




ELSEVIER

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2009 Special Issue

Integrated feature and parameter optimization for an evolving spiking neural network: Exploring heterogeneous probabilistic models

Stefan Schliebs^{a,*}, Michaël Defoin-Platel^b, Sue Worner^c, Nikola Kasabov^a

^a Knowledge Engineering and Discovery Research Institute (KEDRI), Auckland University of Technology, New Zealand

^b Biomathematics and Bioinformatics at Rothamsted Research, United Kingdom

^c Lincoln University, Centre for Bioprotection, New Zealand

ARTICLE INFO

Article history:

Received 6 May 2009

Received in revised form 3 June 2009

Accepted 25 June 2009

Keywords:

Evolving spiking neural network

Quantum-inspired evolutionary algorithm

Multiple probabilistic model

Estimation of distribution algorithm

ABSTRACT

This study introduces a quantum-inspired spiking neural network (QisSNN) as an integrated connectionist system, in which the features and parameters of an evolving spiking neural network are **optimized** together with the use of a quantum-inspired evolutionary algorithm. We propose here a novel **optimization** method that uses different representations to explore the two search spaces: A binary representation for **optimizing** feature subsets and a continuous representation for evolving appropriate real-valued configurations of the spiking network. The properties and characteristics of the improved framework are studied on two different synthetic benchmark datasets. Results are compared to traditional methods, namely a multi-layer-perceptron and a naïve Bayesian classifier (NBC). A previously used real world ecological dataset on invasive species establishment prediction is revisited and new results are obtained and **analyzed** by an ecological expert. The proposed method results in a much faster convergence to an optimal solution (or a close to it), in a better accuracy, and in a more informative set of features selected.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Recently spiking neural networks (SNN) (Gerstner & Kistler, 2002; Izhikevich, 2003) have been developed as biologically plausible connectionist models, which use trains of spikes for internal information representation. It was argued that SNN have at least similar computational power than the traditional Multi-Layer-Perceptron derivatives (Maass, 1999). Nowadays many studies attempt to use Spiking Neural Networks (SNN) for practical applications, some of them demonstrating very promising results on solving complex real world problems. Substantial progress has been made in areas like speech recognition (Verstraeten, Schrauwen, & Stroobandt, 2005), learning rules (Bohte, Kok, & Poutré, 2002), associative memory (Knoblauch, 2005), and function approximation (Iannella & Kindermann, 2005), just to name a few. Based on Kasabov (2007) an evolving spiking neural network was proposed and applied to audio-visual pattern recognition (Wysoski, Benuskova, & Kasabov, 2006, 2008). A similar type of network was later used in the context of a taste recognition task (Soltic, Wysoski, & Kasabov, 2008).

With encouraging results, spiking neural networks were **presented** in the context of a feature selection problem (Schliebs, Defoin-Platel, & Kasabov, 2009). In this work a binary state-of-art **optimization** algorithm, namely the Versatile Quantum-inspired Evolutionary Algorithm (vQEA) (Defoin-Platel, Schliebs, & Kasabov, 2007), was combined with an Evolving Spiking Neural Networks (eSNN) (Wysoski et al., 2006). Through implementing quantum principles, vQEA evolves in parallel a number of independent probability vectors, that may interact at certain intervals with each other, forming a multi-model Estimation of Distribution Algorithm (EDA) (Defoin-Platel, Schliebs, & Kasabov, in press).

Following the wrapper approach, vQEA was used to identify relevant feature subsets and simultaneously evolve an optimal eSNN parameter setting. This extended architecture was referred to as the Quantum-inspired SNN (QisSNN) framework. Applied to carefully designed benchmark data, containing irrelevant and redundant features of varying information quality, the QisSNN-based feature selection led to excellent classification results and an accurate detection of relevant information in the dataset.

The QisSNN framework was used on a case study of ecological modeling (Schliebs, Defoin-Platel, Worner, & Kasabov, in press). Meteorological data, such as monthly and seasonal temperature, rain fall and soil moisture recordings for different geographical sites, were compiled from published results, and each global site was **labeled** according to the presence or absence of the

* Corresponding author.

E-mail addresses: sschlieb@aut.ac.nz, sschliebs@gmail.com (S. Schliebs), michael.defoinplatel@gmail.com (M. Defoin-Platel), worner@lincoln.ac.nz (S. Worner), nkasabov@aut.ac.nz (N. Kasabov).

1 Mediterranean fruit-fly (a serious fruit pest). The study aimed
 2 towards the identification of important features relevant for
 3 predicting the presence/absence of this insect species. Results have
 4 been compared to the classical Naïve Bayesian Classifier (NBC) and
 5 obtained feature subsets were verified by an ecological expert.

6 In this study we want to extend the work presented in Schliebs
 7 et al. (in press) by studying the QiSNN framework on two
 8 benchmark and one real world problem. We will start our analysis
 9 by introducing a novel combined optimization algorithm, which
 10 allows us to explore heterogeneous search spaces simultaneously:
 11 The method uses a binary representation for optimizing feature
 12 subsets and a continuous representation for evolving appropriate
 13 real-valued configurations of a spiking network. Altogether four
 14 methods are here experimentally compared to each other:
 15 QiSNN as presented in Schliebs et al. (in press), the enhanced
 16 QiSNN using the combined optimization algorithm, a multi-layer
 17 perceptron and a classical NBC. A comprehensive analysis of the
 18 results obtained from the benchmarks experiments in terms of
 19 consistency of selected feature subsets, classification accuracy,
 20 computational complexity and evolution of parameters in the
 21 QiSNN framework is presented. Furthermore, we will point out
 22 some significant differences between QiSNN and its enhanced
 23 version. Finally we revisit the ecological dataset used in Schliebs
 24 et al. (in press) and new results are obtained and analyzed by an
 25 ecological expert.

26 In the following sections first the QiSNN framework is sum-
 27 marized. The novel continuous optimization method is introduced
 28 and the simultaneous exploration of a binary and a continuous
 29 search space is discussed. QiSNN is then experimentally studied,
 30 followed by an analysis and discussion of the obtained results.

31 **2. Framework and implementation of QiSNN**

32 Based on our previous results on eSNN and quantum inspired
 33 evolutionary algorithms (Benuskova & Kasabov, 2007; Defoin-
 34 Platel et al., in press; Kasabov, 2007; Wysoski et al., 2008), here
 35 we propose and explore an integrative quantum inspired feature
 36 selection using the eSNN architecture, tightly coupled with the
 37 learning environment (the data).

38 **2.1. The eSNN architecture**

39 The eSNN architecture uses a computationally very simple and
 40 efficient spiking neural model, in which early spikes, received by
 41 a neuron, are more strongly weighted than later ones. The model
 42 was inspired by the neural processing of the human eye, which
 43 performs a very fast image processing. Experiments have shown
 44 that a primate only needs several hundreds of milliseconds to
 45 make reliable decisions about images that were presented in a test
 46 scenario (VanRullen & Thorpe, 2001). Since it is known that neural
 47 image recognition involves several succeeding layers of neurons,
 48 these experiments suggested that only very few spikes could
 49 be involved in the neural chain of image processing. In Thorpe
 50 (1997) a mathematical definition of these neurons was proposed
 51 and tested on some face recognition tasks, reporting encouraging
 52 experimental results. The same model was later extended to eSNN
 53 and used in Wysoski et al. (2006) and Wysoski (2008) to perform
 54 audio-visual face recognition.

55 Similar to other SNN approaches, a specific neural model, a
 56 learning method, a network architecture and an encoding from
 57 real values into spike trains need to be defined in an eSNN model.
 58 The neural model is given by the dynamics of the post-synaptic
 59 potential $u_i(t)$ of a neuron i :

$$u_i(t) = \begin{cases} 0 & \text{if fired} \\ \sum_{j|f(j)<t} w_{ji} m_i^{order(j)} & \text{else} \end{cases} \quad (1)$$




Fig. 1. Evolution of the post-synaptic potential (PSP) of a neural model used in QiSNN for a given input stimulus. If the potential reaches threshold θ a spike is triggered and the PSP set to 0 for the rest of the simulation, even if the neuron is still stimulated by incoming spike trains.

Algorithm 1 Training an Evolving Spiking Neural Network

```

Require:  $m_l \in (0, 1)$ ,  $s_l \in (0, 1)$ ,  $c_l \in (0, 1)$ ,  $l \in L$ 
1: initialize neuron repository  $R_l = \{\}$ 
2: for all samples  $X^{(i)}$  belonging to class  $l$  do
3:    $w_j^{(i)} \leftarrow (m_l)^{order(j)}$ ,
    $\forall j | j$  pre-synaptic neuron of  $i$ 
4:    $u_{\max}^{(i)} \leftarrow \sum_j w_j^{(i)} (m_l)^{order(j)}$ 
5:    $\theta^{(i)} \leftarrow c_l u_{\max}^{(i)}$ 
6:   if  $\min(d(w^{(i)}, w^{(n)})) > s_l$ ,  $w^{(n)} \in R_l$  then
7:      $w^{(n)} \leftarrow \text{merge } w^{(i)} \text{ and } w^{(n)}$ 
8:      $\theta^{(n)} \leftarrow \text{merge } \theta^{(i)} \text{ and } \theta^{(n)}$ 
9:   else
10:     $R_l \leftarrow R_l \cup \{w^{(i)}\}$ 
11:   end if
12: end for
    
```

where w_{ji} is the weight of a pre-synaptic neuron j , $f(j)$ the firing time of j , and $m_i \in (0, 1)$ a parameter of the model, namely the modulation factor. Function $order(j)$ represents the rank of the spike emitted by neuron j . For example, a rank $order(j) = 0$ would be assigned, if neuron j is the first among all pre-synaptic neurons that emits a spike. In a similar fashion the spikes of all pre-synaptic neurons are ranked and then used in the computation of u_i . A neuron i fires a spike when its potential has reached a certain threshold θ . After emitting a spike, the potential is reset to $u_i = 0$. Each neuron is allowed to emit only a single spike at most. The threshold $\theta = c u_{\max}$ is set to a fraction $c \in (0, 1)$ of the maximum potential u_{\max} possible by a neuron. In Fig. 1 the change of the post-synaptic potential for this neural model is presented, when a series of input spikes (stimuli) are presented to the different synapses of this neuron.

An evolving neural network architecture using the above model along with a learning algorithm was proposed in Wysoski et al. (2006, 2008). The method successively creates a repository of trained output neurons during the presentation of training samples. For each training sample a new neuron is trained and then compared to the ones already stored in the repository. If a trained neuron is considered to be too similar (in terms of its weight vector) to the ones in the repository (according to a specified similarity threshold s), the neuron will be merged with the most similar one. Otherwise, the trained neuron is added to the repository as a new output neuron. The merging is implemented as the (running) average of the connection weights, and the (running) average of the two firing threshold. Because of the incremental evolution of output neurons it is possible to accumulate knowledge as it becomes available. Hence a trained network is able to learn new data without the need of re-training already learned samples. The procedure is described in detail in Algorithm 1.

The encoding of input values seems to be a critical factor in all SNN approaches. Several encoding mechanisms for SNN have been proposed, such as frequency mappings, Poisson processes

and rank order encoding. Another approach is the population encoding which distributes a single input value to multiple neurons and hence may cause the excitation and firing of several responding neurons. Our implementation is based on arrays of receptive fields as described in Bohte et al. (2002). Receptive fields allow the encoding of continuous values by using a collection of neurons with overlapping sensitivity profiles. Each input variable is encoded independently by a group of M one dimensional receptive fields. For a variable n an interval $[I_{min}^n, I_{max}^n]$ is defined. The Gaussian receptive field of neuron i is given by its center $\mu_i = I_{min}^n + (2i - 3)/2 * (I_{max}^n - I_{min}^n)/(M - 2)$ and width $\sigma = 1/\beta(I_{max}^n - I_{min}^n)/(M - 2)$, with $1 \leq \beta \leq 2$. Parameter β directly controls the width of each Gaussian receptive field.

2.2. Wrapper approach

The eSNN may be used to address feature subset selection (FSS) problems following the well known wrapper approach. A wrapper contains a general optimization algorithm interacting with an induction method (classifier), also cf. Fig. 2. The optimization task consists in a proper identification of an optimal feature subset, which maximizes the classification accuracy determined by the inductor. An eSNN operates here as the induction method, while, due to its interesting properties in terms of solution quality and convergence speed, the previously proposed Versatile Quantum-inspired Evolutionary Algorithm (vQEA) (Defoin-Platel et al., 2007) was used as the optimization algorithm. A vQEA evolves in parallel a number of independent probability vectors, which interact at certain intervals with each other, forming a multi-model Estimation of Distribution Algorithm (EDA) (Defoin-Platel et al., in press). The binary nature of vQEA fits well to the feature selection problem we want to apply it on.

2.3. Integrated feature and parameter optimization

Manual fine-tuning the neuronal parameters can quickly become a challenging task (Wysoski, 2008). An alternative proposed in Valko, Marques, and Castelani (2005), is to optimize both the set of features and the neuronal parameters, in a simultaneous way. The selection of the fitness function was identified to be a crucial step for the successful application of such an embedded approach. In the early phase of the optimization the parameters are selected randomly. As a result it is very likely that a setting is selected for which the classifier is unable to respond to any input presented, which corresponds to flat areas in the fitness landscape. Hence a configuration that will allow the network to fire (even if not correctly) represents a huge (local) attractor in the search space, which could be difficult to escape in later iterations of the search. In Valko et al. (2005) a linear combination of several sub-criteria was used to avoid a too rugged fitness landscape. Nevertheless we can not confirm that the use of much simpler fitness functions led to any problems in our experiments. Using the classification accuracy on testing samples seemed to work well as it is presented in this and previous papers. All parameters, namely modulation factor m_i , similarity threshold s_i , potential fraction c_i , $\forall i \in L$ of eSNN were included in the search space of the optimization method.

Due to its binary nature, vQEA required the conversion of bit strings into real values. It was experimentally shown that a small number of Gray-coded bits seemed sufficient to approximate meaningful parameter configurations of the eSNN method. Nevertheless the use of a binary optimizer for a real-valued search space appears unsatisfactory. Each real-valued parameter needs to be encoded by a number of bits. For the mapping of bit strings into a real value additional computational resources are necessary. Furthermore a granularity is introduced into the search interval.

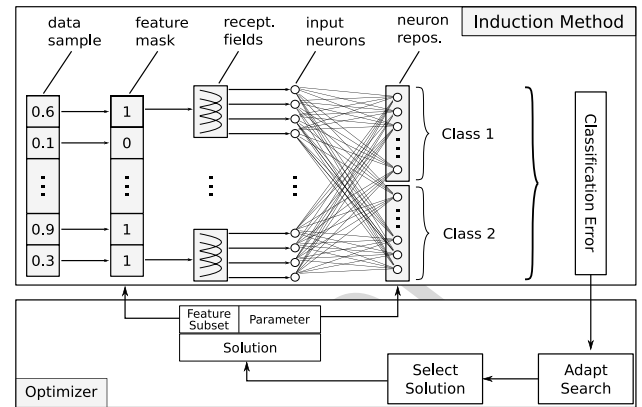


Fig. 2. The QiSNN framework of tightly coupled feature selection and parameter optimization of eSNN, integrated with the data.

Since a single continuous variable is represented by many bits, a binary optimization method has to operate on more variables, compared to a continuous optimizer. Thus scaling problems can be expected for binary optimization, especially in the context of high-dimensional problems, that need a precise optimization of real-valued search variables. Furthermore neighboring solutions in the continuous domain might not be neighbors in their binary representation. Exploring the local neighborhood of a solution may require the optimizer to flip many bits at the same time, which will encourage premature convergence and promote the phenomenon of hitch-hiking. Hence we will extend vQEA in this study towards continuous search spaces and use a combined representation for the simultaneous exploration of a binary landscape and the continuous landscape.

The complete QiSNN framework used in this study is summarized in Fig. 2.

3. Extending vQEA for continuous optimization

Based on the vQEA (Defoin-Platel et al., 2007), we propose here an extension of the algorithm to allow exploration of continuous search spaces. vQEA has been developed as a binary optimization method, which employs some quantum computing principles to enhance classical evolutionary algorithms. The method was deeply studied in Defoin-Platel et al. (in press) and it was shown that vQEA belongs to the class of Estimation of Distribution Algorithms (EDA). The study revealed that the core of vQEA maintains a multiple probabilistic model, which is quite in contrast to other typical EDA, like e.g. Probabilistic Incremental Learning (PBIL), Univariate Marginal Distribution Algorithm (UMDA), and compact Genetic Algorithm (cGA), see Lozano, Larra naga, Inza, and Bengoetxea (2006) for an excellent overview for these algorithms. In these methods only a single probabilistic model is evolved during the optimization process. In vQEA each model explores the search space independently, but it may exchange information at pre-defined intervals with the other models. The algorithm is population based and each individual manages its own probabilistic model. The individuals itself are organized in groups, hence introducing an important structure into the population.

Several advantages of this multi-model and the structured population have been identified. vQEA uses an implicit adaptive learning rate, which makes it robust to its parameter configuration. It was shown that a certain parameter setting is suitable for a variety of problem classes and sizes. Furthermore, the multi-model approach allows a finite number of decision errors, which makes vQEA robust against fitness noise. It was demonstrated that vQEA performs better in terms of speed and solution quality than other first-level EDA, especially when links are introduced between

variables (epistasis). The method was compared to a number of Evolutionary Algorithms on several different benchmark problems. Finally, using several probabilistic models allows a more diverse exploration of the search space than just using a single one.

In the following we want to extend the binary multi-model EDA (i.e. vQEA) towards the area of continuous search spaces. Since all key characteristics of vQEA will be still present in the proposed algorithm, we expect similar advantages of this method in comparison to other evolutionary methods, such as Genetic Algorithm (GA), Particle Swarm Optimization (PSO) and Differential Evolution (D/E).

3.1. Continuous-value multi-model EDA

The probabilistic model in vQEA is based on a Bernoulli random variable for each bit, which is referred to as a qbit according to the used quantum computing metaphor. Sampling from such a string of qbits results in the creation of a bit string, which in turn can be evaluated by the corresponding fitness function. Since we want to consider continuous search spaces now, we have to replace the Bernoulli distribution by a continuous one, such that it becomes possible to sample real values instead of discrete ones. A number of approaches concerning how to employ such distributions and how to model them have been studied in literature. Generally they are based on Gaussian distributions (Bosman & Thierens, 2000; Gallagher & Frean, 2005; Gallagher, Frean, & Downs, 1999; Mininno, Cupertino, & Naso, 2008; Yuan & Gallagher, 2003), histograms (Yuan & Gallagher, 2003), or interval representations (Servet, Travé-Massuyès, & Stern, 1998).

We consider a continuous EDA based on Gaussian distributions here. For each dimension j of the continuous search space and for each probabilistic model i , a random variable following a Gaussian distribution is evolved. Therefore the distribution is fully described by two parameters: The mean $\mu_i^{(j)}$ and the standard deviation $\sigma_i^{(j)}$. In each generation samples are drawn forming real-valued vectors, whose quality can be evaluated by the fitness measure. An update rule is then applied to update $\mu_i^{(j)}$ and $\sigma_i^{(j)}$ to move the search towards promising areas in the search space, making higher quality solutions more likely to be sampled in the next generation. We will first describe the basic structure of algorithm in detail, followed by the presentation of the chosen update rule.

The overall structure of the proposed extension is almost identical to vQEA. Like vQEA also the continuous version is a population-based search method. Its behavior can be decomposed in three different interacting levels: Individual, group and population level.

Individuals. The lowest level corresponds to *individuals*. An individual i at generation t contains a probabilistic model $P_i(t)$ and two real-valued strings $C_i(t)$ and $A_i(t)$. More precisely P_i corresponds to a string of N pairs of values $(\mu_i^{(j)}, \sigma_i^{(j)})$:

$$P_i = P_i^1 \dots P_i^N = \begin{bmatrix} \mu_i^{(1)} & \dots & \mu_i^{(N)} \\ \sigma_i^{(1)} & \dots & \sigma_i^{(N)} \end{bmatrix}. \quad (2)$$

The pair $(\mu_i^{(j)}, \sigma_i^{(j)})$ corresponds to the parameters of the distribution of the j th variable of the i th probabilistic model. Each variable in P_i is sampled according to $\mu_i^{(j)}$ and $\sigma_i^{(j)}$, so that C_i represents a configuration in the search space whose quality can be determined using a fitness function f . In most continuous optimization problems, the variables have a specific domain of definition. Without loss of generality we assume each $c_i^{(j)} \in C_i$ to be defined in the interval $[-1, 1]$. As a consequence, each $c_i^{(j)} \in C_i$ follows a *truncated normal distribution* in the range $[-1, 1]$. Truncated normals can be sampled using a simple

numerical procedure and the technique is widely adopted in pseudo-random number generation, see e.g. Geweke (1991) for an efficient implementation.

To each individual i a solution A_i is attached acting as an attractor for P_i . Every generation C_i and A_i are compared in terms of their fitness. If A_i is better than C_i (i.e. $f(A_i) > f(C_i)$ assuming a maximization problem), an update operation is applied on the corresponding model P_i . The update will move the mean values of the probabilistic model P_i slightly towards the attractor A_i . The choice of a suitable model update operation is critical for the working of the algorithm. We will elaborate the details of the model update in Section 3.1.1.

The update policy of an attractor A_i can follow two distinctive strategies. In the original QEA (Han & Kim, 2002) an *elitist* update strategy was used, in which the attractor A_i is replaced by C_i only if C_i is better than A_i in terms of fitness. In a *non-elitist* update strategy (firstly introduced in Defoin-Platel et al. (2007)) C_i replaces A_i at every generation. The choice of the update policy has great consequences for the algorithm and changes its behavior completely. To emphasize the importance of the update rule the non-elitist version of QEA has been proposed as *Versatile QEA* (vQEA) as the attractors are able to change every generation and therefore demonstrate a very high volatility. Since no experimental condition could be identified that favored the elitist attractor update policy, we will concentrate on the non-elitist version during the course of this paper.

Groups. The second level corresponds to *groups*. The population is divided into g groups each containing k individuals having the ability of synchronizing their attractors. For that purpose, the best attractor (in terms of fitness) of a group, noted B_{group} , is stored at every generation and is periodically distributed to the group attractors. A parameter S_{local} is introduced, which controls the phase of local synchronization, i.e. a local synchronization event is triggered in every S_{local} -th generation.

Population. The set of all $p = g \times k$ individuals forms the *population* and defines the topmost level of the multi-model approach. As for the groups, the individuals of the population can synchronize their attractors, too. For that purpose, the best attractor (in terms of fitness) among all groups, noted B_{global} , is stored every generation and is periodically distributed to the group attractors. A parameter S_{global} is introduced, which controls the phase of global synchronization, i.e. a global synchronization event is triggered in every S_{global} -th generation.

3.1.1. Model update

The update of the probabilistic model is particularly interesting, since it governs how the search space is explored by the algorithm. Several continuous EDA have been proposed in literature (Bosman & Thierens, 2000; Mininno et al., 2008; Mühlenbein, Mahnig, & Rodriguez, 1999; Sebag & Ducoulombier, 1998; Yuan & Gallagher, 2003), along with a number of different update rules, e.g. Gallagher and Frean (2005) and Yuan and Gallagher (2003). The common principle of all these continuous EDA is based on the sampling of a population. In vQEA (and thus also its extension) the situation is very different, since only a *single* solution (for each probabilistic model) is sampled in every iteration. Hence the model update can not rely on the density of a population, but has to use a single attractor instead to perform the desired update.

We formulate here an appropriate update rule for the probabilistic models. Updating the mean $\mu^{(j)}$ in the Gaussian variable j appears to be straight-forward. We adopt a mean shift towards the value of the current attractor $a^{(j)}$ at location j , which is quite similar to the mean update used in methods mentioned

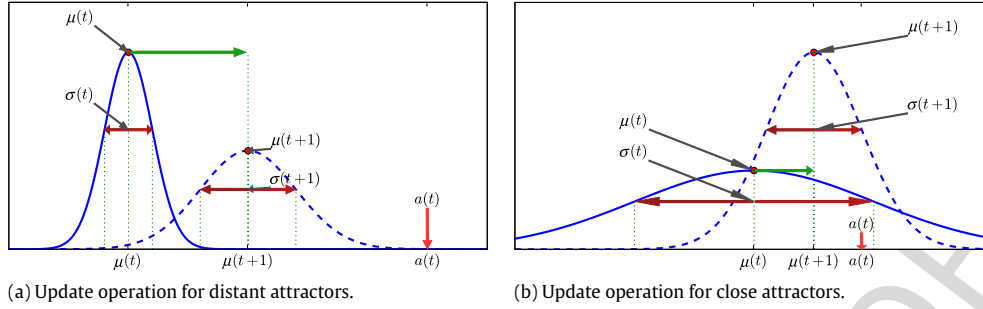


Fig. 3. The figure presents the update operation for a single Gaussian random variable. For each update the distance $d = a(t) - \mu(t)$ between the attractor $a(t)$ and the mean $\mu(t)$ of the Gaussian is computed at generation t . (a) If $d \geq \sigma(t)$ the attractor is considered distant. We interpret that situation by assuming that $\mu(t)$ does not represent a promising area in the search space. In this case the mean $\mu(t)$ is strongly shifted towards the attractor, while at the same time the standard deviation $\sigma(t)$ is increased to allow a wider search in the fitness landscape. (b) On the other hand, if the attractor is inside the boundaries defined by $\sigma(t)$, i.e. $d < \sigma(t)$, we assumed that $\mu(t)$ is already in a promising area of the search space. The algorithm starts to localize the search by shifting $\mu(t)$ only slightly towards the direction of the attractor, while decreasing $\sigma(t)$ at the same time.

above. Depending on the distance $d = a^{(j)} - \mu^{(j)}$ a shift $\Delta\mu^{(j)}$ is computed at generation t :

$$\Delta\mu^{(j)}(t) = \frac{2}{1 + e^{-5d}} - 1 \quad (3)$$

which is then used to perform the update:

$$\mu^{(j)}(t+1) = \mu^{(j)}(t) + \theta_\mu \Delta\mu^{(j)}(t). \quad (4)$$

In Eq. (4) a parameter θ_μ is introduced, which we will refer to as the learning rate of the mean. We note that θ_μ corresponds to the maximum mean shift in a single generation.

For the update of the standard deviation $\sigma^{(j)}$ we will exploit the idea that $\sigma^{(j)}$ should decrease whenever $\mu^{(j)}$ represents a “promising” area in the fitness landscape. We assume $\mu^{(j)}$ to be “fit” when $|d| < \sigma^{(j)}$. Thus, if the attractor $a^{(j)}$ is close to $\mu^{(j)}$ (within the boundaries defined by $\sigma^{(j)}$), the standard deviation $\sigma^{(j)}$ is decreased. It is noteworthy that solutions fulfilling this condition are more likely to be sampled, than other solutions, which means that on average $\sigma^{(j)}$ will decrease. Attractors that are more distant to $\mu^{(j)}$ and thus $|d| \geq \sigma^{(j)}$, will cause an increase of $\sigma^{(j)}$, since it can be assumed that $\mu^{(j)}$ does not represent a promising area in the landscape. We define the standard deviation shift $\Delta\sigma^{(j)}$ at generation t as:

$$\Delta\sigma^{(j)}(t) = \frac{1}{1 + e^{-10(\sigma^{(j)}(t) - 0.5)}} \quad (5)$$

and then use it to perform the update:

$$\sigma^{(j)}(t+1) = \begin{cases} \sigma^{(j)}(t) - \theta_\sigma \Delta\sigma^{(j)}(t) & \text{if } |d| < \sigma^{(j)} \\ \sigma^{(j)}(t) + \theta_\sigma \Delta\sigma^{(j)}(t) & \text{otherwise.} \end{cases} \quad (6)$$

In Eq. (6) a parameter θ_σ is used, which we will refer to as the learning rate of the standard deviation. Again we want to note that θ_σ corresponds to the maximum standard deviation shift in a single generation. In order to avoid divergent behavior of the algorithm, i.e. $\sigma^{(j)}$ increases indefinitely, we restrict the domain of $\sigma^{(j)}$ by defining upper and lower bounds, such that $\sigma_{\min} \leq \sigma^{(j)} \leq \sigma_{\max}$.

In Fig. 3 the principle of the update rule is summarized. Distant attractors (relative to the current mean of the PDF) result in a large mean shift, while at the same time the standard deviation is increased, cf. Fig. 3(a). For close attractors the mean shift is small and the standard deviation is decreased, cf. Fig. 3(b).

It is important to note, that the probabilistic update operator described above, is similar to the rotation gate used in QEA. As shown in Defoin-Platel et al. (in press) the size of an update step using the rotation gate depends on the convergence of the probabilistic model. This phenomenon was described as a form of deceleration of the algorithm before convergence. The sigmoid shape of the standard deviation update adopts a similar strategy, since also here the size of the shift $\Delta\sigma^{(j)}$ decreases with increasing convergence of the algorithm.

3.2. Combined search spaces

Many real-world problems require the exploration of combined search spaces: a binary and a continuous space. An example is the parallel evolution of the topology and the weight matrix of a neural network. Here the topology is encoded as a bit string, where “1” represents a present connection between two neurons and “0” encodes its absence. Another example is the wrapper based feature selection, where the presence/absence of a feature requires a binary search space, while appropriate configurations for the classification method may correspond to a continuous landscape.

It is now possible to employ vQEA on combined search spaces with two types of representation. Each representation uses its corresponding update operator to drive the probabilistic model towards promising areas in the search space. In every generation the models are sampled and then evaluated by a single fitness measure. The fitness evaluation uses the sampled binary and continuous solution part to determine the quality of the combined solution. According to the fitness of the obtained solution the models are updated. This extended vQEA allows us to enhance the original QisNN.

We emphasize that the extended vQEA is similar to a collaborative coevolutionary algorithm (Potter & Jong, 2000). The evolution of the two representations proceeds more or less independently from each other. Both use their own solution representations and update operators and may explore their search space with different learning rates. Despite their independent evolution, both representations share a single fitness function. The binary and continuous sub-solutions are the components of a combined solution, and both parts need to collaborate in order to maximize their fitness.

4. Experiments

We study the enhanced version of QisNN on two benchmark problems. The first benchmark is referred to as the two-spiral problem, on which the original QisNN was investigated before, cf. Schliebs et al. (2009). This problem is composed of two-dimensional data forming two intertwined spirals and was firstly introduced in Lang and Witbrock (1988). It requires learning of a highly non-linear separation of the input space. The data was frequently used as a benchmark for neural networks, including the analysis of the eSNN method itself (Schliebs et al., 2009; Wysocki, 2008). Since the data contains only two relevant dimensions it was extended by adding redundant and random information. The importance of the redundant features was varied: Features range from mere copies of the original two spirals to completely random ones. The inherent information of a feature decreases when stronger noise is applied. A detailed description of the data

generation can be found in Schliebs et al. (2009). The dataset contains seven redundant two-dimensional spiral points x'_i, y'_i , for each point a different noise strength is used, totalling in 14 redundant features. An additional four random features r_1, \dots, r_4 were included. Together with the two relevant features of the spirals (x and y) the dataset contained 20 features in 400 samples.

The second benchmark is the uniform hypercube dataset, to our best knowledge firstly introduced in Estevez, Tesmer, Perez, and Zurada (2009). The problem consists of two classes of 400 samples. For each sample a five-dimensional vector (r_1, \dots, r_5) is drawn from a uniform distribution. A given pattern belongs to class 1 if $r_i < \gamma^{i-1} * \alpha$ for $i = 1, \dots, 5$ and to class 2 otherwise. The parameters were chosen to be $\gamma = 0.8$ and $\alpha = 0.5$. The entire dataset consists of 40 features, five relevant, 30 random and five redundant ones. The latter are linear combinations of the relevant features perturbed by additive Gaussian noise of increasing level. The dataset was balanced. A more detailed explanation of the data generation can be found in Estevez et al. (2009).

4.1. Setup

For the combined optimization method we chose a population structure of ten individuals organized in a single group, which is globally synchronized every generation. This setting was reported to be generally superior for a number of different benchmark problems (Defoin-Platel et al., in press). In the case of the spiral dataset the learning rate for the binary rotation gate was set to $\theta = \pi/50$. For the rate of the mean and standard deviation shift we chose $\theta_\mu = 0.1$ and $\theta_\sigma = 0.1$ respectively. The algorithm was allowed to evolve over a total number of 400 generations. Due to its larger problem size 1000 generations were computed for the hypercube problem, using once more $\theta = \pi/50$ for the binary learning rate and $\theta_\mu = 0.1$ and $\theta_\sigma = 0.05$ for the continuous update operator.

In order to allow a fair comparison between the classification methods used in this study, one has to decide for an appropriate parameter configurations for each classifier. In contrast to the other classifiers, NBC does not require the tuning of any parameters. To setup the MLP we have experimented on a subset of the datasets containing the relevant features only. By changing the number of hidden neurons, the learning rate, and the momentum term a satisfying configuration, in terms of classification accuracy, was experimentally obtained by systematic trial and error. The results of the parameter study for the MLP on the spiral dataset are presented in Fig. 4. The finally chosen setting is based on a tradeoff between computational cost and classification accuracy. The additional cost of more hidden neurons is not worth the slight increase of accuracy reported in Fig. 4. Using 10-fold cross-validation the chosen configuration of MLP achieved a satisfying accuracy of 0.849 (standard deviation 0.0634) on the spiral dataset containing the two relevant features only. When applied to the full dataset using all 20 features, the same configuration resulted in an accuracy of 0.611 (0.0608). Thus, appropriate feature selection does improve the performance of MLP, which is the key principle exploited in the wrapper approach. Finding an appropriate setting for the spiral problem appeared to be more difficult, in contrast to the other benchmark. For the latter problem changes in the configuration did not seem to impact the performance of the classifier too much. Thus we decided to use the same parameter setting for both problems. The common error back-propagation learning algorithm was used to train the network, connection weights were initialized with small uniform random numbers in the range $[-0.25, 0.25]$.

Most of the parameters of QiSNN are optimized during the evolutionary process. For each class $l \in L$ three parameters exist: The modulation factor m_l , the similarity threshold s_l , and

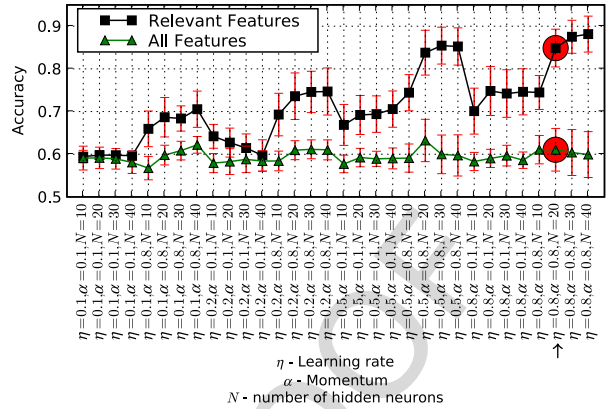


Fig. 4. The figure shows the accuracy levels achieved by 32 different configurations of a multi-layer perceptron on the two-spiral dataset. Each point represents the average of the accuracies obtained in a 10-fold cross-validation experiment, error bars indicate the standard deviation. All configurations use neurons with sigmoid transfer functions, trained in 500 epochs. The lower curve (green triangles) represents the accuracy of the MLP when all 20 features are included in the dataset, the upper curve (black squares) the accuracy when only the relevant features are used. The circles (red) indicate the finally chosen configuration for the experiments performed in this study, which is a satisfying compromise between computational cost and classification quality.

the proportion factor c_l . Since both problems contain two classes, six parameters are involved in the QiSNN framework used here. In terms of the population encoding we found especially that the number of receptive fields needs careful consideration, since it affects the resolution for distinguishing between different input variables. After some preliminary experiments we decided for 20 receptive fields in case of the spiral data and five receptive fields for the hypercube. The Gaussian centers were uniformly distributed over the search interval and the variance was set to $\beta = 1.5$.

In order to guarantee statistical relevance, 30 independent runs for each investigated classification method were performed. In every generation all samples of the dataset were randomly shuffled and divided into training and testing samples, according to a ratio of 0.75. For the computation of the classification error we determined the ratio between correctly classified samples and the total number of testing samples.

4.2. Results

We discuss the results on the two-spiral problem first, cf. Fig. 5. Fig. 5(a)–(d) present the evolution of the average best feature subset in every generation using the enhanced and original QiSNN, MLP and NBC respectively. The color of a point in these diagrams reflects how often a specific feature was selected at a certain generation: The lighter the color the more often the corresponding feature was selected. It can clearly be seen that independent of the used algorithm a large number of features has been discarded during the evolutionary process. Furthermore all algorithms clearly identify the features x and y to be relevant. All methods except the enhanced QiSNN select some redundant and/or irrelevant features, too.

Particularly interesting is the order in which the features have been discarded by each algorithm. Both QiSNN (Fig. 5(a) and (d)) rejected the four random features r_1, \dots, r_4 containing no information almost immediately in less than 20 generations. The redundant features x'_i, y'_i were then rejected one after the other, according to the strength of the inherent noise: The higher the noise, the earlier a feature is identified as irrelevant. We note the improved performance of the enhanced QiSNN, which is clearly able to reject all redundant features in most of the runs. Fig. 5(e) compares the evolution of the number of selected features during

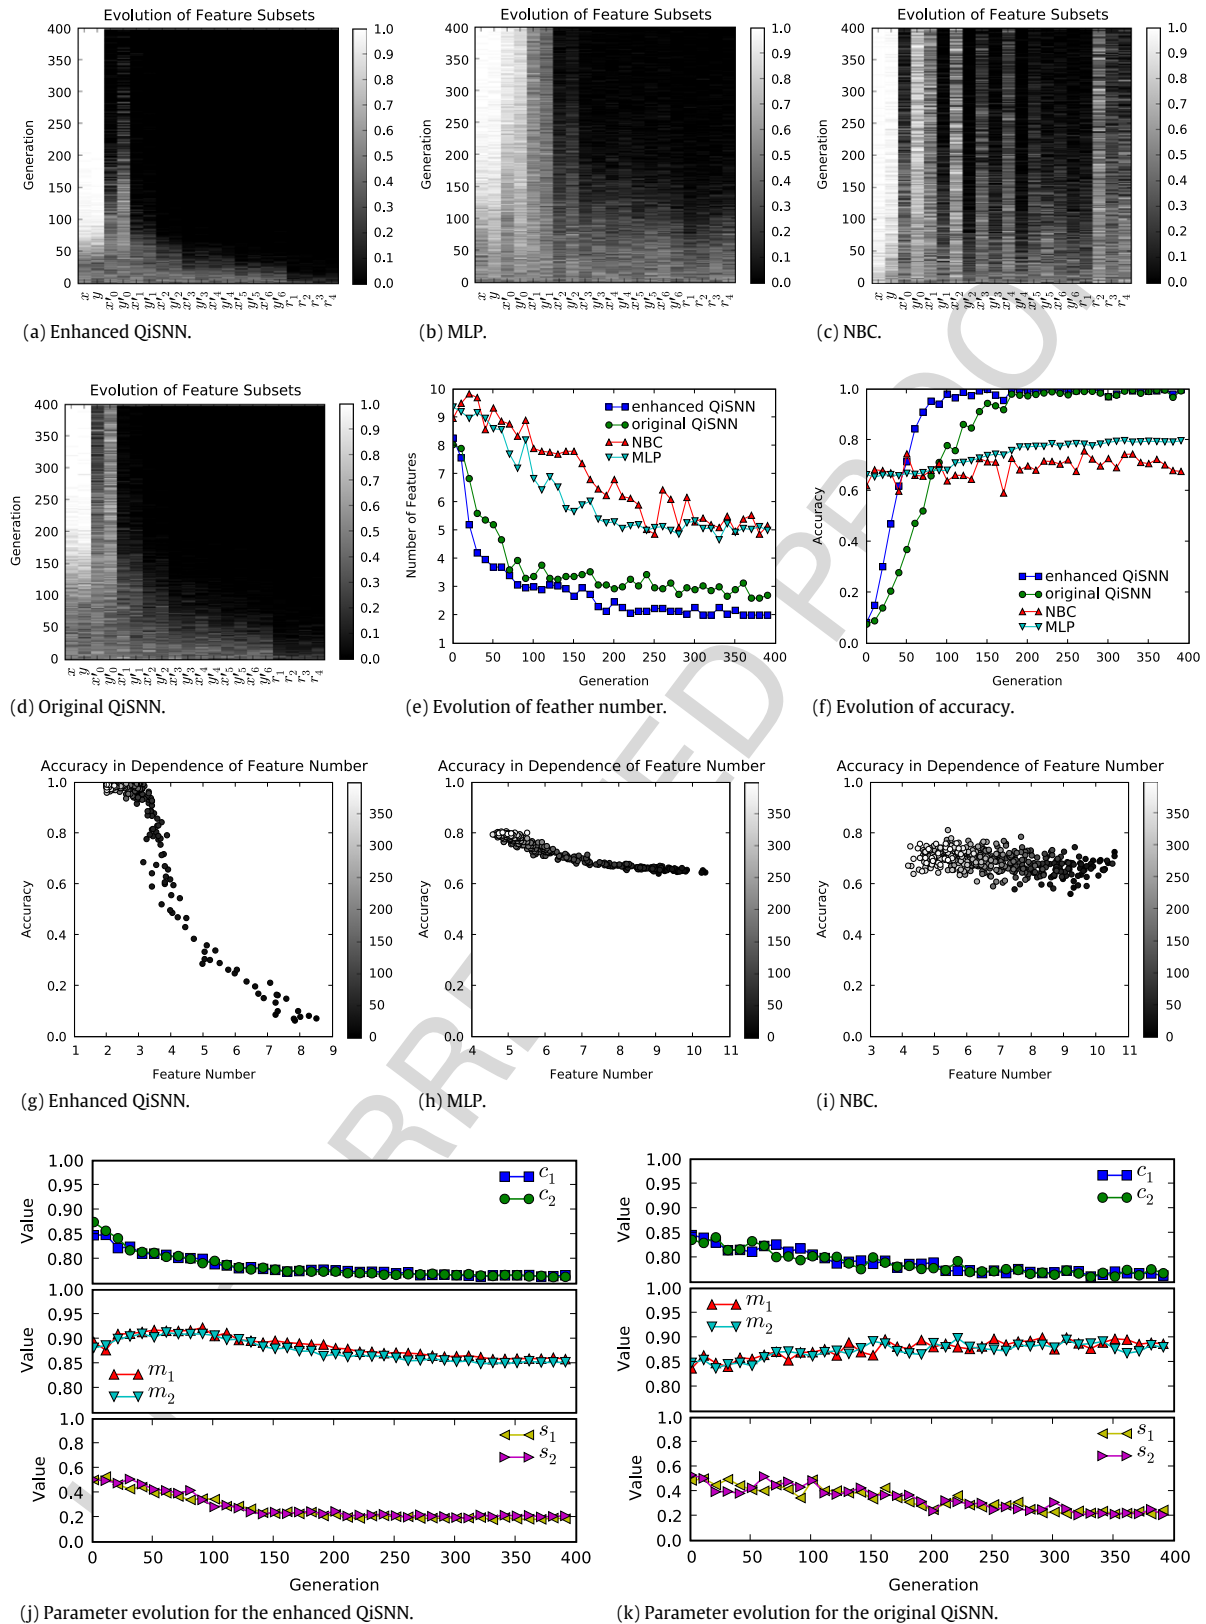


Fig. 5. Results on the two-spiral dataset (see the text for explanation).

1 each generation. While both QiSNN based methods clearly select
 2 less features than their classical competitors at any stage of the
 3 optimization, the original QiSNN is outperformed by the proposed
 4 enhanced version.

It is also interesting to compare the evolution of the classifica-
 tion error for each algorithm, cf. Fig. 5(f). The gradient in the fitness
 landscape defined by eSNN appears to be much steeper compared
 to all other algorithms, ranging from completely unfit solutions at

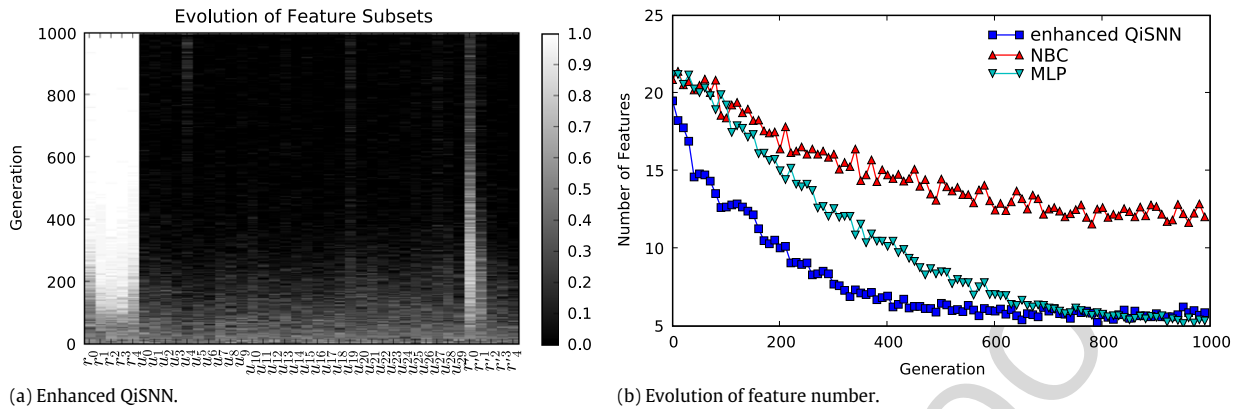


Fig. 6. Results on the hypercube dataset (see the text for explanation).

the beginning of the evolutionary run, towards high quality solutions in later generations. MLP and NBC display a more flat fitness evolution. We believe that fitness gradient is partially responsible for turning eSNN in a very good quality measure for feature subsets.

According to the presented results on the QiSNN, a strong correlation between classification accuracy and feature number appears advantageous in the context of a feature selection task. Fig. 5(g)–(i) present this dependence for the investigated induction methods. Each point in the diagram corresponds to a tuple (accuracy, feature number) obtained from the generational best individual of every generation. The color indicates the generation itself, the lighter the color the later the generation in which a given tuple was obtained. In the case of QiSNN (cf. Fig. 5(g)) a strong relationship between feature number and accuracy can be observed.¹ Even for small decreases of the feature number significant accuracy improvements are reported. The strong correlation between feature number and classification accuracy introduces a gradient and partially reduces *neutrality* in the fitness landscape. Removing a redundant or irrelevant feature from the selected subset corresponds to a fitness gain for QiSNN, which may not necessarily be true for the other two tested methods. If the feature removal does not lead to a certain fitness gain and thus two solutions may have the same fitness value, the fitness landscape has a neutral dimension at the corresponding parameter. Due to genetic drift the neutral parameter will converge randomly, which means a random selection or non-selection of the encoded feature. In the fitness landscape defined by eSNN neutral dimensions are replaced by a fitness gradient, which allows the identification and exclusion of low quality features from the current subset. As a result the fitness landscape can be easily climbed by the optimization algorithm, leading to faster and more consistent convergence towards the optimal feature subset.

Fig. 5(j) and (k) present the evolution of the eSNN parameters for the two versions of QiSNN. Although both methods have evolved similar final parameter configurations, the continuous exploration is much smoother compared to the binary one and allowed a finer parameter tuning. Due to the balanced nature of the dataset the parameter setting for the two classes have evolved to be approximately identical, i.e. $c_1 \approx c_2$, $m_1 \approx m_2$ and $s_1 \approx s_2$.

A similar analysis was done for the second benchmark dataset. We want to note that this dataset was very easy to solve by any of the tested algorithms. Even without feature selection MLP and NBC reported a very high classification accuracy. Nevertheless

we have decided to present these results here, since they show the proper functioning of all tested methods on an additional independent benchmark problem. Fig. 6 summarizes the results on the hypercube problem. In Fig. 6(a) the evolution of the average selected feature subset is shown. We have included the diagram for the enhanced QiSNN only, since the other methods report very similar results. Thus QiSNN was chosen as a representative of all tested algorithms. Similar to the figures presented on the spiral data above, different gray levels reflect how often a specific feature was selected at a certain generation. In this diagram the first five features correspond to the relevant features, followed by 30 irrelevant and finally five redundant features. All methods clearly identify the five relevant variables. QiSNN was capable of decreasing the feature number faster than NBC and MLP, cf. Fig. 6(b). NBC reports optimal classification accuracy without removing all irrelevant and redundant features. Without the presence of any selective pressure some features converge randomly due to genetic drift, which has resulted in the selection of some irrelevant features in case of NBC.

We also want to discuss the computational complexity for each of the presented algorithms here. The fitness evaluation of a feature subset is clearly the most costly part in the wrapper. Depending on the dataset an MLP requires the construction of a rather large neural network, followed by the training of each data sample for 500 epochs using a costly back-propagation procedure and is thus by far the most complex method in this study. The eSNN classifier implements a fast one-pass learning, but additional overhead is required for transforming each data sample into a spike sequence and computing the spike propagation in the network. Due to the simple topology of the network an efficient spike simulation is possible. The NBC requires the lowest computational resources, each training sample is investigated only once and only minimal overhead is necessary, allowing very fast classification.

5. Ecological modeling revisited

In Schliebs et al. (in press) the original QiSNN framework was applied on an ecological modeling problem. Because of the promising results obtained from the benchmark studies before, we wanted to revisit the ecological data using the enhanced QiSNN for feature selection. For many invertebrate species, little is known about their response to environmental variables over large spatial scales. That knowledge is important in order to predict the establishment of a species, that has the potential to cause great environmental harm. The usual approach to determine the importance of a range of environmental variables, that explain the global distribution of a species, is to train or fit a model to its known distribution using environmental parameters measured in areas where the species is present and where it is absent. In this study,

¹ Both versions of QiSNN show a similar behavior here, thus we have chosen to present the enhanced QiSNN as a representative of the two.




Fig. 7. Results on the ecological data set averaged over 30 independent runs. The lighter the color of a point in the diagram, the more often a specific feature was selected at the given generation.

meteorological data that comprised 68 monthly and seasonal temperature, rainfall and soil moisture variables for 206 global geographic sites were compiled from published records (CABI, 2003). These variables were correlated to global locations where the Mediterranean fruit-fly (*Ceratitis capitata*), a serious invasive species and fruit pest, were recorded at the time of the study, as either present or absent. The dataset is balanced having equal number of samples for each of the two classes. Previous use of MLP on the data results in a classification accuracy of approximately 71% (Watts & Worner, 2006).

The experimental setup defined in Schliebs et al. (in press) was kept mostly unchanged here to allow some comparison to previous results: Ten individuals are allowed to evolve in 4000 generations, statistical relevance is guaranteed by performing 30 independent runs and averaging the results. The additional parameters for the mean and standard deviation shift were set to $\theta_\mu = 0.1$ and $\theta_\sigma = 0.01$ respectively, the learning rate of the binary model was $\theta = \pi/100$.

Fig. 7 presents the results of the revisited experiment. Similar to the figures before the evolution of the average best feature subset is shown, where the color reflects how often a specific feature was selected at a certain generation. The comparison between NBC and the original QISNN was discussed in great detail in Schliebs et al. (in press), thus we will concentrate on the discussion of the performance of the two QISNN only. Fig. 8 clearly shows the similarity of the feature subsets obtained by both QISNN. Nevertheless the enhanced version reports greater consistency in the feature rejection. Also the enhanced QISNN selected significantly less features than the original QISNN, cf. Fig. 8: On average 14 features were selected using QISNN, 9 in case of the enhanced QISNN and 18 using NBC. Compared to the original QISNN the enhanced version additionally rejected the following features: temp1, temp3, TAut2, TSpr1, Tannual, rain10, RSumR2, PEAnnual.

The overall classification accuracy was similar between all tested algorithms. From an ecological point of view the evolved feature subsets are coherent with the current knowledge in this area. Winter temperatures, autumn rainfall and the degree-days (DD5 and DD15) were particularly strong features. Degree-days are the accumulated number of degrees of temperature above a threshold

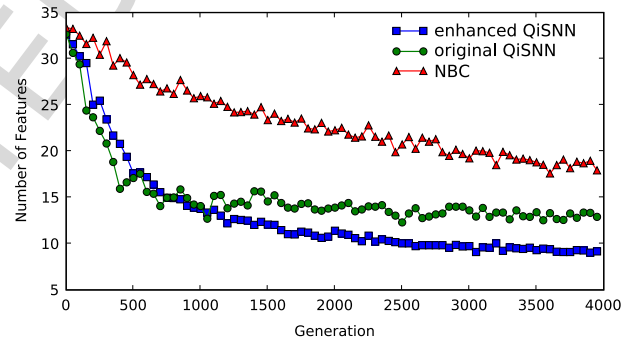


Fig. 8. In all algorithms the number of features decreases with increasing generations, the enhanced QISNN being noticeably faster than the original QISNN and NBC. All classifiers report a similar accuracy after the evolution of 4000 generations.

temperature (5° and 15° in this case) over time (in this dataset over the whole year). It would be expected that the latter two variables would be closely correlated. These results correspond to other analysis where more conventional statistical and machine learning methods were used to identify the contribution of environmental variables to *C. capitata* presence or absence (Worner, Leday, & Ikeda, 2008). While there is no indication from this analysis whether the features have a negative or positive effect on the distribution of the species, it is known that *C. capitata* is limited by the severity of temperatures in the winter and extremes of wet or dry conditions in the summer and autumn (Vera, Rodriguez, Segura, Cladera, & Sutherst, 2002).

6. Conclusions and future directions

In this study we have presented an enhanced QISNN framework by proposing an extension for the used optimization method. The optimizer simultaneously evolves a feature subset along with an parameter configuration for a spiking network by using two separate probabilistic representations: A binary representation for optimizing the feature search space and a continuous one for exploring the parameter space. We have pointed out the similarity to coevolutionary algorithms. The enhanced QISNN was studied on

two benchmark datasets and an ecological data modeling problem was revisited. Results have been compared to the traditional NBC, an MLP and also to the original version of QiSNN. The enhanced QiSNN was shown to be faster, more consistent and reliable than its predecessor. It also compared well to the classical methods tested here. Further development of the QiSNN is planned where the presence/absence of spikes at a given time in a QiSNN will also be represented stochastically as suggested in Kasabov (2008a, 2008b).

References

- Benuskova, L., & Kasabov, N. (2007). *Computational neurogenetic modelling*. NY: Springer.
- Bohte, S. M., Kok, J. N., & Poutré, J. A. L. (2002). Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing*, 48(1–4), 17–37.
- Bosman, P. A., & Thierens, D. (2000). Expanding from discrete to continuous estimation of distribution algorithms: The idea. In *In parallel problem solving from nature - PPSN VI* (pp. 767–776). Springer.
- CABI. (2003). *Crop protection compendium* (5th ed.). Global module.
- Defoin-Platel, M., Schliebs, S., & Kasabov, N. (2007). A versatile quantum-inspired evolutionary algorithm. In *IEEE congress on evolutionary computation, 2007*. (pp. 423–430).
- Defoin-Platel, M., Schliebs, S., & Kasabov, N. (2009). Quantum-inspired evolutionary algorithm: A multimodel eda. *IEEE Transactions on Evolutionary Computation* (-) (in press).
- Estevez, P., Tesmer, M., Perez, C., & Zurada, J. (2009). Normalized mutual information feature selection. *IEEE Transactions on Neural Networks*, 20(2), 189–201.
- Gallagher, M., & Frean, M. (2005). Population-based continuous optimization, probabilistic modelling and mean shift. *Evolutionary Computation*, 13(1), 29–42.
- Gallagher, M., Frean, M., & Downs, T. (1999). Real-valued evolutionary optimization using a flexible probability density estimator. In *Proceedings of the GECCO 1999 Genetic and Evolutionary Computation Conference* (pp. 840–846). Morgan Kaufmann Publishers.
- Gerstner, W., & Kistler, W. M. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge, MA: Cambridge University Press.
- Geweke, J. (1991). Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints and the evaluation of constraint probabilities. In *Computing science and statistics: proceedings of the 23rd symposium on the interface* (pp. 571–578).
- Han, K.-H., & Kim, J.-H. (2002). Quantum-inspired evolutionary algorithm for a class of combinatorial optimization. *IEEE Transactions on Evolutionary Computation*, 6(6), 580–593.
- Iannella, N., & Kindermann, L. (2005). Finding iterative roots with a spiking neural network. *Information Processing Letters*, 95(6), 545–551. (Applications of Spiking Neural Networks).
- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6), 1569–1572.
- Kasabov, N. (2007). *Evolving connectionist systems: The knowledge engineering approach* (Second ed.). Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Kasabov, N. (2008a). Evolving intelligence in humans and machines: Integrative connectionist systems approach (feature article). *IEEE Computational Intelligence Magazine*, 3(3), 23–37.
- Kasabov, N. (2008b). Integrative connectionist learning systems inspired by nature: Current models, future trends and challenges. *Natural Computing*, 8(2), 199–218.
- Knoblauch, A. (2005). Neural associative memory for brain modeling and information retrieval. *Information Processing Letters*, 95(6), 537–544.
- Lang, K. J., & Witbrock, M. J. (1988). Learning to tell two spirals apart. In *Proceedings of the 1988 connectionist models summer school san mateo, Morgan Kaufmann (Ed.)* (pp. 52–59).
- Lozano, J. A., Larra naga, P., Inza, I. N., & Bengoetxea, E. (2006). *Towards a new evolutionary computation: Advances on estimation of distribution algorithms (studies in fuzziness and soft computing)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Maass, W. (1999). *Computing with spiking neurons*.
- Mininno, E., Cupertino, F., & Naso, D. (2008). Real-valued compact genetic algorithms for embedded microcontroller optimization. *IEEE Transactions on Evolutionary Computation*, 12(2), 203–219.
- Mühlenbein, H., Mahnig, T., & Rodriguez, A. O. (1999). Schemata, distributions and graphical models in evolutionary optimization. *Journal of Heuristics*, 5(2), 215–247.
- Potter, M. A., & Jong, K. A. D. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8, 1–29.
- Schliebs, S., Defoin-Platel, M., & Kasabov, N. (2009). Integrated feature and parameter optimization for an evolving spiking neural network. In *LNCS: Vol. 5506/5507. International conference on neural information processing, ICONIP 2008, Springer*.
- Schliebs, S., Defoin-Platel, M., Worner, S., & Kasabov, N. (2009). Quantum-inspired feature and parameter optimization of evolving spiking neural networks with a case study from ecological modeling. In *IEEE international joint conference on Neural Networks (in press)*.
- Sebag, M., & Ducoulombier, A. (1998). Extending population-based incremental learning to continuous search spaces. In *PPSN V: Proceedings of the 5th international conference on parallel problem solving from nature* (pp. 418–427). London, UK: Springer-Verlag.
- Servet, I., Travé-Massuyès, L., & Stern, D. (1998). Telephone network traffic overloading diagnosis and evolutionary computation techniques. In *AE'97: Selected papers from the third European conference on artificial evolution* (pp. 137–144). London, UK: Springer-Verlag.
- Soltic, S., Wysoski, S., & Kasabov, N. (2008). Evolving spiking neural networks for taste recognition. In *IEEE world congress on computational intelligence*, Hong Kong.
- Thorpe, S. J. (1997). How can the human visual system process a natural scene in under 150 ms? Experiments and neural network models. In *ESANN*.
- Valko, M., Marques, N. C., & Castalani, M. (2005). Evolutionary feature selection for spiking neural network pattern classifiers. In B, et al., (Eds.), *Proceedings of 2005 Portuguese conference on artificial intelligence* (pp. 24–32). IEEE.
- VanRullen, R., & Thorpe, S. J. (2001). Is it a bird? is it a plane? ultra-rapid visual categorisation of natural and artificial objects. *Perception*, 30, 655–668.
- Vera, M. T., Rodriguez, R., Segura, D. F., Cladera, J. L., & Sutherst, R. W. (2002). Potential geographical distribution of the Mediterranean fruit fly, *Ceratitis capitata* (diptera: tephritidae) with emphasis on argentina and australia. *Environmental Entomology*, 31(6), 1009–1022.
- Verstraeten, D., Schrauwen, B., & Stroobandt, D. (2005). Isolated word recognition using a liquid state machine. In *ESANN* (pp. 435–440).
- Watts, M., & Worner, S. (2006). Using mlp to determine abiotic factors influencing the establishment of insect pest species. In *International joint conference on Neural Networks* (pp. 1840–1845). Vancouver, Canada: IEEE.
- Worner, S., Leday, G., & Ikeda, T. (2008). Uncertainty analysis and ensemble selection of statistical and machine learning models that predict species distribution. In *Ecological informatics*, Cancun, Mexico.
- Wysoski, S. G. (2008). *Evolving spiking Neural Networks for adaptive audiovisual pattern recognition*. Ph.D. thesis, Auckland University of Technology. <http://hdl.handle.net/10292/390>.
- Wysoski, S. G., Benuskova, L., & Kasabov, N. (2006). On-line learning with structural adaptation in a network of spiking neurons for visual pattern recognition. In *ICANN (1)* (pp. 61–70).
- Wysoski, S. G., Benuskova, L., & Kasabov, N. (2008). Fast and adaptive network of spiking neurons for multi-view visual pattern recognition. *Neurocomputing*, 71(13–15), 2563–2575.
- Yuan, B., & Gallagher, M. (2003). Playing in continuous spaces: Some analysis and extension of population-based incremental learning. In *The 2003 congress on Evolutionary computation: Vol. 1* (pp. 443–450).