Applying the El-Farol Protocol to Strategically Sample Spin Models

Jared Rosenblitt

Department of Chemistry

McGill University

Montréal, Québec, Canada

April 15th, 2021

A thesis presented for the degree of Master of Chemistry (Theoretical)

Abstract

Exploring the features of complex energy landscapes is an important challenge in the computational simulation of molecular systems. A simplified and more insightful view of the system is often obtained by locating and mapping out the free energy minima and barriers. In this thesis we propose to use an approach that is predominantly used in the field of economics, an approach called the 'El-Farol Bar Problem', in order to explore model systems of broad physical-chemical interest: spin models. The 'El-Farol Bar Problem' falls under the 'inductive reasoning' or 'bounded rationality' approaches to the modeling of market dynamics. At its core there is a machine learning algorithm in which the parameters of the 'strategies' are learned in a discontinuous and stochastic fashion via minimizing the objective that is set by the collective behavior of a large number of independent agents acting together. Here we look at the application of the El-Farol protocol to computationally simulate Ising-type spin systems. Our adapted El-Farol protocol shows an interesting way to sample spin models in a strategic manner. We present the El-Farol protocol and apply it in a proof-of-concept study to the problem of global minimization in two interacting Ising-type

spin models. We demonstrate that in our calculations the El-Farol protocol is an effective approach for exploring the energy landscape through control of the magnetisation order parameter. Moreover, we compare it with a naive implementation of Stochastic Tunnelling (STUN) Monte Carlo to show that it can produce relatively comparable results.

Abrégé

L'exploration des caractéristiques des paysages énergétiques complexes est un défi important dans la simulation informatique des systèmes moléculaires. Une vue simplifiée et plus perspicace du système est souvent obtenue en localisant et en cartographiant les minimums d'énergie libre et ses barrières. Dans cette thèse, on propose d'utiliser une approche principalement utilisée dans le domaine de l'économie, appelée «Problème du Bar d'El-Farol», afin d'explorer des systèmes de modèles présentant un large intérêt physico-chimique: les modèles de spin. Le «Problème du bar d'El-Farol» relève des approches de «raisonnement inductif» ou de «rationalité limitée» de la modélisation de la dynamique du marché. À la base, il y a un algorithme d'apprentissage automatique dans lequel les paramètres sont appris de manière discontinue et stochastique en minimisant l'objectif fixé par le comportement collectif d'un grand nombre d'agents indépendants agissant ensemble. On examine dans ce texte l'application du protocole d'El-Farol pour simuler par calcul des systèmes de spin de type Ising. Le protocole d'El-Farol adapté montre une manière intéressante d'échantillonner des modèles de spin de manière

stratégique. On présente le protocole d'El-Farol et l'applique dans une étude de preuve de concept au problème de la minimisation globale dans deux modèles de spin de type Ising en interaction. On démontre que dans ces calculs, le protocole d'El-Farol est une approche efficace pour explorer le paysage énergétique en contrôlant le paramètre d'ordre de magnétisation. De plus, on le compare à une implémentation naïve de Stochastic Tunneling (STUN) Monte Carlo pour montrer qu'il peut produire des résultats relativement comparables.

Acknowledgements

I would like to thank Professor Lena Simine and the members of the Simine lab for their insight, kindness and patience. The world is lucky to have such gifted individuals as them, and without whom this would not have been possible. I would especially like to thank Professor Simine for her expertise and enthusiasm throughout this project and her constant encouragement through thick and thin. Thank you.

Contents

List of Tables

Chapter 1

Introduction

The sampling of complicated energy landscapes is a central problem in the simulation of molecular systems. There are many algorithms that have been developed for this purpose including Monte Carlo methods, multi-dimensional sampling, methods involving simulated annealing (just to name a few) and there are so many others, yet new ideas are often still needed to tackle this large and varied problem [1–8]. This thesis takes an interdisciplinary approach and proposes exploiting the potential of a protocol that has been borrowed from the field of economics, called the 'El-Farol Bar Problem' and to apply it as a means of sampling energy landscapes in problems of physical-chemical interest [9]. The 'El-Farol Bar Problem' or 'El-Farol' (for short) is an unusual machine learning algorithm, one in which the parameters are learned in a discontinuous and stochastic fashion with the constraints imposed through the resulting emergent behavior of a large number of sub-systems or agents acting

together yet independently from each other, each following an independent set of learned 'strategies'. This approach is broadly referred to as 'inductive reasoning' and 'bounded rationality' by the problem's founder Brian Arthur [9, 10]. 'Strategic sampling' in the title of the thesis refers to the idea of using the 'strategies' devised by the agents in the El-Farol scheme to propose states of a system. In its original context, it refers to the number of patrons that an agent predicts will attend a bar in a given week and this prediction will make them ultimately decide for themselves whether to go or not. It is this approach that is taken and re-worked for use in predicting spin configurations of an Ising system that will ultimately provide the narrative for exploring energy states and sampling from the energy landscape in order to find a global minimum. The details of the El-Farol protocol will be presented in Chapter 3.

The systems of interest considered in this thesis are two variants of the Ising Model, which is a particular example of a spin-model system. Spin-models are broadly used across the computational Chemical and Physical sciences: from magnetic systems and quantum information, to modelling of protein folding, to neural networks [11–16]. Traditionally, the Ising model has been used to investigate phase transitions and ferromagnetism [17]. It was ground-breaking for statistical mechanics at the time and has been applied in a variety of ways, from use in simple gases to systems as detailed as magnetic interactions [12, 15]. The simplicity of the model, however, has shown that in can be applied across and studied over a wide range of fields, for example in quantum information it has been useful for

exploring particular Hamiltonians, their corresponding eigenstates and entanglement as well as particular aspects such as quantum Fisher information [18,19]; in the modelling of protein folding, for repeat protein folding for conformational transitions in linear polymers and more general single site contact models among others [20, 21]; even in computer science, where the model can represent nodes on a graph sharing information, or the "spin sites" can take on more than two values, with such an abstraction being referred to as the "Potts" model, and can even be used as training data for reinforcement learning [11,14,16,22,23]. For these reasons a generic Ising-model-like energy landscape is the perfect toy-model to be explored for algorithmic testing.

This thesis is structured in the following way: in Chapter 2 we present key background information about the El-Farol bar problem, and the details of the algorithms used in this thesis; the upgraded El-Farol bar protocol and the Stochastic Tunneling Monte Carlo algorithm (STUN). Chapter 3 presents a brief review of the Ising model and the adaptation of the El-Farol protocol to the simulation of Ising-type spin models. Chapter 4 presents the proof of concept results for using the El-Farol protocol to simulate the interacting Ising model while exerting control over the magnetisation order parameter. Chapter 5 presents conclusions and outlook.

Chapter 2

Background

In this chapter we present the historical El-Farol bar problem, and its modern implementation. Furthermore we present the Stochastic Tunneling Monte Carlo algorithm which will be used in Chapter 4 to gauge the performance of the El-Farol protocol at the task of locating the global minima in Ising-type models.

2.1 The El-Farol Bar Problem

The El-Farol bar problem is attributed to noted economist Brian Arthur. Arthur developed this concept in his 1994 paper "Inductive Reasoning and Bounded Rationality", the key concept of which was players, dubbed "agents", who do not make completely rational or optimal choices when participating in a game [9]. Arthur himself was not satisfied with the then-current treatment of agents in economic contexts as players that always play completely

rationally with full knowledge of all possible outcomes, something that rarely reflects real-world market dynamics.

His idea was to introduce models of bounded rationality. Most theoretical models treat the participating parties in the model as rational beings who always make the optimal choice. However, according to Arthur this does not always reflect reality. Using a chess example, he explains that while theoretically one can determine all possible moves at a given starting point, it is impossible for the human brain to do so and instead it simplifies by thinking in terms of what to expect as the "most likely strategy" to be played by the opponent. In this way he suggests that humans think "inductively" [9].

In the El-Farol bar game there are a set of "agents" which can be thought of as people, or any being that cannot fathom a complete picture of their scenario (they need not even be 'conscience', as they can be 'zero-intelligence agents' simply acting as automation, like spins) [24]. Instead they have an "idea" about how they "think" or "expect" their surroundings to behave in response to their actions. Ideas can be defined as a set of hypotheses or predictors about how the agent expects their environment to operate. The agents then employ a "strategy" which is the "best action" to take among all possible hypotheses/predictors they hold about the environment. Through learning, they change their strategy and use a different hypothesis/predictor if the strategy fails a consecutive number of times. Basically, an agent chooses what, at the given moment, is statistically the most likely predictor that will result in the best outcome for them and sometimes even for everyone (depending on the circumstance).

The El-Farol bar problem is described in Arthur's work as follows: suppose $N = 100$ people who live in Santa Fe, New Mexico, independently, each week, decide whether or not to go to a bar. The reason is that this bar plays Irish jig music on Thursday evenings which is considered by the local population a very enjoyable experience. However, the bar only has a capacity for a certain amount of people. When this capacity is exceeded it becomes crowded and the experience becomes unenjoyable to attendees. Let us set this capacity threshold to be 60 people maximum in the bar. Each week, on the Thursday right before the music starts, the 100 Santa-Fe residents each have a choice of whether to attend the bar or not. Their dilemma is as follows: if they go and the bar is under capacity they will have a good time but if they go and it is over capacity they will not have a good time. At the same time if they choose not to go and the bar is over capacity then they will have "dodged a bullet" and missed out on a bad experience, while if they choose to stay home but the bar is under capacity then they have missed out on a good time. They all have access to the full historic record of the bar attendance (only up to a certain number of weeks previous) but they do not coordinate their actions with each-other. The question is, how does the attendance number change from week to week given this prerogative [9]?

Remarkably, Brian demonstrated that despite the agents' inability to communicate and with very limited choice in the strategies they may employ to make their decisions, the average attendance at the bar robustly fluctuates around the capacity threshold of 60 [9]. This shows that through independent and individual interpretation of historic information

regarding bar attendance the agents gain the ability to self-organize around the capacity threshold despite having no communication with each other. We will exploit this feature of the protocol in our simulations of spin models where we control the magnetisation. In a way we will discuss in Chapter 3 the magnetisation in spin-$\frac{1}{2}$ models is analogous to attendance in the original El-Farol bar problem.

Although Arthur proposes a number of different types of strategies in his paper, there is no clear definition and there is still some ambiguity regarding how a strategy should be defined and interpreted. There have also been different directions in which the El-Farol bar problem has been taken. The most studied of these areas is the **minority game**, a broader generalization of the El-Farol bar problem and a subject which many physicists and economists have devoted time to over the past two decades [24–38]. The El-Farol bar problem itself, however, still provides avenues of research still worth exploring.

2.2 Implementation of the El-Farol Bar problem

Generally, the El-Farol bar problem involves the following:

There are N agents/players in the game and will be indexed over a discrete time-step t_l where $l = 0, 1, ..., T$ and each t_l represents a particular week. Then one has:

i) Agents' actions

At each time step t_l each agent i takes a binary decision $b_i(t_l) \in \{0, 1\}$ where 1 represents going to the bar and 0 is not going. The overall amount in attendance $A(t)$ at the bar for a

particular time-step/week t_l is given by

$$A(t_l) = \sum_{i=1}^{N} b_i(t_l) \tag{2.1}$$

where the summation runs over all agents in the system.

ii) Public Information

Agents base decisions on attendance history over the M (sometimes denoted 'm') most recent steps, given by the string

$$(A(t_{l-M}), ..., A(t_{l-1})) \in \{0,1\}^M \tag{2.2}$$

iii) Agents' strategies

Each agent i has at most S strategies $\mathbf{R}^{ia}$, with $a = 1, ..., S$. A strategy $\mathbf{R}^{ia}$ defines a mapping from attendance information strings to recommended action:

$$\mathbf{R}^{ia} : \{0,1\}^M \to \{0,1\} \tag{2.3}$$

In the original work [9] a strategy worked as a look up table with 2^M entries each being 0 or 1 which remain fixed throughout the game. In later developments these entries are generated by comparing the hypothesis or a predictor $\mathbf{x}^{ia}$ for this week's attendance (this is an integer between 0 and N) and assigned a 1 (attend) if it is **less than or equal to the**

attendance threshold and 0 (stay home) if it is **larger than the attendance threshold**. Agent i using strategy a at time-step t_l acts deterministically according to

$$b_a(t_l) = \mathbf{R}^{ia}(A(t_{l-M}), ..., A(t_{l-1})) \tag{2.4}$$

Obtaining closed equations means specifying how an agent determines which of their personal strategies to use. A strategy is one that allows them to be in the desirable group most of the time. Therefore a strategy $\mathbf{R}^{ia}$ at step t_l is good if the square error of the hypothesis/predictor $\mathbf{x}^{ia}$ on which the decision is based is small relative to the emergent attendance in that time-step :

$$|A(t_l) - \mathbf{x}^{ia}(A(t_{l-M}), ..., A(t_{l-1}))|^2 < \epsilon \tag{2.5}$$

where ϵ is a small number. This leads to "Strategy valuations".

iv) Strategy valuations

Each agent i keeps track of a valuation p_{ia} for each of their strategies, measuring track records (irrespective of used or not):

$$p_{ia}(t_{l+1}) = p_{ia}(t_l) + |A(t_l) - \mathbf{x}^{ia}(A(t_{l-M}), ..., A(t_{l-1}))|^2 \tag{2.6}$$

This equation is used in the calculation of mean square errors in our implementations of this algorithm.

v) Dynamic Selection Strategy: At each time step, each agent i will select a best strategy $a_i(t_l)$ at that stage of the process, defined by

$$\arg \max_{a \in \{1,\ldots,S\}} \{p_{ia}(t_l)\} \tag{2.7}$$

The basic microscopic degrees of freedom are the strategy valuations $\{p_{ia}\}$. These valuations are seen to evolve according to the coupled non-linear equations given above. Finally, if two or more different strategies have the same success/valuation, i.e. $p_{ia}(t_l) = p_{ia'}(t_l)$ for two strategies a, a', then select the next strategy to use, $a_i(t_l)$ randomly with equal probability between a, a'.

While this is the near-original formulation of the El-Farol bar problem, some suggestions have been made to further expand the scope of the problem so that it may, for example, incorporate a stochastic component. Other suggestions on modifying the problem have resulted in entirely new areas of study in both economics and statistical physics [25–27, 39–46]. One such field is the "minority game" [28,29]. The minority game became a phenomenon of interest to both physicists and economists at the turn of the century due to its ability to exhibit 'phase-transition'-like behaviour.

It is here that a brief comment must be made regarding the use of the word 'phase transition' and the exact connotations and assumptions under which it is made throughout the rest of this text. Historically, phase transitions referred to the physical phenomena in

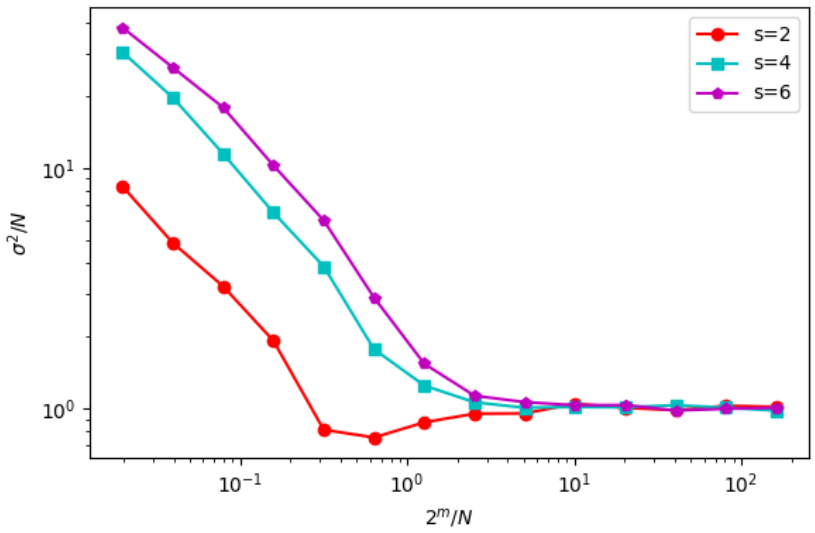

Figure 2.1: Normalized variance of the 'attendance' in the Minority game exhibits a 'phase transition' as a complicated function of history depths and the number of agents. The parameter 's' reflects the number of strategies in each agent's strategy pool. This figure was generated by the author following Ref. [29].

which solids would deform into liquids, gases and vice-versa. Attempts to study this natural phenomena yielded new insights into their classification schemes. Some of the earliest attempts are attributed to Paul Ehrenfest, who proposed that phase transitions be defined in terms of the discontinuities of the first, second and greater derivatives of the free energy, with each corresponding discontinuity of the derivative a 'phase transition'. A first-order phase transition, for example, was therefore the point at which the function describing the free energy was discontinuous with respect to one of its thermodynamic variables [47, 48]. A more detailed albeit brief account may be found in Charles Kittel's *Introduction to Solid-State Physics*. This scheme was useful, but was difficult to include other phenomena into the theory that were being investigated at the time, chiefly

magnetism and especially ferromagnetism. Since then, many schemes have been proposed, with one of the most successful being the 'critical exponents' method. This method, verified by experiments, shows that as a physical quantity approaches the phase transition at the critical temperature T_c, viewed in terms of the expression $|T - T_c|^{\lambda}$, the exponent λ somehow attains a universal value regardless of the material under investigation [15]. Therefore the various universal exponents for different classes of phase transitions are used to define the type of phase transition. A more detailed explanation is available in [15]. The most important aspect of this research in terms of the minority game was not the physical phenomena of phase transitions themselves, but the mathematics that had been developed and used to describe phase transitions. Physicists and economists merely took these same mathematical tools used in statistical mechanics and applied it to the study of minority games. The phrase 'phase transition' in this context does not refer to the physical phenomena of a transition, but instead refers to the mathematical/statistical mechanical tool, in this case of the critical exponent method for phase transitions, which was used to characterize similar behaviours observed when using the same tools for minority games. There is no physical way to interpret a 'phase transition' among agents in the minority game. Instead, 'phase transition' refers to the fact that there exists a critical value (more specifically a particular exponent embedded within this value when it is calculated), often referred to as α_c, for which agents switch from using intelligent strategies to purely random strategies, and this same value is observed in many different versions of the minority game

even when certain parameters of the game have been adjusted, in a similar comparison to the universality of the critical exponents, independent of the material [28, 29, 38].

Figure 2.1 is an example of the strange "phase transition" behaviour in a model closely related to El-Farol Bar problem but slightly different. We will not discuss the differences but use it to illustrate the idea of the phase transition which may be supported by the El-Farol dynamics as well. In the figure, the abscissa is given by $\frac{2^M}{N}$, where M is the "history-depth" and N is the number of agents in the simulation [29]. One then measures the variance in the attendance A of the bar, $\sigma = \sqrt{\langle A^2 \rangle - \langle A \rangle^2}$ as a function of the ratio $2^M/N$ for a fixed "history" and changing number of agents, each time with a different number of strategies S an agent uses per simulation [29]. The cut-off that agents tend towards at $\sigma = 1$ indicates purely random decisions. This figure shows that there is some critical point α_c in which turn-over occurs in which the agents transform to accurately predicting attendance at the bar and make wise decisions that improve everyone's experience [29]. However increasing N shows that as more people are added, decisions seem to resemble purely random behaviour, even if strategies are being employed. The key difference in minority game evaluation is the analysis of expressions for when $N \to \infty$ (or various other quantities, explained later).

This field has opened many new insights to the El-Farol bar problem and its generalization, the minority game and although it will not be discussed here, it is worth noting that this is a successful example of the applications of methods of statistical mechanics to economics. This thesis aims to take these newly developed insights gained

while studying a problem of economics and reapply it in a chemical/physical environment to see if any new insight may be gained on older and potentially newer problems in chemistry and physics.

2.3 Stochastic Implementation of the El-Farol Bar Problem

An interesting proposal for the implementation of the El-Farol bar problem comes from Fogel et al. [10]. The idea lies chiefly in the stochastic way of generating strategies. Rather than provide all agents with a fixed set of strategies that they will use throughout the entire game, Fogel et al. instead introduce the idea that agents can continuously change strategies. However, if a strategy is working well and continues to work well, then the agent will "keep" such a strategy for further use until it fails a sufficient number of times to justify discarding it. Equations (2.1) - (2.7) are therefore "combined" together in a way to form a single equation, called the agent's "predictor", which will determine the agent's prediction of the attendance for a given week. This predictor is given by

$$\hat{x}^i_{j,l^i_j}(t) = round\left(\left|a^i_{j,0} + \sum_{k=1}^{M \leq l^i_j \leq NS} a_{j,k}(t)A_k(t)\right|\right) \tag{2.8}$$

where i indicates the particular agent who is predicting an attendance using their given strategies at a particular time-step t, j is the particular strategy chosen to be used (note

however that all strategies are used to make a prediction but only the "best" strategy's prediction is utilized to determine attending the bar), $a^i_{j,0}$ is a bias term (in our simulations taken to be 0) while the general $a_{j,i}(t)$ is referred to as a *lag coefficient* and is a real random number selected with an unbiased probability from $[-1, 1]$ while $A_i(t)$ is the attendance in the bar on the k-th week back from the current week/time-step and can go from a minimum of M weeks into the past to a maximum of NS weeks. However, the authors introduce a term $l_{i,j}$ which is the number of lag coefficients allowed in the sum and it is chosen at random with uniform probability from $\{0, ..., k\}$ for each instance of the predictor. The actual sum term in Equation 2.8 is more recognizable by the name "autoregressive" function, in other words functions of the form

$$\sum_{k=1}^{m} f_k A(t-k)$$

where f_k are real numbers, and $A(t-k)$ are integers [10].

Thereby Fogel et al. developed a partially stochastic algorithm to simulate Brian Arthur's original El-Farol bar problem that in this instance allow the agents to discard strategies that do not work well and instead are randomly assigned new strategies, should said strategies continually perform poorly (as tallied by 2.6) [10]. In practice, before predicting the attendance in the first week, each agent evolves a set of models for 10 generations. Generations here refers to a small set of time-steps which the agents use to refine their existing models.

The overall flow of the algorithm including the refinement step is as follows:

1. Each agent i of the N agents creates an "off-spring" that has the same number of strategies NS as the "parent" agent/individual and in this case $NS = 10$.

    - The number of lag terms for the offspring i is chosen with equal probability from the number of lag terms in its parent j from $\{l_j^i - 1, l_j^i, l_j^i + 1\}$, where l_j^i is assigned uniformly at random from $\{1, ..., k\}$ at the start (in our case $k = 10$). If $l_j^i - 1$ is smaller than 1 or $l_j^i + 1$ is larger than 10, they are rounded up or down, respectively.
    - Once the number of lag terms needed is known, each of the coefficients $a_{j,k}(t), ..., a_{j,k}(t - l_j^i)$ are assigned by taking the coefficient of the parent and adding a zero-mean Gaussian random variable with standard deviation 0.1 to it (i.e. $N(0, 0.1)$). If the offspring has less terms than the parent, then coefficients are assigned zero in place of what comes from the parent, and if more, then they are given new coefficients taken from $N(0, 0.1)$.

2. The 20 ($2k = 20$) total models/number of strategies are evaluated based on the sum of each strategy's mean square error in predicting bar attendance from the last 12 weeks. The hyper-parameters here are chosen based on experimental evidence of good performance.

3. The 10 best models, that is the ones with the lowest sum of their mean square errors, are chosen as parents for the next generation. See Equation 2.6 for definition of the

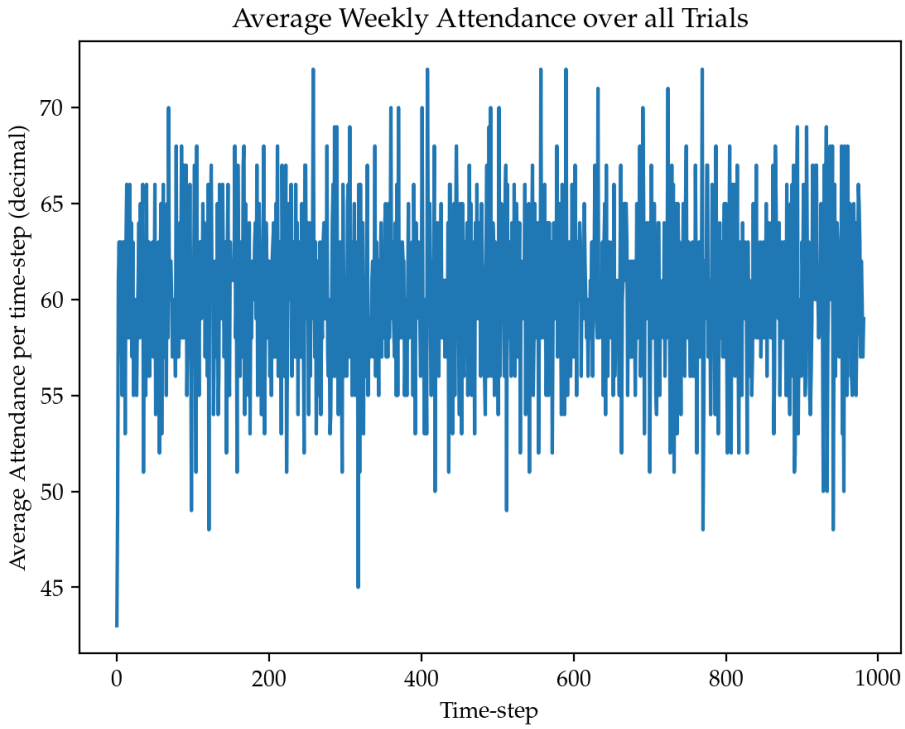

Figure 2.2: Mean Attendance across all 20 trials.

error function.

4. Steps 1-3 are repeated if ten generations had not yet been completed; otherwise each agent used its current best strategy (one with lowest sum of mean square error) to predict an attendance and decide whether or not to go to the bar.

5. If the maximum number of weeks/time-steps specified is reached, simulation is halted and the data is designated as a single 'trial'. The process would then start again for the next trial until a specified number of trials are simulated.

The original aim of Fogel et al's paper was to get rid of "dead-weight" strategies, strategies

which consistently fail for predicting the right attendance for the agent [10]. If a strategy is working, even if it fails occasionally, on the whole it will be useful to continually use it. However, if a strategy consistently fails, then the agent is better off replacing it with a new one.

It is also important to note the choice of simple linear functions for the predictors and the importance of considering the history depth as a variable parameter. The former will be elaborated on more below when evaluating the computational complexity of the Ising El-Farol model while the latter is important for investigating multiple ways the agents may formulate strategies based on available information. The extreme cases are when the history depth is 0 and the agents make purely random decisions. Treating the history depth as variable allows for the extraction and interpretation of the agent's behaviour, such as characterizing the 'phase transition' at which agents may switch from informed to seemingly randomly chosen strategies [29]. It is also noteworthy that with the increasing relevance of machine learning and neural networks in science whether the Ising El-Farol can perform a minimum energy search as well as a trained neural network due to the use of low-computationally-costing linear functions.

We have implemented a modified version of this algorithm. In our version any predictions made by an agent that were greater than the total number of agents, as predicted by an agent using Equation 2.8 were instead changed so that the agent's prediction was the same prediction *modulo* N instead of simply rounding down to the number of agents total (agents

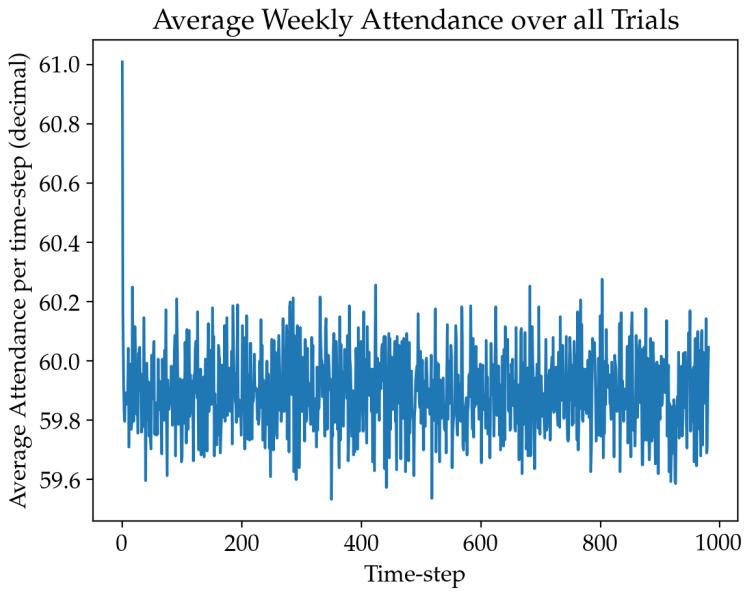

Figure 2.3: Benchmark results for the Stochastic implementation of the El-Farol protocol: Mean Attendance across 20 trials.

cannot predict an attendance number more than the maximum number of agents partaking in the simulation/game). This "tweak" did not change the essence of the dynamics but in our computational experiments it provided a better agreement with the original work in the limiting probability distribution and allowed more varied and different attendance states to be explored as per Figures 2.3, 2.4, 2.5 and 2.6. We attribute the deviations between our results and the results reported in Ref. [10] to minor differences in the functional forms of the strategies and our implementation of the algorithm due to a partially incomplete presentation of the algorithm in the original paper. All results in this paper were generated with our in-house version of this stochastic El-Farol simulation.

Figure 2.4 is a typical single simulation of the El-Farol bar problem. Each time-step

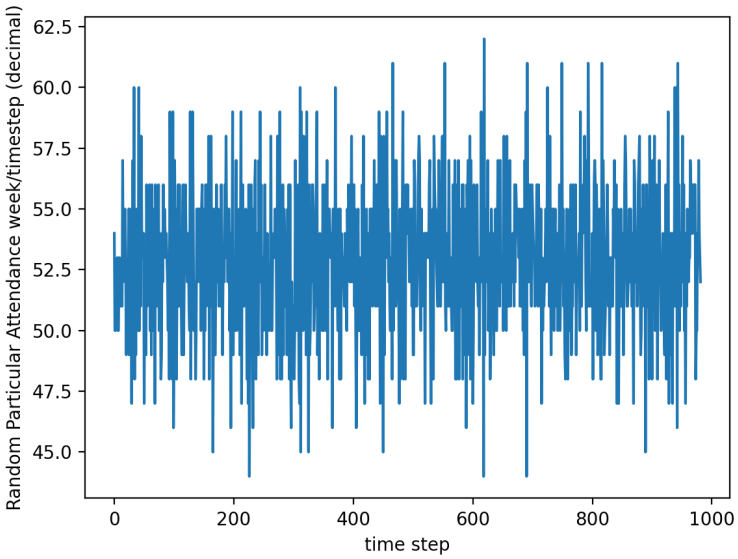

Figure 2.4: Benchmark results for the Stochastic implementation of the El-Farol protocol: Attendance in a representative single trial (no averaging).

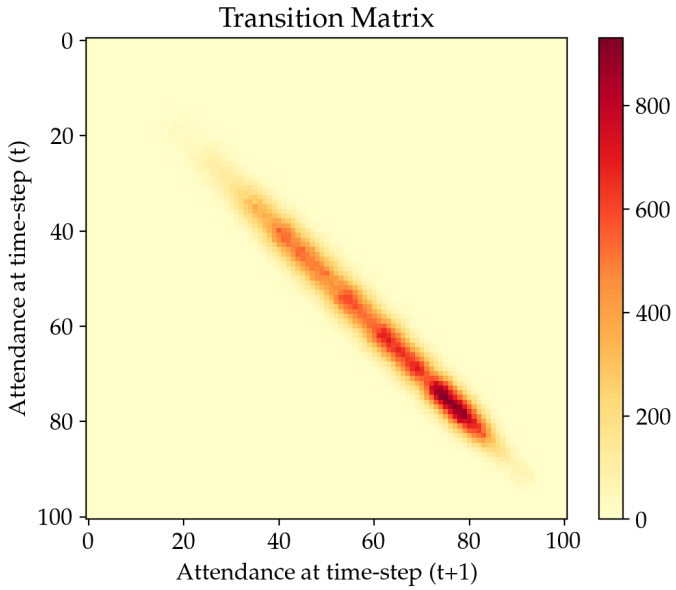

Figure 2.5: Benchmark results for the Stochastic implementation of the El-Farol protocol: Attendance sequence transition matrix (not normalized, side-bar represents frequency of occurrence of attendance number).

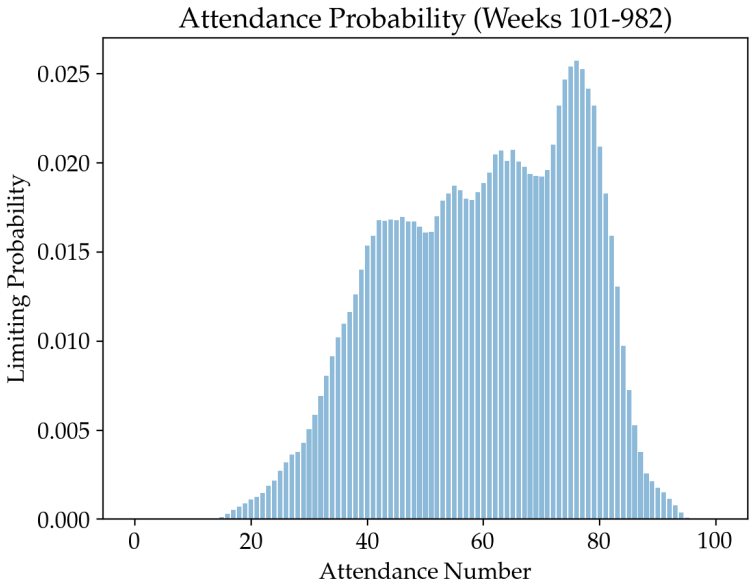

Figure 2.6: Benchmark results for the Stochastic implementation of the El-Farol protocol: The limiting probability distribution of Attendance.

(week), the attendance of the bar is recorded for that given week. The simulation progresses each time-step and another attendance is recorded. The figure demonstrates that the attendance of the bar fluctuates quite drastically, in some weeks reaching as low as 40 people and other weeks as high as 62. The single simulation itself is therefore quite noisy, with a variance of about 10 people. Figure 2.3 shows the result of averaging the attendance over each week results in (across all weeks) many different simulations similar to figure 2.4. Despite the "noisiness" of each individual simulation due to the action of each agent, on the whole an average behaviour is exhibited, as seen in the figure, with the mean attendance across the different simulations being nearly 60 people in attendance, in other words, the chosen threshold. The major difference between this result and those presented

in [10] are the fact that Fogel et al averaged over 300 simulations while here, an average was taken over 20 simulations, more-so to demonstrate that the essential elements needed for constructing the Ising El-Farol program based on the El-Farol scheme were present. The essential elements in terms of the average attendance and threshold are therefore present.

Figure 2.5 represents a non-normalized transition matrix showing how likely the system of agents, given the current attendance of the week (vertical side), will transition to a particular attendance the following week. As an example, if the current attendance for the week was 60, then it is unlikely that the attendance in the following week will be 20 (light yellow) while the attendance the following week is very likely to be somewhere within the range of 55 to 65 people attending. Figure 2.6 gives their limiting probabilities, in other words, how likely a particular attendance is obtained on any given week. As an example, an attendance of 10 is extremely unlikely while an attendance of 45 is much more likely. As noted, discrepancies between the results obtained here and compared to [10] are attributed to minor differences in implementation, but the core element of each simulation is maintained, that is, fluctuation around an average due to 'intelligent' choice of strategy.

In what follows, we focus our attention on the possibility of using this version of the El-Farol protocol to simulate the dynamics of the Ising-type models and locate the global energy minimum by controlling the magnetisation order parameter. Before we present the analogy between the El-Farol agents and the Ising model spins and explore this direction,

we introduce a popular Monte Carlo based algorithm for global minimization. STUN will be used in Chapter 4 to compare the performance of the El-Farol protocol to a popular algorithm commonly used for the task of finding minima of energy functions.

2.4 Stochastic Tunnelling with Markov Chain Monte Carlo (STUN)

One of the biggest problems being tackled by the field of computational chemistry is the determination of energy minima [49,50]. It often occurs in nature that the geometry of a given molecule is one that happens to be optimized with respect to the energy; in other words, it's the geometry that requires the least amount of energy to form or maintain [49, 50]. This has created a wide variety of fields of investigation, namely into understanding the fundamentals of this "optimization process". However, to find an optimum is a maximization, or in this case, minimization problem, namely minimizing the energy. There is, however, no single, uniform way to approach or solve this problem. Therefore the study of energy minimization is a wide and varied problem, with many different theories and algorithms being developed for it and to solve this problem of energy minimization.

One popular approach used in the determination of energy minima for various models, including the Ising model, is an algorithm called ***Stochastic Tunneling*** Monte Carlo, or ***STUN***. Its name is derived from the fact that the algorithm, which has aspects based on the

model of simulated annealing, can become "trapped" in an energy well; however, to escape, the algorithm can "tunnel" out and go into another well with a deeper minimum [51]. This is demonstrated in Figure 2.7. The algorithm itself behaves similarly to the Metropolis-Hastings (MH) algorithm used in Markov Chain Monte Carlo methods, however, wherein MH uses a fixed temperature, STUN uses simulated annealing to bring the temperature down [52]. There are also large differences in the finer details, such as which function is used to calculate the acceptance probability and what parameters must be pre-defined to allow the simulation flexibility to explore different configurations. Such differences are discussed below.

Figure 2.7 shows an interpretation of the potential energy landscape using STUN to find the lowest energy minimum [51]. The system is given an initial "guess" of the landscape (given by the shape in blue) and continually explores and improves its guesses based on previous guesses and previous information obtained from its "exploration" until it finds better energy minima (red line). A new energy minimum is found each time by "tunnelling" to the next well. The estimated PES effectively "shrinks" as new energy minima are found as one continues to tunnel to new minima until it comes close to resembling the actual potential energy surface (black) [51].

The STUN algorithm is based on Monte Carlo methods, specifically the Metropolis-Hastings (MH) algorithm, one key feature of which is the probability of accepting or rejecting a newly proposed state. The algorithm in a given time-step works as follows [51] :

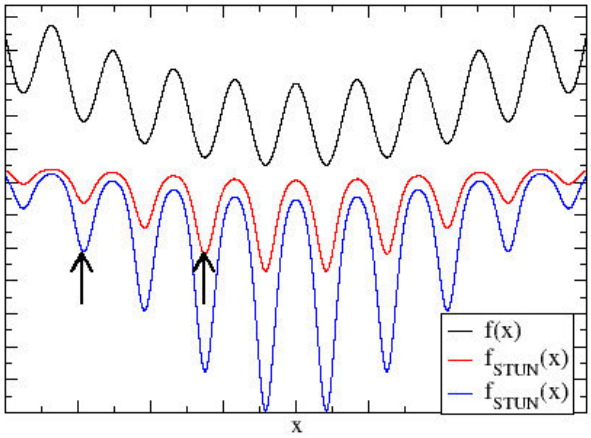

Figure 2.7: An example of STUN, better minima are found (red) from previous minima (blue) until the correct one has been found (black), (© K.Hamacher (2004) / Wikimedia Commons / CC-BY-SA-3.0, from https://commons.wikimedia.org/wiki/File:Stun.jpg).

1. Propose a new state of the system and extract the corresponding variable which forms the basis of the application of the algorithm. For the case of the Ising model, this will be to propose a new spin configuration of the system, whose current energy is then calculated.

2. Using the smallest energy found so far, denoted E_0, calculate f_{STUN}, defined as $f_{STUN} \coloneqq 1 - \exp(-\gamma(E_t - E_0))$ where E_t is the current calculated energy at the given time-step t (refer to Table 2.1), E_0 is the lowest energy found thus far, and γ is a 'tunneling' parameter.

3. The probability of accepting this new state is $\min(1; \exp(\beta \cdot \Delta f_{STUN}))$ where Δf_{STUN} is the change in f_{STUN} from the previous to the current time-step and β is the usual inverse temperature defined by $\beta(T_t) \coloneqq 1/(k_B T_t)$, where k_B is the usual Boltzmann's

constant, taken to be 1 in our case and T_t is the temperature at the current time-step t (refer to Table 2.1 for values used in this experiment).

It is of interest to point out that the requirement for β as a parameter is that STUN behaves in a similar respect to simulated annealing where the temperature is slowly changed, often lowered, at each time-step, hence the temperature T being dependent on the current time-step t (refer to Table 2.1 for values used in this experiment). In this case, $\beta(T)$ was increased by 0.1 (since it is inverse temperature, this is equivalent to lowering the temperature) every time-step and re-set after every 2000 time-steps [51].

γ	Starting $\beta(T)$	f_{STUN}
0.01	1.0	$1 - \exp(-\gamma(E_t - E_0))$

Table 2.1: Parameters used for STUN simulations (note all are unit-less for convenience).

STUN has been successfully applied to a variety of problems in a variety of fields, from statistical physics to protein folding in biochemistry to minimal solution methods in computer science [51–54]. It is therefore a useful benchmark in evaluating the El-Farol algorithm. We implement a naive version of STUN in this thesis to use as such a benchmark.

Chapter 3

Simulating Ising-type Models using the El-Farol Bar Dynamics

3.1 The Ising Model (A Brief Review)

The Ising model is a system of interacting spins first investigated by Ernst Ising in his 1925 thesis [12]. Since then, Ising-type models have become very popular systems of study for fields ranging from physics to computer science [11, 14–17, 54, 55].

Traditionally, the Ising model was a model used to study spin systems and their corresponding total energy via the Hamiltonian

$$\mathcal{H} = J \sum_{<ij>} s_i s_j + h \sum_i s_i \tag{3.1}$$

where s_i represents the direction of the i-th spin (+1 or -1), the sums are over all spins in the model, $< ij >$ is over nearest-neighbors pairs, J is the exchange interaction from the spin pairs and h is a presence from an external magnetic field [15,54] Figure 3.1. is an example of an Ising-type model with more that one nearest neighbor interactions shown for generality. The connecting lines represent the spins' "coupling together" as nearest neighbours and the strength of this interaction is J_{ij}. The models studied often made various assumptions so that they could be simplified, such as constant magnetic fields (so that h is the same for every spin) or constant Js to make obtaining analytic solutions easier. This model has been studied to various degrees in physics as a model for phase transitions [15]. Ernst Ising determined no phase transitions occur in the 1-D model, but they do in two and three dimensions. Many scientists have studied these varying models, most famously Lars Onsager determining an exact solution in 1944 for $h = 0$ for the two-dimensional model [13].

However, the Ising model has become an interest to fields far from traditional physics. The progress of computational techniques and faster machines has enabled studies of a much larger range of Ising models beyond the traditional systems. Over time, the Ising model was developed further into formulations such as various spin models and the Hopfield Network. These models allow for non-uniform exchange interactions and/or external magnetic fields; in other words, they permit varying the interaction and local field parameters for spins separately. The Hamiltonian given in (3.1) can therefore be rewritten as

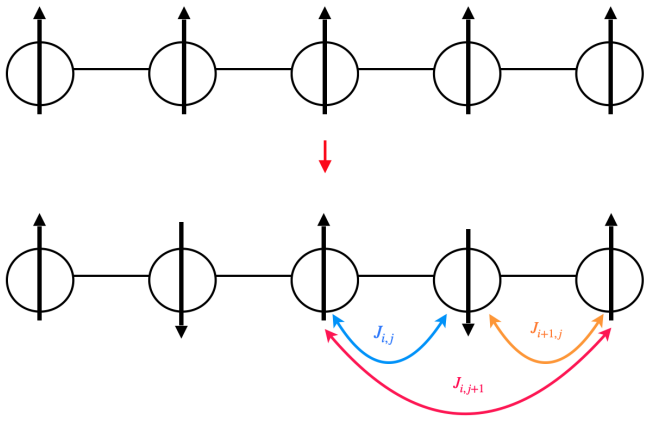

Figure 3.1: An illustration of the spin-flipping process in an Ising-type model with N=5 spins (see the change from top row to bottom row). The interactions are shown to be extended beyond the nearest neighbors. Here the i, j indices denote the nearest neighbors, and $j + 1$ indicates a farther neighbor.

$$\mathcal{H} = \sum_{<ij>} J_{ij} s_i s_j + \sum_i h_i s_i \tag{3.2}$$

Although the "physical" intuition behind the Ising model can become lost since the magnetic field is non-uniform, this Hamiltonian shows that the Ising model can be extended and studied far beyond its traditional scenarios and is a subject worth studying even now, nearly one hundred years later, as new insights regarding the model become relevant.

Since analytic solutions of the Ising model are few and far between, one of the great triumphs of computers is the ability to study Hamiltonians of more complex Ising models given by (3.2) via the use of simulation techniques. One of the most popular methods for studying phase transitions in Ising models has been through the use of Markov Chain

Monte Carlo simulations. The Metropolis-Hastings algorithm is often a popular choice for simulating Ising models and their associated phase transitions [56].

Here, we will use the El-Farol Bar Problem protocol to explore and control the sampling of the energy landscape of two different interacting Ising models with $N = 100$ spins each. We have chosen to focus on models with moderate correlation between the magnetisation and the total energy because in models with strong correlation the configuration of the global minimum is easy to guess.

Next we will explain how we use the El-Farol protocol to scan through the magnetisation of the Ising models. We will present our results which explore the energy space associated with each magnetisation value for the two models in Chapter 4. We will then show that this approach is effective for finding low energy minima by comparing its performance to STUN - a popular Monte Carlo based approach for finding global minima discussed in Chapter 2.

3.2 Simulating the Ising Model Using the El-Farol Protocol

In order to apply the El-Farol protocol to the simulation of the Ising model we need to establish the equivalence between the actions of the El-Farol bar attendees and the dynamics of the Ising spins. This is achieved by adapting and combining the Ising model and the El-Farol protocol. Each 'week' in the El-Farol simulations will now be referred to as a 'time-

step'. The 'agents' in the original El-Farol bar problem, which referred to patrons of the bar, are now the spins in the Ising model. The decision now becomes whether the agents (spins) should flip 'up' (+1) or 'down' (−1) based on whether they correctly 'predict' the controlling factor of the simulation. In the El-Farol bar problem, this control parameter was based on the public information about attendance of the bar that was known and available to every agent of a certain number of previous weeks prior. For the Ising model, attendance is replaced with magnetisation. The controlling factor or the 'attendance threshold', simultaneously, will also be based on the magnetisation: 'magnetisation threshold'.

Thus, the "El-Farol" dynamics for the Ising model boils down to a collection of agents (spins) which at each time-step (week) use strategies to predict the magnetisation for the time-step, and flip their spin based on whether they successfully predicted the magnetisation to be below/above a certain magnetisation threshold. The "El-farol" dynamic works and can successfully be applied even to agents of "zero intelligence" like spin, as already justified by Challet et al [24].

The algorithm, originally presented for the El-Farol bar problem, is then modified as follows:

1. Each spin i of the N *spins* creates an "off-spring" that has the same number of strategies NS as the "parent" spin and in this case $NS = 10$.

    - The number of lag terms for the offspring i was chosen with equal probability from the number of lag terms in its parent j from $\{l_j^i - 1, l_j^i, l_j^i + 1\}$, where l_j^i is

assigned uniformly at random from $\{1, ..., k\}$ at the start (once again, $k = 10$ in our case). If $l_j^i - 1$ was smaller than 1 or $l_j^i + 1$ was larger than 10, they were rounded up or down, respectively.

- Once the number of lag terms needed was known, each of the coefficients $a_{j,k}(t), ..., a_{j,k}(t - l_j^i)$ were assigned by taking the coefficient of the parent and adding a zero-mean Gaussian random variable with standard deviation 0.1 to it (i.e. $N(0, 0.1)$). If the offspring had less terms than the parent, then coefficients were assigned zero in place of what comes from the parent, and if more, then they were given new coefficients taken from $N(0, 0.1)$.

2. The 20 ($2k = 20$) total models/number of strategies were evaluated based on the sum of each strategy's mean square error in predicting system *magnetisation* from the last 10 weeks.

   - 10 is arbitrary, it is just deemed sufficient enough to avoid insufficient sample size

3. The 10 best models, that is the ones with the lowest sum of their mean square errors, were chosen as parents for the next generation.

4. Steps 1-3 are repeated if ten generations had not yet been completed; otherwise each spin used its current best strategy (one with lowest sum of mean square error) to predict a *magnetisation* and decide whether or not to go to flip its spin *up/down*. The total energy of the system for that time-step was then calculated using (3.2)

5. If the maximum number of *time-steps* specified was reached, simulation was halted and the data was designated as a single 'trial'. The process would then start again for the next trial until a specified number of trials were simulated.

Chapter 4

Results. El-Farol-based Search for the Global Minimum in Two Model Ising Spin Systems

The problem of locating the global minimum of a high-dimensional function is very important because virtually any problem may be recast as a problem of minimization/maximization of some quantity of interest. Since the analogy between the Ising model dynamics and the El-Farol Bar Problem established in previous chapter has been made clear, we are now in a position to apply the El-Farol protocol to the problem of finding the global minimum in the energy landscape of an Ising-type spin models. It is to this end we will use the control over the magnetisation achieved in El-Farol protocol

(discussed in the next section) in order to systematically explore the configurational space of our model systems for a given value of magnetisation. The search will be done by varying the magnetisation control 'threshold' in our simulation and thereby driving the system towards increasingly stronger alignment starting from a fully disoriented state. The next section will discuss the relationship between the control of magnetization and the control of the energy in the Ising-type spin systems. Results for the minima search in two model Hamiltonians will follow.

4.1 The El-Farol Control of Magnetisation

The magnetisation control within the El-Farol protocol is very simple: the 'magnetisation threshold' is varied as a function of time. In our case this is done in a straight line connecting the starting point of vanishing magnetisation to the final point of perfect alignment. The strong degree of the achieved control is most surprising due to the simple nature of how it is obtained. Other formulations of the El-Farol bar problem keep the "attendance threshold" fixed [24, 25, 28–31].

One major concern in using El-Farol for the purpose of locating the energy minima is the difficulty of control afforded to the simulation. In the absence of a strong correspondence between magnetisation and the energy of the spin system the quality of sampling depends entirely on the ability of the El-Farol strategies to explore sufficient configurational space. We will demonstrate empirical evidence that this protocol is quite

successful at this task. Furthermore, we note in passing but we do not show that the variance of sampled configurations may be in principle controlled, as discussed in Chapter 2.

To make these ideas more concrete we will use a few examples to discuss the correlation between magnetization and energy in several Ising-type systems. The correlation is very strong in systems whose Hamiltonians contain only a small number of interacting spin pairs (i, j), dubbed "nearest neighbour interactions". These Hamiltonians follow the form of Equation (3.2) but contain much fewer interaction pairs. As an example, for only nearest neighbour interactions, any spin of the system can interact with, at most, only the spin on either side of itself, hence "nearest neighbours". Since the system used here is a 1-Dimensional "chain" of spins, the interaction takes place such that any given spin interacts solely with its neighbours to its left and right (note there are no periodic boundary conditions, so atoms at the opposite ends of the chain will only interact with atoms next to them and not each other). A visual representation in shown by the blue and orange arrows in figure 3.1, where the an example of a "nearest-neighbour" interaction is the spin-down site interacting with its neighbours to the left and right (blue and orange arrows, respectively). A similar Hamiltonian is made for maximum two nearest neighbours, denoted $\mathcal{H}_{NN2}$ and for 3, denoted $\mathcal{H}_{NN3}$, respectively. Table 4.1 shows various Pearson correlation coefficients between the energy of the system and the magnetisation obtained using a simplified driven El-Farol sampling procedure for the Hamiltonians $\mathcal{H}_{NN1}$, $\mathcal{H}_{NN2}$,

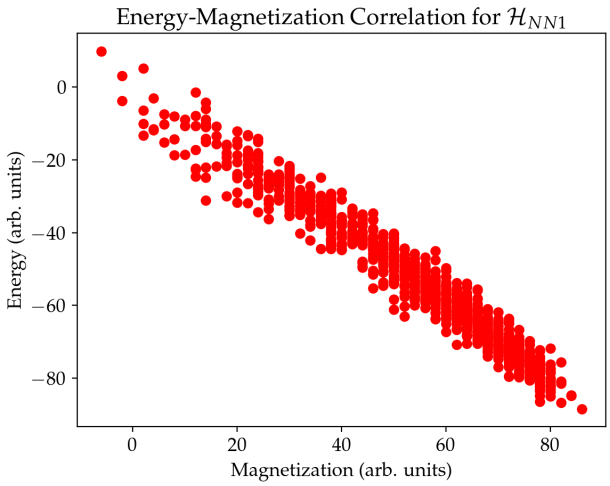

Figure 4.1: In H_{NN1} the correlation between the energy and magnetisation is strong due to small values of J's and less pair interactions.

$\mathcal{H}_{NN3}$,$\mathcal{H}_1$, and $\mathcal{H}_2$ respectively. The last two in this list will be used in our simulations: they contain N=100 spins with $N/4$ and $N/2$ number of nearest neighbours each (with solid cut-off boundaries at either end) and the correlation between the magnetization and the energy is quite poor.

Nearest Neighbours	Correlation (arb. units)
1	0.977
2	0.913
3	0.947
$N/4 = 25$	0.500
$N/2 = 50$	0.655

Table 4.1: Comparison of correlation of energy and magnetization found by El-Farol for various number of nearest-neighbour interactions

This makes sense from a basic standpoint as well because following Equation (3.2), if there are few nearest neighbour interactions involved, then the dominant term in the Hamiltonian

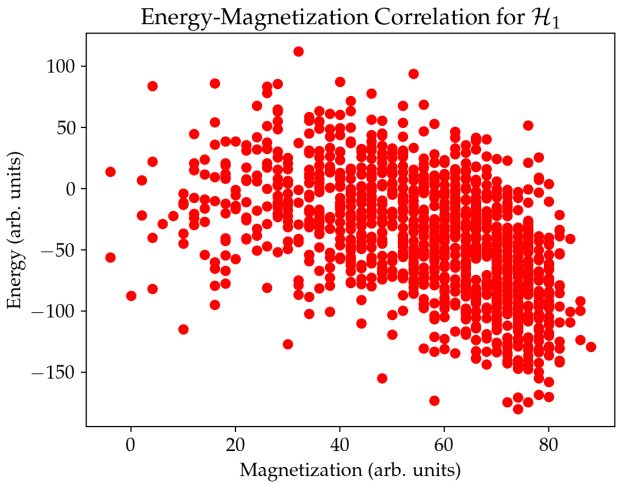

Figure 4.2: In $\mathcal{H}_1$ the correlation between the energy and magnetisation is moderate (0.5) due to larger values of J's and longer range of interactions.

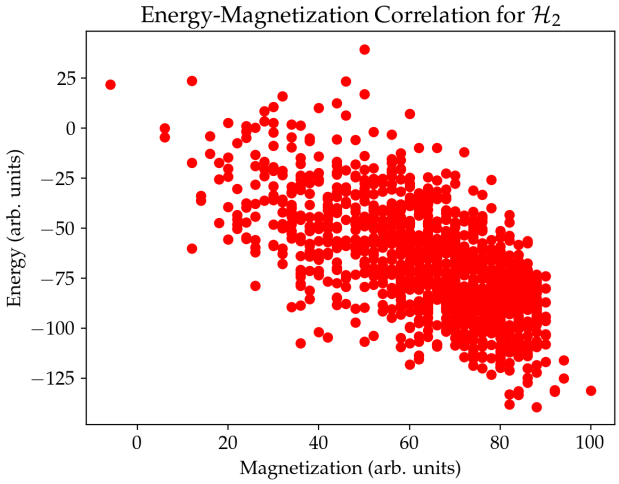

Figure 4.3: In $\mathcal{H}_2$ the correlation between the energy and magnetisation is moderate (0.65) due to more and longer-range interactions

will be $\sum_i h_i s_i$, a single body term that is related to the magnetisation $\sum_i s_i$. However if $|J|$ is increased and/or the number of nearest neighbour interactions is increased, then $\sum_{<ij>} J_{ij} s_i s_j$ becomes more dominant and the correlation between magnetization and energy breaks down.

Interestingly, despite this obstacle, we will show that El-Farol performs well even when compared with a naive STUN implementation (see Chapter 2 for details) and even with the increased difficulty of control, as each spin in $\mathcal{H}_1$ has up to 25 nearest neighbours and in $\mathcal{H}_2$ each spin has 50 nearest neighbors. As seen from the correlation coefficients for these two cases: 0.5 and 0.65 respectively (see Table 4.1) the correlation between magnetisation and energy in the two systems we chose is moderate.

4.2 Results for Hamiltonian 1: Strong Medium-ranged Exchange Interactions

The Ising-type Hamiltonian presented in this section follow the form of (3.2):

$$\mathcal{H} = \sum_{<ij>} J_{ij} s_i s_j + \sum_i h_i s_i$$

where s_i represents the direction of the i-th spin (+1 or -1), the sums are over all spins in the model, $< ij >$ is over all possible interacting pairs, J_{ij} is the exchange interaction from the spin pairs and h_i is a local field.

One interesting Hamiltonian to use in order to compare the global minimization performance of the El-Farol protocol to STUN is a Hamiltonian $\mathcal{H}_1$ with a different exchange interaction terms J_{ij} for every pair of spins i, j such that:

1. $J_{i,j} \neq J_{k,l}$ for any pair of pairs of spins (i, j) and (k, l), $i \neq k, l$ or $j \neq k, l$
2. $0 \leq |J_{ij}| \leq 3$ assigned randomly (note these values in this example and below stay fixed throughout the simulation once assigned).
3. $h_i = 1$ for each i.

This Hamiltonian will be denoted as $\mathcal{H}_1$, with the number of nearest neighbours/interactions for each spin going to a maximum of $N/4$ where in this simulation, $N = 100$ spins (providing a number of interactions for up to the maximum of 25 nearest neighbouring spins). The link to the online arxiv containing python binaries with the specific $\mathcal{H}_1$ 100x100 array used in our numerical examples is provided in the Appendix.

An important note should be made about the rate of change of the threshold from time-step to time-step. It is crucial to pick a rate of change for the threshold that is not too low, or else the algorithm will become inefficient, as the same state will likely end up being explored a fair number of times before the system explores states with potentially lower minima. This would cause the algorithm, and the Ising El-Farol scheme, to perform extremely slowly. Simultaneously, if the rate of change is too large, that is the change is the threshold is too drastic from time-step to time-step, the system no longer has enough to explore different

states and several interesting and even crucial states with possibly lower energy minima than those found closer to the end of the simulation may be missed. Therefore an investigation into determining whether there is an optimal rate of change for the threshold, as well as studying the effects of varying the rate of change for the threshold, would likely be very fruitful. In the present case of the El-Farol Ising model, a conservative rate of change was chosen so that the system had more time to explore states in order to have a better chance at finding a truly global minimum, but no investigation was made into whether this was the optimal choice for the rate of change of the threshold.

The driving of the system by driving the magnetisation threshold from 0 to 100 towards lower energies is the general trend of Figure 4.4. There is a downward trend in which the system begins to explore lower energy states as the simulation progresses. This is the result of the magnetisation manipulation implemented in our El-Farol algorithm which opens a degree of control on which configurations and related energies the system can explore by exploiting the threshold which is based on the magnetisation. Thus by changing the threshold at each time-step, in this case by changing the magnetisation of the system in a controlled way, the El-Farol algorithm is able to be "directed" towards a goal such as exploring and finding minimum energy states and provides a great degree of control. This is compared to Figure 4.5 where it is clear that such a degree of control is not necessarily as evident in this implementation of the naive STUN protocol.

The results of our simulations show that the El-Farol protocol finds a close but slightly

lower minimum energy in a significantly smaller number of time-steps than this implementation of STUN. The lowest energy minimum found by either algorithm can be found in Table 4.2. However, what is of more interest in Figures 4.4 and 4.5 are the fact that one can make two very clear and distinct observations. The first one is that the number of time-steps required to find an energy minimum using El-Farol is significantly smaller than the number needed to find it using this naive implementation of STUN: Figure 4.4 shows that its smallest minimum was found at time-step 942 while Figure 4.5 only finds such a minimum at time-step $420,011$. This is a difference of two orders of magnitude increase in the number of steps needed by this naive implementation of STUN as compared to El-Farol. The second advantage is that Ising El-Farol finds a lower global minimum than this implementation of STUN as shown in Table 4.2. This simulation illustrates the potential utility of the El-Farol protocol for finding energy minima. The complexity analysis is left for future work.

Algorithm	Minimum Energy (arb. units)
El-Farol	-180.124
STUN	-179.986

Table 4.2: Energy minima found by El-Farol and STUN for $\mathcal{H}_1$.

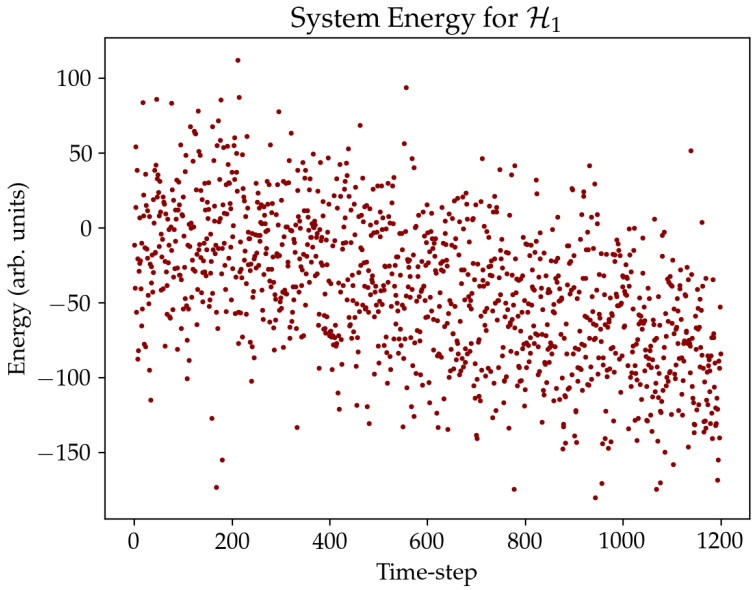

Figure 4.4: Energies for configurations of N=100 spin system with Hamiltonian $\mathcal{H}_1$ visited by the driven El-Farol protocol for 1200 time-steps.

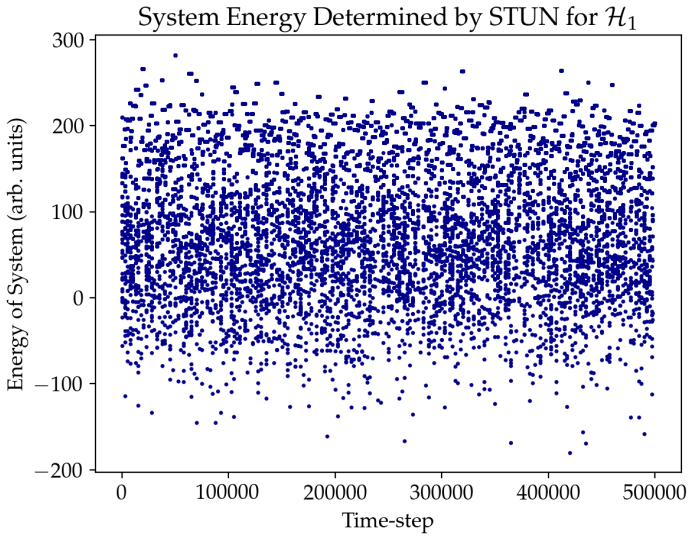

Figure 4.5: Energies for configurations of N=100 spin system with Hamiltonian $\mathcal{H}_1$ visited by STUN in $500,000$ time-steps.

4.3 Results for Hamiltonian 2: Medium Strength Long-range Exchange Interactions

Another interesting Hamiltonian to compare the performance of the El-Farol protocol to our naive implementation of STUN is a Hamiltonian $\mathcal{H}_2$ with a **large** number (N/2 = 50) of different exchange interaction terms J such that:

1. $J_{i,j} \neq J_{k,l}$ for any pair of pairs of spins (i,j) and (k,l), $i \neq k,l$ or $j \neq k,l$
2. $0 \leq |J_{ij}| \leq 1$ assigned randomly
3. The number of nearest neighbours/interactions for each spin is a maximum of $N/2$ where in this simulation, $N = 100$ spins (providing a number of interactions for up to the maximum of 50 nearest neighbouring spins).
4. $h_i = 1$ for each i.

Algorithm	Minimum Energy (arb. units)
El-Farol	-139.268
STUN	-119.118

Table 4.3: Comparison of minimum energy found by El-Farol and STUN for $\mathcal{H}_2$.

This Hamiltonian will be denoted as $\mathcal{H}_2$. The link to the online arxiv containing python binaries with the specific $\mathcal{H}_2$ 100x100 array used in our numerical examples is provided in the Appendix. Table 4.3 gives the minimum energy found by both algorithms.

A similar conclusion as before may be drawn in that the El-Farol algorithm is able to find the minimum energy and finds a similar one just like STUN in a feasible amount of time.

The driving of the system towards lower energies is the general trend of Figure 4.6. Once again, a downward trend in which the system begins to explore lower energy states as the simulation progresses is observed. This is again the result of the magnetisation manipulation implemented in our El-Farol algorithm which opens a degree of control on which configurations and related energies the system can explore by exploiting the threshold which is based on the magnetisation. This is once again compared to Figure 4.7 where it is clear that such a degree of control is not nearly as evident.

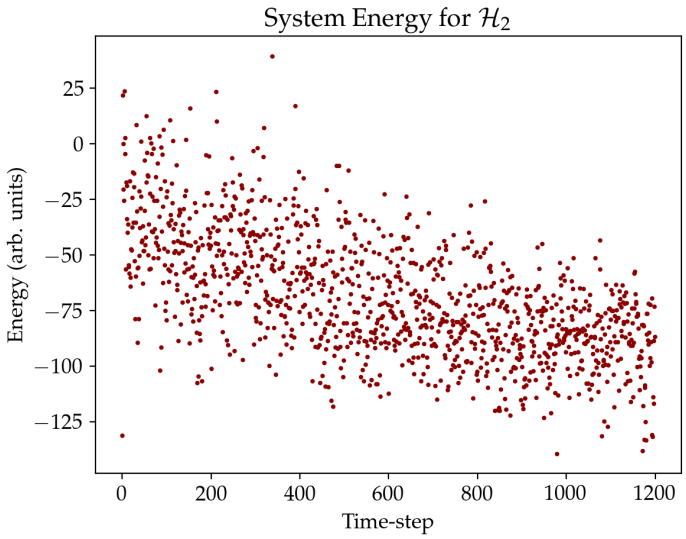

Figure 4.6: Energies for configurations of N=100 spin system with Hamiltonian $\mathcal{H}_2$ visited by the driven El-Farol protocol for 1200 time-steps.

In contrast to the minor observation in the previous example based on Hamiltonian $\mathcal{H}_1$ these results show that the El-Farol protocol finds a significantly lower minimum energy in

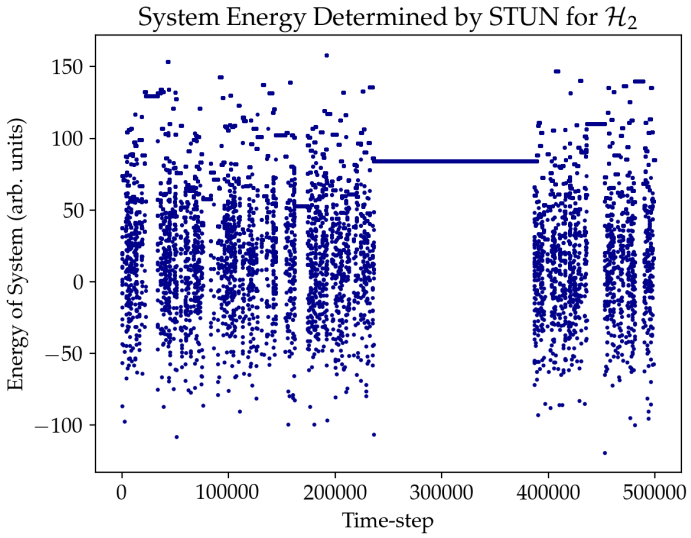

Figure 4.7: Energies for configurations of N=100 spin system with Hamiltonian $\mathcal{H}_2$ visited by STUN in $500,000$ time-steps.

a significantly smaller number of time-steps than this implementation of STUN. The lowest energy minimum found by either algorithm can be found in Table 4.3. Figures 4.6 and 4.7 reaffirm that the number of time-steps required to find an energy minimum using El-Farol is orders of magnitude smaller than the number of steps needed to find it using this naive version of STUN. Figure 4.6 shows that its smallest minimum was found at time-step 979 while Figure 4.7 only finds such a minimum at time-step $453,035$. This simulation illustrates the potential utility of the El-Farol protocol for finding energy minima even in complex energy landscapes such as $\mathcal{H}_2$. The full complexity analysis of the El-Farol energy minimization scheme is left to future work using principles as laid out in [57–61].

Chapter 5

Conclusions and Outlook

5.1 Conclusions

There are different algorithms available for finding the energy minima of particular systems and many algorithms are suitably optimized for a particular class of systems [42]. The investigation of the El-Farol bar protocol and its subsequent application to the problem of determining global energy minima in Ising-type models has proven quite fruitful and promising. Our preliminary calculations suggest that the El-Farol protocol performs at competitive levels and is capable of locating deep minima in Ising-type model energy landscapes. It has been compared with and tested against the widely-implemented STUN Monte Carlo algorithm and has shown to be able to find energy minima in much fewer number of time-steps than our single conventional implementation of STUN. Whether it is

comparable on the whole to the class of STUN methods themselves remains a further topic of research. It also provides a reasonable degree of a control over systems with good correlation between the system's magnetization and total energy and still performs competitively with other algorithms for systems in which the correlation is much weaker. It therefore marks itself as a useful algorithm for studying these class of systems, especially systems with moderate correlation between energy and magnetization.

Finally, it is important to stress the Ising model's utility as a system worth studying in general, especially, but not limited to, models in theoretical and real chemical domains. Since it has been used originally to investigate phase transitions and ferromagnetism [17], future studies of other chemical systems modeled by spin systems may benefit from estimating their energy using the Ising El-Farol protocol. The 1-dimensional Ising model has found interesting applications today in studies of nano-scale phenomenon, therefore it is important to have models that can accurately predict the energy of a given spin system, especially when comparing with other analytic models. [62–64] Once again though, the simplicity of the model has shown that in can be applied across and studied over a wide range of fields, quantum information [18, 19]; in the modelling of protein folding [20, 21]; even in computer science [11, 14, 16, 22, 23]. For these reasons a generic Ising-model-like energy landscape is the perfect toy-model to be explored for algorithmic testing.

5.2 Outlook

One interesting feature of the above implementation of El-Farol lies in the fact that the threshold was continually changed at each time-step to control the configurations explored by the system [24, 32–34]. We have explained before that this is not typically part of the El-Farol bar problem, nor for minority games in general, as the threshold is assumed to be fixed. This has interesting implications, especially for the analysis of the statistical mechanics behind the El-Farol problem. The works of Challet et al find that agents playing the game with random information as per determination of strategies via Equation 2.8 (i.e. Equation 2.3) will play strategy s with frequency $f_{i,s}$ (i is the i-th agent) that minimize the quantity ([24])

$$H = \sum_{\nu=1}^{P} \rho^{\nu}(\langle A|\nu\rangle - L)^2 \tag{5.1}$$

where P is an integer arising from the fact that because it is random information, any of the 2^m histories may be selected, L is the *threshold*, ρ^{ν} is the probability that the information (that is, the recorded history of how well a strategy $a_{i,s}^{\mu(t)}$ did in predicting the attendance for the current time-step 't') $\mu(t) = \nu$, or $\rho^{\nu} = Prob\{\mu(t) = \nu\} = 1/P$ and

$$\langle A|\nu\rangle = \sum_{i=1}^{N}\sum_{s_i}^{S} f_{s,i}a_{s,i}^{\nu}$$

which is the average of attendance $A(t)$ conditional on the event $\mu(t) = \nu$ [24].

They derive the following by treating Equation 5.1 as a Hamiltonian with minima. They found that the predictability is given by

$$\frac{H}{N} = \frac{\sqrt{l(1-l)}}{2} \cdot \frac{1 + Q(\zeta) + 2\gamma^2/\alpha}{(1+\chi(\zeta))^2} \tag{5.2}$$

and the fluctuation by

$$\sigma^2 = H + N\frac{\sqrt{l(1-l)}}{2}(1 - Q(\zeta)) + N\Sigma \tag{5.3}$$

Σ, $Q(\zeta)$ and $\chi(\zeta)$ are values given by equations in [24] and dependent on the parameter γ and α, where $\alpha = \frac{P}{N}$ but are inconsequential to the analysis at hand. The main concern is l, given by

$$l \equiv \frac{L}{N} \tag{5.4}$$

as Equations 5.2 and 5.3 are usually analysed in the limit $N \to \infty$ and in the presence of small fluctuations around L [24]. An important limit mentioned but not completely analysed in [24] is a continually varying threshold, as many of the key results derived in that paper focus on a fixed value of L. An analysis may be taken for non-fixed L and even one step further in the limit of $L \to \infty$ so that l may no longer be finite.

Appendix A

Numerical parameters for the spin models used in this thesis

A.1 Hamiltonian 1

The matrix equation that a matrix containing all values used for the exchange interactions in subsequent calculations between any sites i and j ($i \neq j$, since self-interactions were not allowed) is given by

$$(\mathcal{H}_1)_{i,j} = I_{+/-}X + 2I_{+/-}X \tag{A.1}$$

where $I_{+/-} \in \{-1, 1\}$ and $X \in [0, 1)$.

The full matrix with all values is available for download on figshare server:

https://doi.org/10.6084/m9.figshare.14416094.v1

A.2 Hamiltonian 2

The matrix equation that a matrix containing all values used for the exchange interactions in subsequent calculations between any sites i and j ($i \neq j$, since self-interactions were not allowed) is given by

$$(\mathcal{H}_1)_{i,j} = 2I_{+/-}X + 2I_{+/-}X \tag{A.2}$$

where $I_{+/-} \in \{-1, 1\}$ and $X \in [0, 1)$.

The full matrix with all values is available for download on figshare server https://doi.org/10.6084/m9.figshare.14416094.v1

Bibliography

[1] F. Cazals, T. Dreyfus, D. Mazauric, C.-A. Roth, and C. H. Robert, "Conformational ensembles and sampled energy landscapes: Analysis and comparison," *Journal of computational chemistry*, vol. 36, no. 16, pp. 1213–1231, 2015.

[2] J. M. Carr, D. Mazauric, F. Cazals, and D. J. Wales, "Energy landscapes and persistent minima," *The Journal of chemical physics*, vol. 144, no. 5, p. 054109, 2016.

[3] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, "Transition path sampling: Throwing ropes over rough mountain passes, in the dark," *Annual review of physical chemistry*, vol. 53, no. 1, pp. 291–318, 2002.

[4] F. Pietrucci, "Strategies for the exploration of free energy landscapes: Unity in diversity and challenges ahead," *Reviews in Physics*, vol. 2, pp. 32–45, 2017.

[5] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*. Elsevier, 2001, vol. 1.

[6] D. Wales *et al.*, *Energy landscapes: Applications to clusters, biomolecules and glasses*. Cambridge University Press, 2003.

[7] D. J. Wales and T. V. Bogdan, "Potential energy and free energy landscapes," *The Journal of Physical Chemistry B*, vol. 110, no. 42, pp. 20 765–20 776, 2006.

[8] D. Devaurs, K. Molloy, M. Vaisset, A. Shehu, T. Siméon, and J. Cortés, "Characterizing energy landscapes of peptides using a combination of stochastic algorithms," *IEEE transactions on nanobioscience*, vol. 14, no. 5, pp. 545–552, 2015.

[9] W. B. Arthur, "Inductive reasoning and bounded rationality," *American Economic Review*, vol. 84, no. 2, pp. 406–11, 1994. [Online]. Available: https://EconPapers.repec.org/RePEc:aea:aecrev:v:84:y:1994:i:2:p:406-11

[10] D. B. Fogel, K. Chellapilla, and P. J. Angeline, "Inductive reasoning and bounded rationality reconsidered," *IEEE transactions on evolutionary computation*, vol. 3, no. 2, pp. 142–146, 1999.

[11] K. Mills, P. Ronagh, and I. Tamblyn, "Controlled online optimization learning (cool): Finding the ground state of spin hamiltonians with reinforcement learning," *arXiv preprint arXiv:2003.00011*, 2020.

[12] E. Ising, "Beitrag zur theorie des ferromagnetismus," *Zeitschrift für Physik*, vol. 31, no. 1, pp. 253–258, 1925.

[13] L. Onsager, "Crystal statistics. i. a two-dimensional model with an order-disorder transition," *Physical Review*, vol. 65, no. 3-4, p. 117, 1944.

[14] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the national academy of sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.

[15] M. Toda, R. Kubo, N. Saitō, N. Hashitsume *et al.*, *Statistical physics I: equilibrium statistical mechanics*. Springer Science & Business Media, 1991, vol. 1.

[16] P. P. Martin, *Potts models and related problems in statistical mechanics*. World Scientific, 1991, vol. 5.

[17] D. Ruelle, "Statistical mechanics of a one-dimensional lattice gas," *Communications in Mathematical Physics*, vol. 9, no. 4, pp. 267–278, 1968.

[18] P. Štelmachovič and V. Bužek, "Quantum-information approach to the ising model: Entanglement in chains of qubits," *Physical Review A*, vol. 70, no. 3, p. 032313, 2004.

[19] C.-c. Liu, D. Wang, W.-y. Sun, and L. Ye, "Quantum fisher information, quantum entanglement and correlation close to quantum critical phenomena," *Quantum Information Processing*, vol. 16, no. 9, pp. 1–15, 2017.

[20] T. Aksel and D. Barrick, "Analysis of repeat-protein folding using nearest-neighbor statistical mechanical models," *Methods in enzymology*, vol. 455, pp. 95–125, 2009.

[21] A. Bakk and J. S. Høye, "One-dimensional ising model applied to protein folding," *Physica A: Statistical Mechanics and its Applications*, vol. 323, pp. 504–518, 2003.

[22] R. B. Potts, "Some generalized order-disorder transformations," in *Mathematical proceedings of the cambridge philosophical society*, vol. 48, no. 1. Cambridge University Press, 1952, pp. 106–109.

[23] F.-Y. Wu, "The potts model," *Reviews of modern physics*, vol. 54, no. 1, p. 235, 1982.

[24] D. Challet, M. Marsili, and G. Ottino, "Shedding light on el farol," *Physica A: Statistical Mechanics and Its Applications*, vol. 332, pp. 469–482, 2004.

[25] M. Marsili, D. Challet, and R. Zecchina, "Exact solution of a modified el farol's bar problem: Efficiency and the role of market impact," *Physica A: Statistical Mechanics and its Applications*, vol. 280, no. 3-4, p. 522–553, Jun 2000. [Online]. Available: http://dx.doi.org/10.1016/S0378-4371(99)00610-X

[26] B. K. Chakrabarti, "Kolkata restaurant problem as a generalised el farol bar problem," in *Econophysics of Markets and Business Networks*. Springer, 2007, pp. 239–246.

[27] A. J. Collins, "Strategic group formation in the el farol bar problem," in *Complex Adaptive Systems*. Springer, 2019, pp. 199–211.

[28] D. Challet, M. Marsili, and Y.-C. Zhang, "Minority games: Interacting agents in financial markets," Oxford University Press, Tech. Rep., 2004.

[29] A. Coolen, *The Mathematical Theory of Minority Games. Statistical Mechanics of Interacting Agents*. Oxford University Press, 2005.

[30] M. De Cara, O. Pla, and F. Guinea, "Competition, efficiency and collective behavior in the "el farol" bar model," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 10, no. 1, pp. 187–191, 1999.

[31] H. Luş, C. O. Aydın, S. Keten, H. İ. Ünsal, and A. R. Atılgan, "El farol revisited," *Physica A: Statistical Mechanics and Its Applications*, vol. 346, no. 3-4, pp. 651–656, 2005.

[32] D. Challet, M. Marsili, and R. Zecchina, "Statistical mechanics of systems with heterogeneous agents: Minority games," *Physical Review Letters*, vol. 84, no. 8, p. 1824, 2000.

[33] D. Challet, M. Marsili, and Y.-C. Zhang, "Stylized facts of financial markets and market crashes in minority games," *Physica A: Statistical Mechanics and its Applications*, vol. 294, no. 3-4, pp. 514–524, 2001.

[34] M. Marsili, D. Challet, and R. Zecchina, "Exact solution of a modified el farol's bar problem: Efficiency and the role of market impact," *Physica A: Statistical Mechanics and its Applications*, vol. 280, no. 3-4, pp. 522–553, 2000.

[35] D. Challet and Y.-C. Zhang, "On the minority game: Analytical and numerical studies," *Physica A: Statistical Mechanics and its applications*, vol. 256, no. 3-4, pp. 514–532, 1998.

[36] E. Moro, "The minority game: an introductory guide," *arXiv preprint cond-mat/0402651*, 2004.

[37] A. Cavagna, "Irrelevance of memory in the minority game," *Physical Review E*, vol. 59, no. 4, p. R3783, 1999.

[38] D. Challet and M. Marsili, "Phase transition and symmetry breaking in the minority game," *Physical Review E*, vol. 60, no. 6, p. R6271, 1999.

[39] D. Whitehead, "The El Farol Bar Problem Revisited: Reinforcement Learning in a Potential Game," Edinburgh School of Economics, University of Edinburgh, Edinburgh School of Economics Discussion Paper Series 186, Sep. 2008. [Online]. Available: https://ideas.repec.org/p/edn/esedps/186.html

[40] J. Tanimoto, *Fundamentals of evolutionary game theory and its applications*. Springer, 2015.

[41] N. F. Johnson and P. M. Hui, "Crowd-anticrowd theory of collective dynamics in competitive, multi-agent populations and networks," *arXiv preprint cond-mat/0306516*, 2003.

[42] K.-H. Hoffmann and M. Schreiber, *Computational statistical physics: from billiards to Monte Carlo*. Springer Science & Business Media, 2002.

[43] S. St. Luce and H. Sayama, "Phase spaces of the strategy evolution in the el farol bar problem," in *Artificial Life Conference Proceedings*. MIT Press, 2020, pp. 558–566.

[44] M. Hart, P. Jefferies, N. Johnson, and P. Hui, "Crowd–anticrowd theory of the minority game," *Physica A: Statistical Mechanics and its Applications*, vol. 298, no. 3-4, pp. 537–544, 2001.

[45] T. Kalinowski, H.-J. Schulz, and M. Briese, "Cooperation in the minority game with local information," *Physica A: Statistical Mechanics and its Applications*, vol. 277, no. 3-4, pp. 502–508, 2000.

[46] N. Johnson, P. Hui, D. Zheng, and C. Tai, "Minority game with arbitrary cutoffs," *Physica A: Statistical Mechanics and its Applications*, vol. 269, no. 2-4, pp. 493–502, 1999.

[47] C. Kittel, "Introduction to solid state physics eighth edition," pp. 467–480, 2005.

[48] G. Jaeger, "The ehrenfest classification of phase transitions: introduction and evolution," *Archive for history of exact sciences*, vol. 53, no. 1, pp. 51–81, 1998.

[49] F. Jensen, *Introduction to computational chemistry*. John wiley & sons, 2017.

[50] B. T. Sutcliffe, "The idea of a potential energy surface," *Molecular Physics*, vol. 104, no. 5-7, pp. 715–722, 2006.

[51] W. Wenzel and K. Hamacher, "Stochastic tunneling approach for global minimization of complex potential energy landscapes," *Physical Review Letters*, vol. 82, no. 15, p. 3003, 1999.

[52] K. Hamacher, "Adaptation in stochastic tunneling global optimization of complex potential energy landscapes," *EPL (Europhysics Letters)*, vol. 74, no. 6, p. 944, 2006.

[53] A. Schug, T. Herges, and W. Wenzel, "Reproducible protein folding with the stochastic tunneling method," *Physical review letters*, vol. 91, no. 15, p. 158102, 2003.

[54] S. Friedli and Y. Velenik, *Statistical mechanics of lattice systems: a concrete mathematical introduction*. Cambridge University Press, 2017.

[55] F. Barahona, "On the computational complexity of ising spin glass models," *Journal of Physics A: Mathematical and General*, vol. 15, no. 10, p. 3241, 1982.

[56] K. Binder, "Applications of monte carlo methods to statistical physics," *Reports on Progress in Physics*, vol. 60, no. 5, p. 487, 1997.

[57] L. T. Wille and J. Vennik, "Computational complexity of the ground-state determination of atomic clusters," *Journal of Physics A: Mathematical and General*, vol. 18, no. 8, p. L419, 1985.

[58] J. T. Ngo and J. Marks, "Computational complexity of a problem in molecular structure prediction," *Protein Engineering, Design and Selection*, vol. 5, no. 4, pp. 313–321, 1992.

[59] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671–680, 1983.

[60] A. M. Gutin, V. I. Abkevich, and E. I. Shakhnovich, "Chain length scaling of protein folding time," *Physical Review Letters*, vol. 77, no. 27, p. 5433, 1996.

[61] K. Hamacher and W. Wenzel, "Scaling behavior of stochastic minimization algorithms in a perfect funnel landscape," *Physical Review E*, vol. 59, no. 1, p. 938, 1999.

[62] T. Kaneyoshi, "Compensation point phenomena in nanoscale ising particles with high critical temperature," *Phase Transitions*, vol. 93, no. 8, pp. 826–842, 2020.

[63] A. Kaplan and S. Volkov, "Nanoscale stratification of local optical fields in low-dimensional atomic lattices," *Physical review letters*, vol. 101, no. 13, p. 133902, 2008.

[64] ——, "Nanoscale stratification of optical excitation in self-interacting one-dimensional arrays," *Physical Review A*, vol. 79, no. 5, p. 053834, 2009.