

Training HMM Structure with Genetic Algorithm for Biological Sequence Analysis

Kyoung-Jae Won¹, Adam Prügel-Bennett¹, Anders Krogh²

¹ISIS Group, ECS, University of Southampton, SO17 1BJ, United Kingdom and ²Bioinformatics Centre, University of Copenhagen, DK-2100 Copenhagen, Denmark

ABSTRACT

Motivation: Hidden Markov Models (HMMs) are widely used for biological sequence analysis because of their ability to incorporate biological information in their structure. An automatic means of optimising the structure of HMMs would be highly desirable. However, this raises two important issues; firstly, the new HMMs should be biologically interpretable, and secondly we need to control the complexity of the HMM so that it has good generalisation performance on unseen sequences. In this paper we explore the possibility of using a Genetic Algorithm (GA) for optimising the HMM structure. GAs are sufficiently flexible to allow incorporation of other techniques such as Baum-Welch training within their evolutionary cycle. Furthermore, operators which alter the structure of HMMs can be designed to favour interpretable and simple structures.

Results: In this paper a training strategy using Genetic Algorithms is proposed, and it is tested on finding HMM structures for the promoter and coding region of the bacterium *C. jejuni*. The proposed Genetic Algorithm for Hidden Markov Models (GA-HMM) allows HMMs with different numbers of states to evolve. To prevent over-fitting a separate data set is used for comparing the performance of the HMMs to that used for the Baum-Welch training. The GA-HMM was capable of finding an HMM comparable to a hand-coded HMM designed for the same task, which has previously been published.

Contact: j.won@ecs.soton.ac.uk

INTRODUCTION

The vast amount of biological sequence data being produced has been accompanied by the development of machine learning techniques to extract as much information from the data as possible. One of the most successful class of techniques for analysing biological sequences has been Hidden Markov Models. Although these models were originally applied to speech recognition, see *e.g.* (Rabiner, 1989), they have proved highly successful for modelling biological sequences, see *e.g.* (Durbin *et al.*, 1998). Also, they have shown good performance in predicting the structure of biological sequences (K.Asai *et al.*, 1993). In this paper, the secondary structure of proteins was predicted by the HMMs. The success of HMMs owes

much to their ability to encode biological information in their structure while allowing many unknown quantities to be learnt through the optimisation of their transition and emission probabilities. The constraints imposed by their structure will often limit excessive over-fitting of the training data, although some over-fitting is still observed when using Baum-Welch training.

Automatic optimisation of the structure of HMMs would potentially be highly beneficial. In many applications the biological mechanism is not well understood. Given sufficient data it would be possible to learn an HMM architecture that encodes some biological information that we are unaware of. However, in learning the structure of an HMM we do not wish to lose the advantages they currently offer. In particular, we wish to control the complexity of the HMM models and if possible retain their biological interpretability. In this paper we investigate the effectiveness of using Genetic Algorithms for optimising the HMM structure. A Genetic Algorithm is a robust general purpose optimisation technique which evolves a population of solutions (Goldberg, 1989). It is easy to hybridise other algorithms such as Baum-Welch training within a GA. Furthermore, it is possible to design operators which favour biologically plausible changes to the structure of an HMM. GAs have been widely used to optimise architecture for Neural Networks (Yao, 1999). However, to the best of our knowledge, GAs have only been used once before to optimise the structure of an HMM. Yada *et al.* (T.Yada *et al.*, 1994) used a GA to find a TATA box model. They included a term in their fitness function to penalise over-complex models, however, their results depended critically on this parameter. In this paper, we explore the use of GAs for evolving HMMs. Our GA differs from Yada *et al.* in that we split the training examples into two. One half is used for training the HMMs using Baum-Welch, the other half is used to evaluate the HMM's fitness. This reduces over-fitting of the training data. In addition, to prevent over-fitting we performed Baum-Welch only on a proportion of randomly selected individuals in the population at each generation.

SYSTEM AND METHODS

Genetic Algorithms for HMM (GA-HMM)

To discover if GAs are potentially useful for evolving HMMs we implemented a standard GA where a population of HMMs

¹To whom correspondence should be addressed

are evolved from one generation to the next. At each generation some proportion of the HMMs are trained with Baum-Welch on a test set. The fitness of the HMMs are measured on a validation set and the fitter members are selected. Finally the members are mutated and crossed-over to form then next generation. This procedure is shown in figure 1. The state labelled *genetic operations* include selection, mutation and crossover.

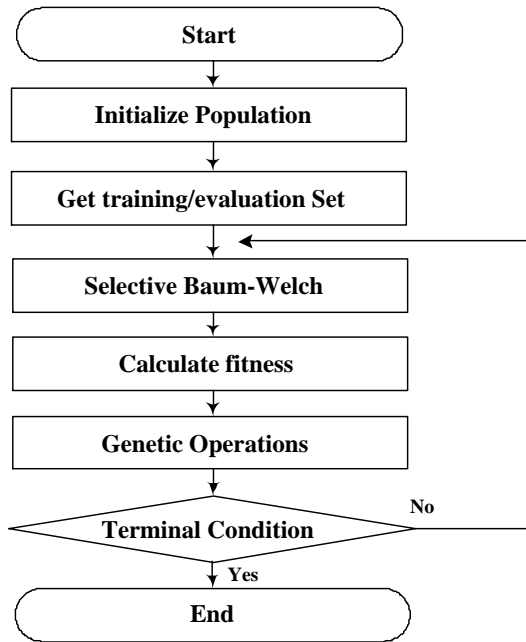


Figure 1: The GA-HMM algorithm. Baum-Welch training is combined with selection, mutation and crossover to evolve HMMs.

In the experiments described below, the initial population consists of HMMs with just two states. The number of states will change due to state insertion and deletion mutations and through crossover. Also as part of the initialisation stage the training data is divided into a set used for training with Baum-Welch and a set used for evaluating the fitness. The algorithm terminates when there is no significant change in the structural model.

Genetic Operations for GA-HMM

The genetic operations consist of selection, mutation and crossover in that order. Selection uses proportional selection with stochastic universal sampling (Baker, 1987) to reduce genetic drift. In both stochastic universal sampling and the more traditional roulette wheel selection, the sampling process can be visualized as assigning the pocket sizes of a roulette wheel to be proportional to the probability of selecting an individual. In roulette wheel selection P games are played independently

to select P individuals. In stochastic universal sampling the roulette wheel is spun once and P individuals are selected at equally spaced intervals around the wheel.

For a mutation to be useful it should make changes which cause minimal disruption so that the new HMM has a high probability of having a fitness close to that of the unmutated HMM. We considered mutations that only change either the number of states or the number of transitions by one. This gave us four mutation operators; insert state, delete state, insert transition and delete transition, which are shown in figure 2. Insertion of a state can happen between any two states or at either end of the chain. When a state is inserted, the states on its right hand side shift by one as shown in figure 2(a). The emission probabilities of the new state are set to randomly selected values. If a state is deleted, all its transitions are removed. Insertion of a transition can happen between any two states and deletion of a transition happens at any state if the state has more than one outgoing transition. Although, these mutations allow highly interconnected HMMs they provide a certain bias towards chain structures because of the state insertion operator.




Figure 2: Four types of mutations. (a) insert state (inserting a state in the second position), (b) delete state (delete the third state), (c) delete transition, (d) insert transition.

Crossover takes place between two HMMs and exchanges states. A number of successive states can be crossed over in one operation. Only outgoing transitions from a state are exchanged during crossover. An example of crossover is shown in figure 3.

Selective Baum-Welch

The Baum-Welch algorithm is commonly used to train HMMs. It estimates model parameters from training sequences while




Figure 3: Crossover. During crossover outgoing transitions move with transition.

maximising the log-likelihood of the model. The log-likelihood of model k is

$$L_k = \log(P(O|\theta_k)), \quad (1)$$

where θ_k denotes the parameters of the k th HMM and O is the training sequences.

Although Baum-Welch maximises the log-likelihood for training sequences, it can over-fit the training sequences and produce inferior results compared with HMMs that have experienced fewer training cycles. Figure 4 shows the negative log-likelihood versus the number of Baum-Welch iterations in an experiment with one of the models described later. The negative log-likelihood measured on the training data decreases monotonically, whereas on unseen data it initially decreases, but then increases again due to over-fitting. Initially, we tested a GA where Baum-Welch was performed at each generation. However, this caused over-fitting as the HMMs were trained too frequently. A selective Baum-Welch scheme was adopted to reduce over-fitting. The generalisation performance was found to improve by changing the algorithm so that Baum-Welch training only occurred with a fixed probability. In this scheme, 20 percent of population are randomly selected in each generation for the Baum-Welch training. In addition, we made sure that the fittest member of the population was never subjected to Baum-Welch training as this could cause a loss of fitness for the best member of the population. The over-fitted HMM structures receive lower fitness value when the HMM is applied to the test set. Those HMMs that had over-fit the Baum-Welch training data will be punished by the selection procedure as it uses a different set of data to evaluate the fitness.

Fitness Value

The performance of the GA was found to depend strongly on the fitness value used. The fitness value used by Yada *et al.* (T.Yada *et al.*, 1994) incorporated a balance factor between complexity and likelihood

$$w_k = \frac{1}{-2\log(L(O; \theta_k)) + 2\lambda p_k}. \quad (2)$$

Here, $L(O; \theta_k)$ is the maximum logarithmic likelihood estimate of the k th HMM, p_k is the number of free parameters in

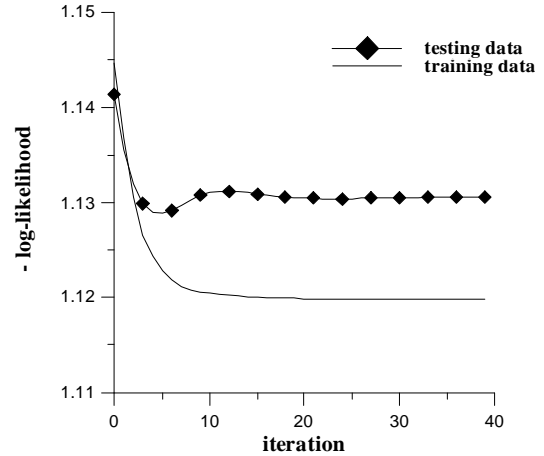


Figure 4: The graph show the negative log-likelihood for an HMM plotted against the number of Baum-Welch iterations. The negative log-likelihood for the training data is monotonically decreased by the Baum-Welch algorithm. However, the negative log-likelihood measured on independent testing data will typically decrease, reach a minimum and then increase again. In this instance the best generalisation performance was found after 5 iterations. The training data used was from *C. jejuni* and is described in experiment I below.

the k th HMM (its complexity) and λ is the balance factor. By adjusting the balance factor, the complexity can be controlled. However, the complexity was found to be very sensitive to the value of λ and it was difficult to choose a value for λ *a priori*.

In this paper, we avoided the problem of setting the balance factor by separating the training data into a set used for Baum-Welch training (O_{train}) and a set used for evaluation of the generalisation ability (O_{eval}). The fitness function used in our GA-HMM algorithm is

$$w_j = \frac{1}{-\log(L(O_{eval}; \theta_j))} \quad (3)$$

This method reduces over-fitting. As the training and testing sets are separated, over-fitted models get lower fitness values in the fitness evaluation stage and will be discarded through the selection operation.

The probability of choosing member k during the selection is then given by

$$F_k = \frac{w_k}{\sum_{j=1}^P w_k} \quad (4)$$

RESULTS AND DISCUSSION

Experiment I : Promoter model of C.Jejuni

Our first experiment was to find an HMM for the RpoD promoter region upstream of the TATA-box, see (Petersen *et al.*,

2003). The full promoter model developed by Petersen *et al.* is shown in figure 5. The RpoD promoter consists of approximately repeated blocks with an observed period of 10.6 nucleotides. A hand crafted model for this region used in (Petersen *et al.*, 2003) is shown in figure 6.




Figure 5: Model for predicting the promoter region of *C. jejuni* from L. Petersen *et al.*, (2003). In experiment I we try to learn the periodic region, starting from HMMs with two states.




Figure 6: The hand constructed HMM model for the periodic promoter region of *C. jejuni* used in L. Petersen *et al.*, (2003).


In our experiment we evolved a population of HMMs for this region using a GA starting from two states models. The parameters controlling the GA are shown in table 1. The data set consists of 175 sequences and 135 sequences were used for training and 40 for testing. Of the 135 training sequences 95 were used for Baum-Welch training and 40 used for measuring the fitness. The experiment was repeated five times (5-fold cross validation) to obtain better statistics. The data used in this experiment can be obtained from <http://www.ecs.soton.ac.uk/~kjlw02r/data.html>.

Table 1: GA-HMM parameters used in the experiment 1


Parameter	value
Population size	30
Offspring size	4
Iteration	300
Crossover rate	0.06
Insert / delete transition	0.06
Insert / delete state	0.12

Figure 7(a) shows the fitness of the best member of the population versus the generation for the first data partition. Only test sequences which are not involved in the training are used in calculating the fitness value (Eq. 3). Because the

GA-HMM reduces over-fitting, the fitness value of the fittest member of the population always increases. The graphs for the other cross validation sets look similar. In figure 7(b) the average number of states in the HMM is shown. We observe that the number of states grows until it reaches around 10 which appears to support the observed periodicity found in this region of the promoter. At this stage the number of states remains roughly fixed while the fitness value continues to rise indicating that the fine structure of the model is still evolving. Occasionally, the number of states found by the GA-HMM was less than 10 states. We believe this is because the GA-HMM found a local optimum around 5-6 states.



(a)



(b)

Figure 7: Results during GA-HMM training: (a) shows the fitness value of fittest individual on each iteration (b) shows average number of states for periodic signal. The GA started with a population consisting of 2 states. After 150 generations the HMM have a length of 10 states. Although the length does not significantly change thereafter the fitness continues to improve indicating that the finer structure is being fine tuned.

An example of one of the HMM structures found by the GA is shown in figure 8. The model is considerably more complex than that proposed by Petersen *et al.* shown in fig-

ure 6. Table 2 shows a comparison of the fitness values (w_k) defined in equation (3) and their variances obtained from the 5 cross-validation experiments for our GA-HMM. In the same table we show the fitness of the model proposed in (Petersen et al., 2003) and trained using Baum-Welch. The hand crafted model is trained only with the Baum-Welch algorithm. Of 175 sequences, for each test 140 sequences are used for the Baum-Welch training and 35 sequences are used to calculate the fitness. Because Baum-Welch algorithm finds a locally optimal solution, the final results depend on the initial values of the transition and emission probabilities. For the hand crafted model and the GA model, we conducted ten experiments with different initial values.




Figure 8: After training the *C. jejuni* sequences, GA-HMM found one structure model for the periodic signal

Table 2: Comparison between Baum-Welch and GA-HMM. The results show the fitness value ($-1/L(O_{eval}; \theta)$) and variance for five different partitionings of the data.

Test	Petersen's model	GA-HMM
1	0.8855(3.0e-6)	0.8863(1.3e-6)
2	0.8854(1.1e-6)	0.8861(1.9e-6)
3	0.8887(3.9e-6)	0.8877(2.4e-6)
4	0.8796(5.6e-7)	0.8797(4.5e-7)
5	0.8713(2.0e-6)	0.8767(4.3e-6)

The performance of the HMM found using the GA and the hand designed HMM are roughly similar. Even though it does not show better performance in test 3, it produces slightly better result in the other tests. The paired t-test was conducted to assess the statistical difference (T.M.Mitchell, 1997). This analysis compares the means of two groups to see if those two groups are statistically different. Let $X_{GA}(i)$ and $X_P(i)$ be the results for the GA-HMM and Petersen et al.. Then, the t-value for the k samples can be defined by

$$t - value = \frac{\bar{d}}{\sqrt{\sigma_d^2/k}} \quad (5)$$

where \bar{d} is the mean of the differences $X_{GA}(i) - X_P(i)$, and σ_d is the standard deviation of this mean.

This statistic has $n - 1$ degrees of freedom. The t-values for the each test set are given in table 2. In this experiment, the degrees of freedom is 4. With this result the significance

level at which two distributions differ can be determined. The t-value obtained is 1.09. By using the t-value and the degrees of freedom the probability of null hypothesis can be obtained from the table of t-distribution. To get more precise probability we interpolate the values in the table of t-distribution. The probability of this result, assuming the null hypothesis, is 0.336. This result says that about 66 times out of 100 cases a statistical difference between two algorithms can be found. From this test we can conclude that the GA-HMM has slightly better performance than the hand-crafted model overall.

Experiment II : Coding region model of C.Jejuni

To see if a GA was capable of finding biologically interpretable solutions we performed a second experiment using 556 sequences from the coding regions from *C. jejuni*. The sequence data comprised a start codon (ATG), some number of codons and stop codon (TAA, TAG or TGA). A simple HMM architecture for detecting this region would consist of a 3 state loop (figure 9). A third order state model was used to model the codon region. That is, instead of a state emitting a symbol from a single base unconditional distribution, the state emits a symbol, which is dependent on the three previous bases in the sequence, through a conditional probability distribution. We again used a 2-state HMM as an initial model for the coding region as shown in figure 10. The transition and emission probability for the start codon and stop codon was not evolved. Each state is third order. The GA was run for 700 generations. Of the 566 sequences 160 sequences were used for testing. The other parameters are set as in experiment I.




Figure 9: HMM architecture for *C. jejuni* coding region




Figure 10: Initial HMM architecture

Figure 11 shows some of the HMM architectures produced by the GA in this experiment. Even though some results were hard to interpret, almost every simulation produced an interpretable model. Figure 11(a) shows the same result as the common model. Figure 11(b) has a path 1-3-4,(2-3-4)—here, bracket indicate the a loop. Figure 11(c) has (1-5-6) and (1-2-3, 4-5-6) loops. Figure 11(d) has 1-2-3, (9-7-8), 9-10-11 or 1-2-3, 4-5-6, (6-6-6), 6-7-8, (9-7-8), 9-10-11 loops. These HMM solutions, although complicated, predominantly

generate triplets. We performed ten experiment and on every experiment with coding region we could find the clean 3 state loops as shown Figure 11 (a)(b)(c). Figure 11(d) shows the worst result we have. The numbers shown are transition probabilities. Other transition probabilities are easily calculated, because the sum of outgoing transitions from a state is always 1. Because of the transition probabilities the most probable loop are 1-2-3, (9-7-8) or 1-2-3, (9-7-8), 9-10-11.




Figure 11: Result of simulation for the *C. jejuni* coding region: (a) has loop (1-2-3), (b) has a path 1-3-4, (2-3-4) (c) has (1-5-6) and (1-2-3, 4-5-6) loops. (d) has 1-2-3, (9-7-8), 9-10-11 or 1-2-3, 4-5-6, (6-6-6), 6-7-8, (9-7-8), 9-10-11.

DISCUSSION

The experiments described here suggest that genetic algorithms are quite capable of finding reasonable HMM architectures for biological sequence analysis. Even with a rather naive implementation, the GA was able to achieve comparable or slightly superior generalisation performance to a hand designed HMMs. Both experiments show that a major drawback of automating the design of HMM architectures is that the resulting model may be difficult to interpret biologically. Although we used genetic operators that favoured the building of chain structures they nevertheless allowed considerable cross linking within the chain. The mutation operators could be modified to strongly bias biologically interpretable models. This is an area we are currently investigating. A drawback of constraining the search is that we may inhibit the GA from discovering completely novel types of architecture. The crossover operators can swap a series of states in one operation. The emission probabilities are also crossed over with the states. This enables the GA-HMM to swap any meaningful part of the HMM structure.

We prevented the GA from producing over-complex models by measuring the fitness on a different set of data from

that used for training Baum-Welch. This ensures reasonable solutions but did not entirely exclude over-complex models from appearing. Yada *et al.* (T.Yada *et al.*, 1994) incorporated a penalty term depending on the number of states in their fitness function. Although, we found such penalty terms difficult to use because they required careful parameter selection to work, they may be beneficial if applied properly. This is another area that deserves further research.

One of the merits of GAs is the capability of dealing with substructures of the solution. The proposed crossover scheme can treat such modularity in some sense, even though the strategy does not completely implement the modular structure. Before our GA-HMM method, the Genetic Programming methods that capture the hierarchical structures were developed (T.Yada, 1998; T.Yada *et al.*, 1995). They used S-expressions to encode the network structures and applied genetic programming to find the optimal structure.

There were some attempt to organize probabilistic model without prior knowledge. Stolcke (Stolcke, 1994) used log-likelihood and posterior probability to merge the HMM states from the initial model. Because this approach was used to get compact HMM architecture by merging states, it does not have the splitting state operation which is useful in dealing with new data. State splitting methods as well as deleting negligible states and transitions were used to find a optimal HMM topology (Y.Fujiwara *et al.*, 1995). They used the transition ambiguity and the expected observation differences to split state and applied this to find a HMM structure of a leucine zipper motif. Those statistical approaches can be used to find a particular pattern like motif. However, those algorithms do not seem to find a structural model shown in the Simulation II.

Our experiments were carried out on short sections of an HMM. It seems very unlikely in the current state of development that a GA would be able to find large HMM structures *ab initio* that are competitive with hand designed architectures. Nevertheless, even in the short term GAs may be able to ‘tune’ a hand designed HMM especially in areas where the biological significance of a region is poorly understood.

Acknowledgements

We would like to thank Lise Petersen for her help with the sequence data, the model and her valuable comments on this work.

REFERENCES

- Baker, J. E. (1987) Reducing Bias and Inefficiency in the Selection Algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms*. Lawrence Erlbaum Associates (Hillsdale).
- Durbin, R., Eddy, S., Krogh, A., and Mitchison, G. (1998) *Biological sequence analysis*. Cambridge. Cambridge University Press.

- Goldberg, D. E. (1989) *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley (Reading, Mass).
- K.Asai, S.Hayamizu, and K.Handa (1993) Prediction of protein secondary structure by the hidden Markov model. *Computer Applications in the Biosciences*, **9**, 141–146.
- Petersen, L., et al. (2003) RpoD promoters in *Campylobacter jejuni* exhibit a strong periodic signal instead of a -35 box. *Journal of Molecular Biology*, **326**, 1361–1372.
- Rabiner, L. R. (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, **77**, 257–286.
- Stolcke, A. (1994) *Bayesian Learning of Probabilistic Language Models*. Ph.D. thesis, University of California at Berkeley.
- T.M.Mitchell (1997) *Machine Learning*. McGraw-Hill Companies, Inc.
- T.Yada (1998) *Stochastic Models Representing DNA Sequence Data - Construction Algorithms and Their Applications to Prediction of Gene Structure and Function*. Ph.D. thesis, University of Tokyo.
- T.Yada, M.Ishikawa, H.Tanaka, and K.Asai (1994) DNA Sequence Analysis using Hidden Markov Model and Genetic Algorithm. In *Genome Informatics Vol.5*, pp.178-179.
- T.Yada, Y.Totoki, K.Asai, and M.Ishikawa (1995) Generation of Hidden Markov Model Describing Complex Motif in DNA Sequences. In *IPSJ Transactions*, pages 750–767.
- Yao, X. (1999) Evolving Artificial Neural Networks. *IEEE Press, USA, Proceedings of the IEEE*, **87**, 1423–1447.
- Y.Fujiwara, M.Asogawa, and A.Konagaya (1995) Motif Extraction using an Improved Iterative Duplication Method for HMM Topology Learning. In *In Pacific Symposium on Biocomputing '96*, pages 713–714.