

# All binary representations are equal: but some are more equal than others

Kai Willadsen <sup>†</sup> and Janet Wiles <sup>†‡</sup>

<sup>†</sup> School of Information Technology and Electrical Engineering

<sup>‡</sup> School of Psychology

The University of Queensland

Queensland, 4072

Australia

{janetw, kaiw}@itee.uq.edu.au

*Abstract*— The original Hinton and Nowlan simulation demonstrating the Baldwin effect is well known and serves as an interesting basis for GA research. Belew varied the original representation to use a binary code, a side effect of which was that learning became a substitute for internalised knowledge. We varied the binary representation such that learning became an expression of uncertainty. This simple modification to the representation resulted in an interesting and non-trivial set of interactions between GA operators and the representation used.

## I. INTRODUCTION

GAs are highly dependant on the problem representation that they are provided with. The choice of representation may determine the ability of an algorithm to solve a given problem, and it is highly likely that different representations of a problem will cause the same problem to be solved in different ways.

This paper investigates the differences between two very simple implementations of a simulation of a well studied problem - the Baldwin effect.

### A. The Baldwin effect

An effect designed to demonstrate the possible influence of learning upon evolution was simultaneously posited by Baldwin[1], Morgan[2] and Osborn[3], but has since been labelled ‘The Baldwin effect’. There is no simple way of testing the theory biologically, and as a result, the Baldwin effect was (and still is) largely discounted by biologists. The Baldwin effect posits that it may be possible for learnt characteristics to become incorporated into a genome through selection for an *ability to learn*. In other words, over evolutionary timescales, the ability to learn fact  $A$  can be transformed into an innate knowledge of fact  $A$ .

### B. Computational models of the Baldwin effect

Hinton and Nowlan [4] were the first to demonstrate the computation feasibility of the Baldwin effect. They achieved this by incorporating a very simple learning mechanism into a population, and demonstrating that such a learning mechanism could turn a needle-in-a-haystack problem into a feasible search. While certain aspects of the simulation have been criticised, (i.e., the learning mecha-

nism) the simulation as a whole still remains a very influential example of the use of computer simulations to investigate theories that could otherwise never be tested. In addition, the Hinton and Nowlan needle-in-a-haystack task has been used to study a variety of effects in GAs such as exploration vs. exploitation, and their effects on residual learning [5].

Neural networks have also been use in the investigation of the Baldwin effect, and there have been studies incorporating Baldwin-style effects in the evolution of neural networks [6]. Such studies demonstrate that the criticisms of the learning framework used by Hinton and Nowlan do not fundamentally undermine the effects shown. In short, the Baldwin effect is portable to more complex learning mechanisms.

The Baldwin effect has been replicated in an artificial life environment [7]. The simulations involved, while based upon the Hinton and Nowlan ones, altered *many* characteristics of the environment. Specifically, there was a dynamic (grid-based) environment, an implicit fitness function and a separation between the genotype and phenotype that was not present in the Hinton and Nowlan simulations.

### C. Further work on the Hinton and Nowlan simulation

The Hinton and Nowlan simulation parameters were such that the probability of obtaining an agent capable of finding the fitness ‘needle’ in the initial generation was high; the number of guesses an agent was allowed was tuned to provide an environment in which learning was possible, but not trivial [8]. The simulation has also been extended to incorporate mutation [9], showing that the primary effect of mutation is to disrupt the solutions present in the population. The phenomena experienced (though not discussed) by Hinton and Nowlan regarding the convergence of the GA and the existence of residual non-innate alleles have been investigated, with causes suggested such as genetic drift [10] and biases of selection algorithms [5].

The final significant influence on the Baldwin simulations in general is the notion of the *necessity* of learning. This involves an expression of the Baldwin effect, not purely as a learning mechanism, but more generally as *plasticity* [7]. In a situation where the ‘goal’ condition is a needle-

TABLE I

GENOTYPE-PHENOTYPE MAPPING IN THE HINTON &amp; NOWLAN REPRESENTATION

Phenotype	Genotype
0	0
1	1
?	2

in-a-haystack problem, the benefit of learning is alter the effective fitness landscape from a rapid-failure situation to a graceful degradation situation [4]. However, in a situation where genotypic and phenotypic fitness are not significantly different this can also lead to learning ‘hiding’ higher fitness levels [11].

## II. CHARACTERISTICS OF VARIOUS REPRESENTATIONS

### A. Ternary representation

The simulation used by Hinton and Nowlan [4] used a ternary representation to represent the three possible values that a locus on the genome may take (see Table I). This representation has the advantages of being simple and easy to code, but it also makes an implicit assumption about the way that the genetic operators work.

In a ternary representation, the fact that each allele value can be considered ‘atomic’ means that crossover and mutation act in predictable ways. In crossover splicing will always occur at the boundary of an allele group, and mutation will always alter one phenotypic allele value to a different phenotypic allele value. While this is the most simple representation from a genetic operator point of view, alternative representations may introduce different behaviour.

This model has one particular characteristic that contrasts with those of the other representations used in this paper.

- In a situation with no mutation, (as in the classic Hinton & Nowlan simulations) the representation has access only to those phenotypic values currently in the population.

### B. Belew representation

Belew [9] introduced a binary representation (see Table II). His representation focuses on implementing the three-allele system in a ‘pure’ manner, by using a binary representation that is translated into the ternary representation for evaluation.

TABLE II

GENOTYPE-PHENOTYPE MAPPING IN THE BELEW REPRESENTATION

Phenotype	Genotype
0	10
1	11
?	00
?	01

TABLE III

MUTATION IN THE BELEW REPRESENTATION

Phenotype	Mutate to	Silent
?	?	Yes
	0	No
?	?	Yes
	1	No
0	?	No
	1	No
1	0	No
	?	No

TABLE IV

CROSSOVER INTERACTIONS IN THE BELEW REPRESENTATION

Original Phenotypes		Final Phenotypes	
Parent 1	Parent 2	Child 1	Child 2
?	1	?	0
?	0	?	1
0	?	1	?
1	?	0	?

As Belew originally observed, this representation produces interesting behaviour that is not observable in the ternary representation.

- In a situation without mutation, the representation can ‘recover’ phenotypic values that have ceased to exist in the population. The allele pairs [0,0] and [1,1] may be crossed over to produce [0,1] and [1,0], representing a transformation from a ‘?’ and ‘1’ to a ‘?’ and ‘0’ (see Table IV). This is a side-effect of allowing crossover to occur at *any* locus, not just on the boundaries between allele pairs.
- In a situation with mutation, the representation allows ‘silent’ mutations. The allele pairs [0,1] and [0,0] both represent a ‘?’, and therefore a change to the second allele in a pair may constitute a silent mutation (see Table III).

### B.1 Problems in the Belew representation

While the interactions introduced by Belew [9] produce similar results to those obtained by Hinton and Nowlan [4], they also alter some of the dynamics of the simulation. In addition, while the use of silent mutations in this representation is interesting, it is also unnecessary in the context of the simulation.

In considering the interactions that Belew’s representation created, it should be clear that the effects of some operators (crossover in particular) are not intuitive. Rephrasing the allele values to a trial-based context,

Parent *A* will eat food *x*, parent *B* might eat food *x*. (1, ?)  
 Of their children, child *C* won’t eat food *x* and child *D* might eat food *x*. (0, ?)

Clearly, while silent mutations are valuable as far as cre-

TABLE V

GENOTYPE-PHENOTYPE MAPPING IN THE CURRENT REPRESENTATION

Phenotype	Genotype
0	00
1	11
?	01
?	10

ating neutrality in a landscape goes, this particular implementation appears to sacrifice conceptual consistency with the underlying model as a side effect.

### C. Current representation

In order to address the problems raised with Belew’s implementation of the Baldwin effect simulation, a new representation was designed (see Table V) which represented the underlying semantics of the system in a way that emphasised the idea of learning as a complement to innate knowledge, rather than as an alternative.

As with Belew’s representation, this representation results in interesting interactions not present in the ternary case.

- In a situation without mutation, the representation can ‘recover’ phenotypic values that have ceased to exist in the population. The allele pairs [0,1] and [1,0] may be crossed over to produce [0,0] and [1,1], representing a transformation from a ‘?’ and ‘?’ to a ‘1’ and ‘0’ (see Table VII). As in the Belew representation, this is a side-effect of allowing crossover to occur at *any* locus, not just on the boundaries between allele pairs.
- The representation does *not* allow silent mutations.

Belew’s representation utilised a single bit as a ‘switch’ for the learnability of a gene. This places an emphasis on learning as an *alternative* to genetic information. In contrast, the proposed representation implies that learning is a mechanism of *uncertainty*. In other words, learning is of use when the genome does not (or cannot) contain the appropriate information. This represents the idea of learning creating a large area of high-fitness in a needle-in-a-haystack fitness landscape.

In the current representation, both the mutation and

TABLE VI

MUTATION IN THE CURRENT REPRESENTATION

Phenotype	Mutate to	Silent
0	?	No
	?	No
?	0	No
	1	No
?	1	No
	0	No
1	?	No
	?	No

TABLE VII

CROSSOVER INTERACTIONS IN THE CURRENT REPRESENTATION

Original phenotypes		Final phenotypes	
Parent 1	Parent 2	Child 1	Child 2
0	1	?	?
?	?	0	1
?	?	1	0
1	0	?	?

crossover operators have less radical behaviour than in the Belew representation. Mutation no longer has the possibility of producing silent mutations (see Table VI), whereas the crossover of parent’s information acts more sensibly from the point of view of confidence in a genome (see Table VII). Repeating the crossover example given earlier for the Belew representation,

Parent *A* will eat food *x*, parent *B* might eat food *x*. (1, ?)

Of their children, child *C* might eat food *x* and child *D* will eat food *x*. (?, 1)

More noticeable however, is the representation’s ability to resolve conflicts between parent genes,

Parent *A* will eat food *x*, parent *B* won’t eat food *x*. (1, 0)

Of their children, child *C* might eat food *x* and child *D* might eat food *x*. (?, ?)

As can be seen from the preceding examples, this representation focuses on learning as a mechanism arising from uncertainty on the part of an organism, rather than as an substitute for internalised knowledge.

## III. EXPERIMENTAL CHARACTERISTICS OF REPRESENTATIONS

### A. Base case

The initial set of simulations duplicate the original Hinton and Nowlan [4] simulations. In many subsequent studies, the results obtained by Hinton and Nowlan have not been exactly duplicated, for reasons that remain unknown. Specifically, the extremely high frequency of ‘?’ alleles (and corresponding low frequency of ‘1’ alleles) in the original results has failed to be duplicated on many occasions [8], [10]. However, the results obtained in this simulation agree with those obtained by other Baldwin-related simulations.

The base set of simulations was run with a high probability of crossover but without mutation for a period of 50 generations with 1000 learning trials per agent. A population’s progress is tracked by the relative proportions of ‘1’, ‘0’ and ‘?’ alleles in the population. In doing this, we observe that all of the simulations tend towards the elimination of ‘0’ alleles early on, matched by a gentle rise in the occurrence of ‘1’ alleles, and a corresponding decrease in the occurrence of ‘?’ alleles (see Figure 1). This represents the incorporation of learnt characteristics into the genome as hard-wired reactions. However, it is also obvi-




Fig. 1. Distribution of alleles in the population for the ternary representation. The results have been averaged over 100 runs.

ous that the incorporation experiences a slowdown, and will likely asymptote before ‘?’ alleles are eliminated from the genome completely. This problem has been investigated before and has been attributed to problems with convergence to monomorphic question marks in the population [10], [5].

The Belew representation has the lowest percentage of ‘1’ alleles present in the genome by the end of the simulation (see Figure 2). This may seem contrary to expectations as the Belew representation has twice as many alleles as the ternary representation; a greater number of alleles should slow convergence in a genetic algorithm. As pointed out by Harvey [10], there will be strong initial selection against ‘0’ alleles, which will be likely to select against *all* innate alleles, reducing the distribution of the ‘non-learning’ genotypic allele in the population. The current representation demonstrates less rapid convergence than both the ternary and the Belew representation.

Another important factor in these simulations is the reliability of the representation in terms of the expected fitness of an agent. In these simulations, the standard deviation of the fitnesses of agents in a population using the Belew-based representation was higher than that of populations using the current representation (see Table VIII).

This implies that the Belew representation had a wider range of allele distributions through the population, meaning that it maintained fitness levels less reliably than the other representations.

TABLE VIII  
COMPARISON OF FITNESSES IN STANDARD SIMULATION (100 RUNS)

	Fitness	Std dev
Ternary	19.227	1.185
Belew	18.935	2.296
Current	19.328	1.370




Fig. 2. Comparison of the distributions of ‘1’ and ‘?’ in the population for Belew’s representation and current representation. The results have been averaged over 100 runs.

### B. The effects of mutation

Belew [9] briefly discusses the use of mutation within the Hinton & Nowlan model, and asserts that low mutation rates make very little difference to a simulation, while high rates of mutation make it impossible for a solution to be maintained. Mutation (in theory) should act to maintain diversity within the population, but at the cost of disrupting successful genomes. As a result, resilience to the detrimental effects of mutation is an important feature of a representation.

The interesting interactions between the mutation operator and the Belew representation (see Table III) suggest a significant difference between the behaviour of the Belew representation and the current representation under mutation. The Belew representation allows a single-bit mutation from a ‘1’ allele directly to a ‘0’ allele, without an inter-

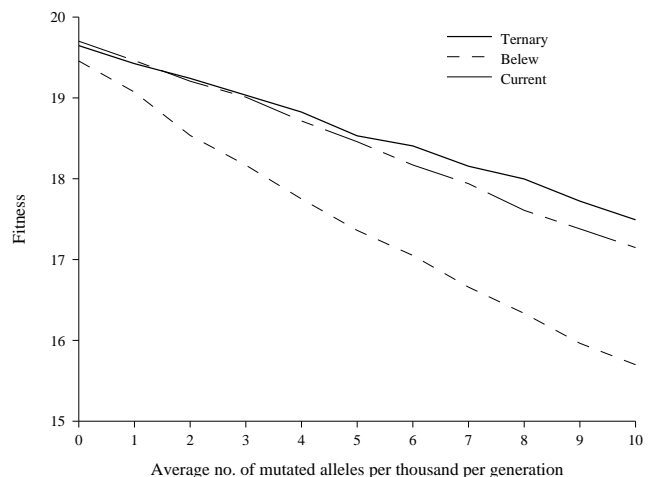


Fig. 3. Even small mutation rates cause significant decreases in the average fitness of a population, with the Belew representation being the worst effected. The results have been averaged over 50 runs.

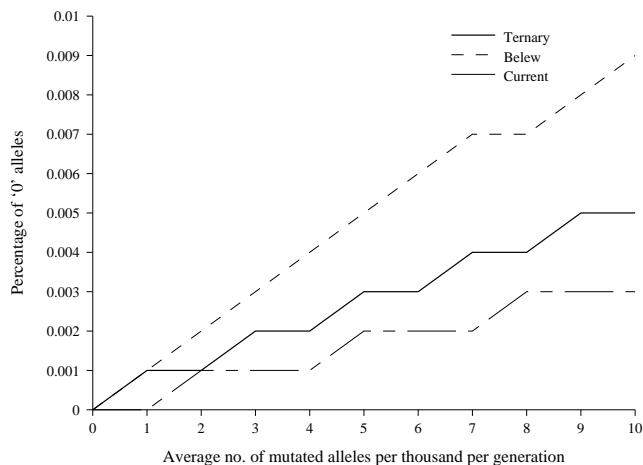


Fig. 4. The Belew-based representation demonstrates more rapid deterioration than either the ternary or current representation as a result of mutation. The results have been averaged over 50 runs.

mediate step, whereas a single-bit mutation in the current representation will only allow mutation from a ‘1’ to a ‘?’ (see Table VI). As a result, the performance of the Belew representation degrades more rapidly because of mutation (see Figure 3) and the occurrences of ‘0’ alleles in the population increases rapidly (see Figure 4).

#### IV. CONCLUSIONS

The overall message conveyed by these results is that representation matters. While this is a widely accepted fact, the question of ‘How does it matter?’ remains a valid one. Here it has been shown that the interactions between the genetic operators of mutation and crossover and a problem representation can be non-trivial. In this case, the complexity arises from the indirect genotype to phenotype mapping used by these representations, though similar effects can be imagined in any situation where there are interactions between loci at a genotypic (or phenotypic) level. As has been shown, the choice of representation can make a difference, not only in performance terms, but also in the way in which it interacts with the underlying ideas of a simulation. Different representations will emphasise different aspects of a simulation, and foreknowledge of the tradeoffs that this implies is vital.

#### REFERENCES

- [1] J. Baldwin, “A New Factor in Evolution,” *American Naturalist*, vol. 30, pp. 441–451, 1896.
- [2] C. Morgan, “On Modification and Variation,” *Science*, vol. 4, pp. 733–740, 1896.
- [3] H. Osborn, “Ontogenic and Phylogenic Variation,” *Science*, vol. 4, pp. 786–789, 1896.
- [4] G. Hinton and S. Nowlan, “How learning can guide evolution,” *Complex Systems*, vol. 1, pp. 495–502, 1987.
- [5] J. Wiles, R. Schulz, J. Hallinan, S. Bolland, and B. Tonkes, “Probing the persistent question marks,” in *Proceedings of the 2001 Genetic and Evolutionary Computation Conference (GECCO-2001)* (L. S. et al., ed.), pp. 710–717, 2001.
- [6] F. Gruau and D. Whitley, “Adding learning to the cellular development of neural networks: Evolution and the baldwin effect,” *Evolutionary Computation*, vol. 1, no. 3, pp. 213–233, 1993.

- [7] R. French and A. Messinger, “Genes, Phenets and the Baldwin Effect: Learning and Evolution in a Simulated Population,” in *Artificial Life IV* (R. B. . P. Maes, ed.), pp. 277–282, 1994.
- [8] R. Belew, “When both Individuals and Populations Search: Adding simple learning to the Genetic Algorithm,” in *Proceedings of the Third International Conference on Genetic Algorithms* (J. D. Schaffer, ed.), pp. 34–41, 1989.
- [9] R. Belew, “Evolution, Learning and Culture: Computational metaphors for adaptive algorithms,” *Complex Systems*, vol. 4, no. 1, pp. 11–49, 1990.
- [10] I. Harvey, “The Puzzle of the Persistent Question Marks: A Case Study of Genetic Drift,” in *Genetic Algorithms: Proceedings of the Fifth International Conference* (S. Forrest, ed.), pp. 15–22, 1993.
- [11] G. Mayley, “Guiding or hiding: Explorations into the effects of learning on the rate of evolution,” in *The Proceedings of the Fourth European Conference on Artificial Life* (P. Husbands and I. Harvey, eds.), 1997.