 2002.
- [7] A. L. Samuel, “Some studies in machine learning using the game of checker II,” in *Computer Games I*, D. L. Levy, Ed. New York: Springer-Verlag, 1998, pp. 366–400, 1967.
- [8] E. Harley, “Book review: *Blondie24, playing at the edge of AI*,” *The IEEE Comput. Intell. Bulletin*, vol. 1, no. 1, pp. 25–27, Dec. 2002.
- [9] K. O. Stanley, B. D. Bryant, and R. Miikkulainen, “Real-time neuroevolution in the NERO video game,” *IEEE Trans. Evol. Comput.*, vol. 9, no. 6, pp. 653–668, Dec. 2005.
- [10] J. E. Stiglitz, “Information and the change in the paradigm in economics,” *Les Prix Nobel*, pp. 472–540, 2001, The Noble Foundation.

- [11] R. Arnott, B. Greenwald, R. Kanbur, and B. Nalebuff, *Economics for An Imperfect World—Essays in Honor of Joseph E. Stiglitz*. Cambridge, MA: MIT Press, 2003.
- [12] A. Newell and H. A. Simon, “GPS, a program that simulates human thought,” in *Computers and Thought*, E. A. Feigenbaum and J. Feldman, Eds. New York: McGraw-Hill, 1963, pp. 279–293.
- [13] M. L. Minsky, *The Society of Mind*. New York: Simon & Schuster, 1986.
- [14] D. S. Weld and J. de Kleer, *Readings in Qualitative Reasoning about Physical Systems*. San Mateo, CA: Morgan Kaufmann, 1988.
- [15] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann, 1988.
- [16] L. A. Zadeh, “Fuzzy sets,” *Information and Control*, vol. 8, pp. 338–353, 1965.
- [17] ———, “Fuzzy sets as a basis for a theory of possibility,” *Fuzzy Sets and Systems*, vol. 1, pp. 3–28, 1978.
- [18] X. Yao, *Evolutionary Computation—Theory and Applications*. Singapore: World Scientific, 1999.
- [19] D. B. Fogel, *Evolutionary Computation: Toward A New Philosophy of Machine Intelligence*, 2nd ed. Piscataway, NJ: IEEE Press, 2000.
- [20] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computation (Natural Computing Series)*. New York: Springer, 2003.
- [21] D. B. Fogel, A. Blair, and R. Miikkulainen, Eds., “Special issue on Evolutionary and Games,” *IEEE Trans. Evol. Comput.*, vol. 9, no. 6, Dec. 2005.
- [22] R. G. Reynolds, “An Adaptive computer model of the evolution of agriculture for hunter-gatherers in the Valley of Oaxaca,” Ph.D. dissertation, Univ. Michigan, Ann Arbor, MI, 1979.
- [23] ———, “An introduction to cultural algorithms,” in *Proc. 3rd Annu. Conf. Evol. Program.*, 1994, pp. 131–139.
- [24] M. Sternberg and R. G. Reynolds, “Using cultural algorithms to support re-engineering of rule-based expert systems in dynamic environments: A case study in fraud detection,” *IEEE Trans. Evol. Comput.*, vol. 1, no. 4, pp. 225–243, 1997.
- [25] R. G. Reynolds, Z. Kobti, and T. A. Kohler, “The effects of generalized reciprocal exchange on the resilience of social networks: An example from the prehistoric Mesa Verde region,” *J. Comput. Math. Organ. Theory*, vol. 9, no. 3, pp. 229–254, 2003.
- [26] R. G. Reynolds, Z. Kobti, T. A. Kohler, and L. Yap, “Unravelling ancient mysteries: Reimagining the past using evolutionary computation in a complex gaming environment,” *IEEE Trans. Evol. Comput.*, vol. 9, no. 6, pp. 707–720, Dec. 2005.
- [27] R. G. Reynolds and S. Saleem, “Function optimization with cultural algorithms in dynamic environments,” in *Proc. IEEE Particle Swarm Optimization Workshop*, 2001, pp. 63–79.
- [28] R. G. Reynolds and B. Peng, “Cultural algorithms: Knowledge learning in dynamic environments,” in *Proc. IEEE Int. Congr. Evol. Comput.*, 2004, pp. 1751–1758.
- [29] R. G. Reynolds and S. Saleem, “The impact of environmental dynamics on cultural emergence,” in *Perspectives on Adaptation in Natural and Artificial Systems—Essays in Honor of John Holland*, L. Booker, S. Forrest, M. Mitchell, and R. Riolo, Eds. London, U.K.: Oxford Univ. Press, 2004, pp. 253–280.
- [30] R. G. Reynolds and D. Ostrowski, “Using cultural algorithms to evolve strategies for recessionary markets,” in *Proc. IEEE Int. Congr. Evol. Comput.*, 2004, pp. 1780–1785.
- [31] R. G. Palmer, W. B. Arthur, J. H. Holland, B. LeBaron, and P. Tayler, “Artificial economic life: A simple model of a stock market,” *Physical D*, vol. 75, pp. 264–274, 1994.
- [32] W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler, “Asset pricing under endogenous expectations in an artificial stock market,” in *The Economy As an Evolving Complex System*. Reading, MA: Addison-Wesley, 1997, vol. XXVII, pp. 15–44.
- [33] S. H. Chen and C. H. Yeh, “Toward an integration of social learning and individual learning in agent-based computational stock markets: The approach based on population genetic programming,” *J. Manage. Econ.*, vol. 25, pp. 363–393, 2001.
- [34] S. Schulenburg and P. Ross, “Strength and money: An LCS approach to increasing returns,” in *Advances in Learning Classifier Systems*, ser. Lecture Notes in Artificial Intelligence, P. L. Lanzi, W. Stolzmann, and S. W. Wilson, Eds. Berlin, Germany: Springer-Verlag, 2001, vol. 1996, pp. 114–137.
- [35] S. H. Chen, *Evolutionary Computation in Economics and Finance*. Berlin, Germany: Physica-Verlag, 2002.
- [36] G. Kendall and S. Yan, “Co-evolution of successful trading strategies in a simulated stock market,” in *Proc. Int. Conf. Mach. Learn. Appl.*, 2003, pp. 200–206.
- [37] ———, “A multi-agent based simulated stock market—testing on different types of stocks,” in *Proc. Congr. Evol. Comput.*, 2003, pp. 2298–2305.
- [38] ———, “Learning with imperfections—A neural-genetic trading system with different levels of social learning,” in *Proc. IEEE Int. Conf. Cybern. Intell. Syst.*, 2004, pp. 47–52.
- [39] D. E. Moriarty and R. Miikkulainen, “Forming neural networks through efficient and adaptive evolution,” *IEEE Trans. Evol. Comput.*, vol. 4, no. 5, pp. 373–399, 1997.
- [40] X. Yao, “Evolving artificial neural networks,” in *Proc. IEEE*, Sep. 1999, vol. 87, no. 9, pp. 1423–1447.
- [41] Z. Bodie, A. Kane, and A. J. Marcus, *Investments*. New York: McGraw-Hill, 2002.



Graham Kendall (M'03) is a member of the Automated Scheduling, Optimization and Planning Research Group, University of Nottingham, Nottingham, U.K. He is a member of the U.K. Engineering and Physical Sciences Research Council (EPSRC) Peer Review College and is an associate editor of three international journals. He also chairs the steering committee of the Multidisciplinary International Conference on Scheduling: Theory and Applications (MISTA) and has chaired several other international conferences. Dr Kendall has been a member of the program (or refereeing) committees of over 70 international conferences over the last few years. During his career, he has edited/authored six books and has published over 70 refereed papers. He has been awarded externally funded grants worth over 5M pounds from a variety of sources including EPSRC and commercial organizations. His other research interests include: adaptive learning (with an emphasis on game intelligence), heuristic development, optimization, scheduling, and artificial intelligence.



Yan Su received the Ph.D. degree from the University of Nottingham, Nottingham, U.K., in 2005, after studying in the Automated Scheduling, Optimization and Planning Research Group.

She is the author of five internationally peer reviewed papers and her research interests include individual and social learning and artificial intelligence.