Evolving Variable Plasticity in Neural Systems

JOHN A. BULLINARIA
Department of Psychology
University of Reading
Reading, RG6 6AH
UK
j.bullinaria@reading.ac.uk http://www.reading.ac.uk/~sxs97jb

Abstract: This paper considers the evolutionary forces that might determine how neural plasticity varies during an individual's lifetime. Explicit simulations of an evolving population of adaptable artificial neural systems suggest that there may be a real advantage of having a plasticity that declines with age, quite independently of the physical overheads of that plasticity.

Key-Words: Plasticity, Adaptation, Evolution, Neural Networks, Baldwin Effect, Control Systems.

1 Introduction

There is an old saying that "you cannot teach old dogs new tricks". This might not be totally true, but there is certainly considerable evidence that neural plasticity does decrease with age [4]. One plausible reason for this is that the necessary resources are depleted or removed early in life, leaving little for the "new tricks" later on. In our evolutionary past this may not have been too detrimental, since individuals could probably cope sufficiently well with what they had learned early on in their short lives. There would consequently have been little pressure to preserve a large potential for neural plasticity late into life with its inherent cost in terms of physical resources. When it comes to building artificial adaptable systems, however, the overheads of learning and accidents of evolutionary history are not so important. The question to be considered in this paper, is whether there is any advantage in an adaptable system of having a plasticity that declines with age, quite independently of the physical overheads of that plasticity.

The approach I shall adopt is to consider a simplified generic neural network control system with four adaptable parameters (i.e. connection weights) each with its own plasticity (i.e. learning rate). It will also have two additional parameters that specify how the plasticity varies with age. The system thus has ten innate parameters: four initial weights, four initial learning rates, and two plasticity variation parameters. We can consider these to specify the genotype of each individual system. Each such system is then expected to adjust its weights by a simple gradient descent learning algorithm so that it performs as best it can on the training data it is given. The important feature of this approach is that we do not just consider one individual system, but a whole population of them that learn, procreate and die over many generations. By implementing reasonably realistic natural selection, procreation and mutation for this process, the population evolves into better and better adaptable controllers. As this process approaches asymptote we shall see what values of the ten parameters are optimal. In particular, we shall see if there is any advantage to having neural plasticities that fall with age in a manner allowed by our parameterization.

2 The Simplified Control Model

The simplified control system that will form the basis of the investigation is shown in Figure 1. The input is a sequence of target responses and a feedback loop allows the determination of an error signal. This signal feeds into simple integral and proportional controllers, the outputs of which are added to bias and tonic signals, and fed into the plant to produce the response. The bias provides an appropriate resting state, and the tonic allows short time-scale adaptation of the resting state during periods of constant demand. The system can be regarded as a fully dynamical network of leaky integrator neurons. In the human accommodation (eye focussing) system, for example, we have blur being processed to generate signals for the ciliary muscles in the eye appropriate for the distance of the visual target. In our model the four adjustable parameters (weights *WC*, *WP*, *WT*, and bias *WB*) are learned by a simple gradient descent algorithm that minimizes a cost function consisting of response error and regularization (smoothing) components which will be readily available to the system. Corresponding to these learnable weights, each instantiation of the model also has four fixed initial

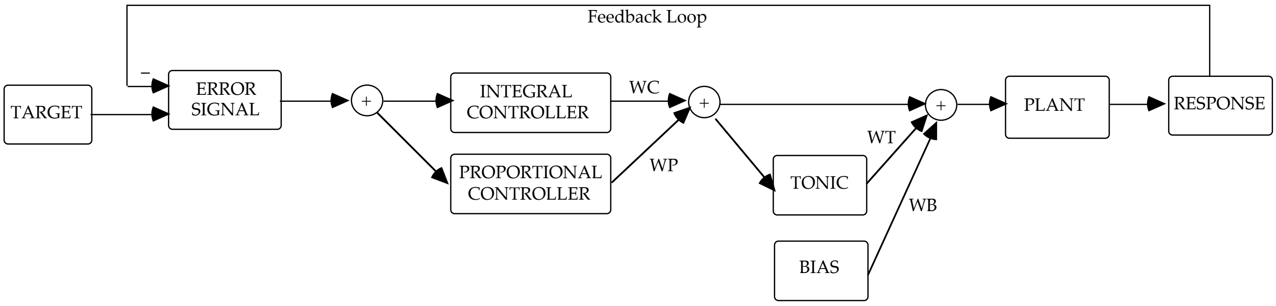

Figure 1: Simplified generic control model with four learnable parameters.

weights (*iWC*, *iWP*, *iWT*, *iWB*), four variable learning rates rates (*eWC*, *eWP*, *eWT*, *eWB*), and four fixed initial learning rates (*ieWC*, *ieWP*, *ieWT*, *eWB*). The model also has various other parameters (neuron time constants, plant characteristics, feedback time delay, and so on) which we take to be the same for all instantiations. Such a system that has learned/evolved a good set of parameters will produce appropriate damped responses to arbitrary discontinuous output requirements (steps) and smooth pursuit of arbitrary continuous output changes (ramps).

This model is a simplified version of a model used elsewhere [3] to explore the Baldwin Effect [1, 2] in adaptable control systems. Here the system has two additional parameters in the genotype that allow the learning rates to vary with age. For completeness, we consider two different parameterizations, one deterministic variation, and one that will adapt appropriately to the environment:

Deterministic: The learning rates are constant till age *pAGE*, after which they increase or decrease exponentially by a constant scale factor 1+*pSF* every ten weeks. This is easily implementable in real or artificial systems, but could potentially leave an individual vulnerable if it needs to adapt to variable environments late in life.

Adaptable: When two weight changes are consistent (i.e. in the same direction) the corresponding learning rate changes by a constant scale factor 1+*pSFI*, and when they are inconsistent (i.e. in opposite directions) the corresponding learning rate changes by a constant scale factor 1+*pSFD*. This is harder to implement than the deterministic case but allows the different learning rates to adapt independently and allows the plasticity to increase again should the environment change.

It is not uncommon in the neural network literature to see modellers varying their network learning rates during the course of training [5]. The rates may be decreased near the end of training to minimise the weight variations seen after each sample in online training, or increased to speed the saturation of sigmoids as the errors become small. The aim here is to see if such strategies evolve naturally, and if different strategies will evolve under different circumstances.

3 Evolving the Model

Simulating an evolutionary process for our model involves taking a whole population of individual instantiations and allowing them to learn, procreate and die in a manner approximating these processes in real (living) systems. The genotype of each new individual will depend on the genotypes of its two parents and random mutation. Then during their life the individuals will learn from the environment how best to adjust their weights to perform most effectively. Finally, each individual eventually dies, perhaps after producing a number of children.

In realistic situations, the ability of an individual to survive or reproduce will rely on a number of factors which can depend in a complicated manner on that individual's performance on a range of related tasks (food gathering, fighting, running, and so on). For the purposes of our simplified model, we shall consider it to be a sufficiently good approximation to assume a simple linear relation between our single task fitness function and the survival or procreation fitness. In fact, any monotonic relation will result in similar evolutionary trends.

It seems appropriate to follow a more natural approach to procreation, mutation and survival than many evolutionary simulations [2]. Rather than training each member of the whole population for a fixed time and picking the fittest to breed and form the next generation, our populations contain competing learning individuals of all ages, each with the potential for dying or procreation at each stage. During each simulated year, each individual learns from their own experience with a new randomly generated common environment (i.e. set of training/testing data) and has its fitness measured. A biased random subset of the least fit individuals, together

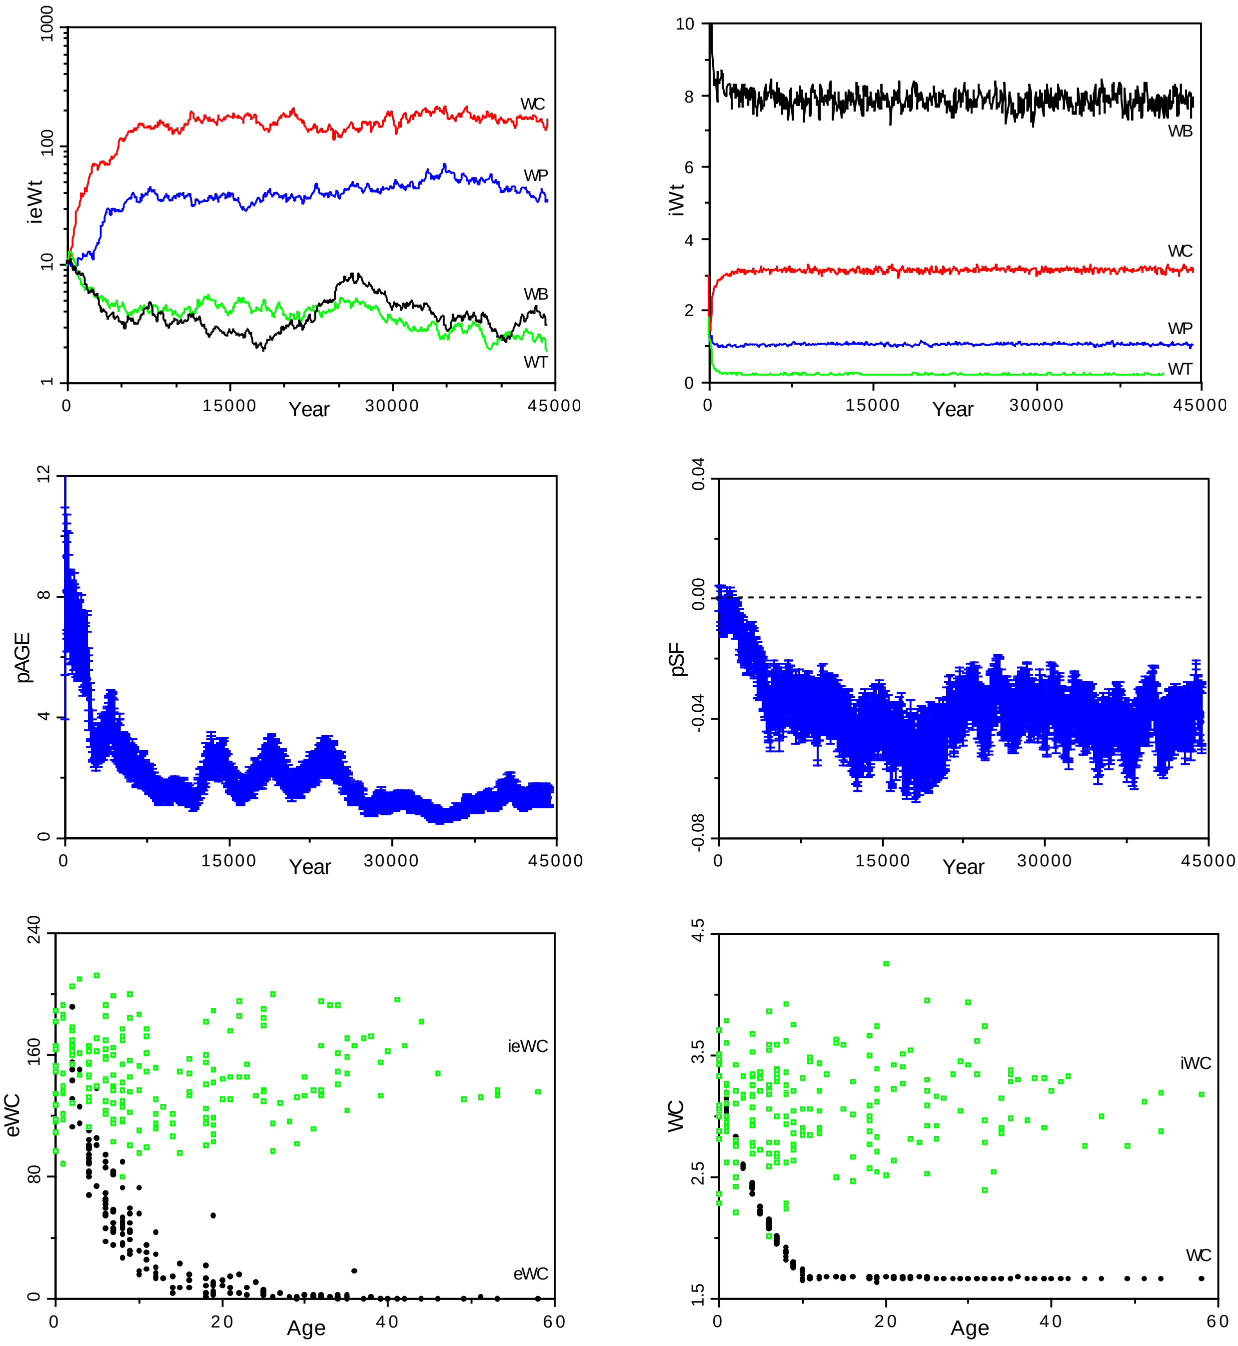

Figure 2: Evolution of the basic system with deterministic plasticity variation.

with a flat random subset of the oldest individuals, then die. These are replaced by children, each having one parent chosen randomly from the fittest half of the population who randomly chooses their mate from the rest of whole population. Each child inherits characteristics from both parents such that each innate free parameter is chosen at random somewhere between the values of its parents, with sufficient noise (or mutation) that there is a reasonable possibility of the parameter falling outside the range spanned by the parents. Ultimately, our simulations might benefit from more realistic encodings of the parameters, concepts such as recessive and dominant genes, learning and procreation costs, different inheritance and mutation details, different survival and procreation criteria, more restrictive mate selection regimes, offspring protection, different learning algorithms and fitness functions, and so on, but for the purposes of this paper, our simplified approach seems adequate.

4 Simulation Results

Even when a good set of innate parameters have evolved, a control system will still benefit from being plastic since that will allow it to fine tune its

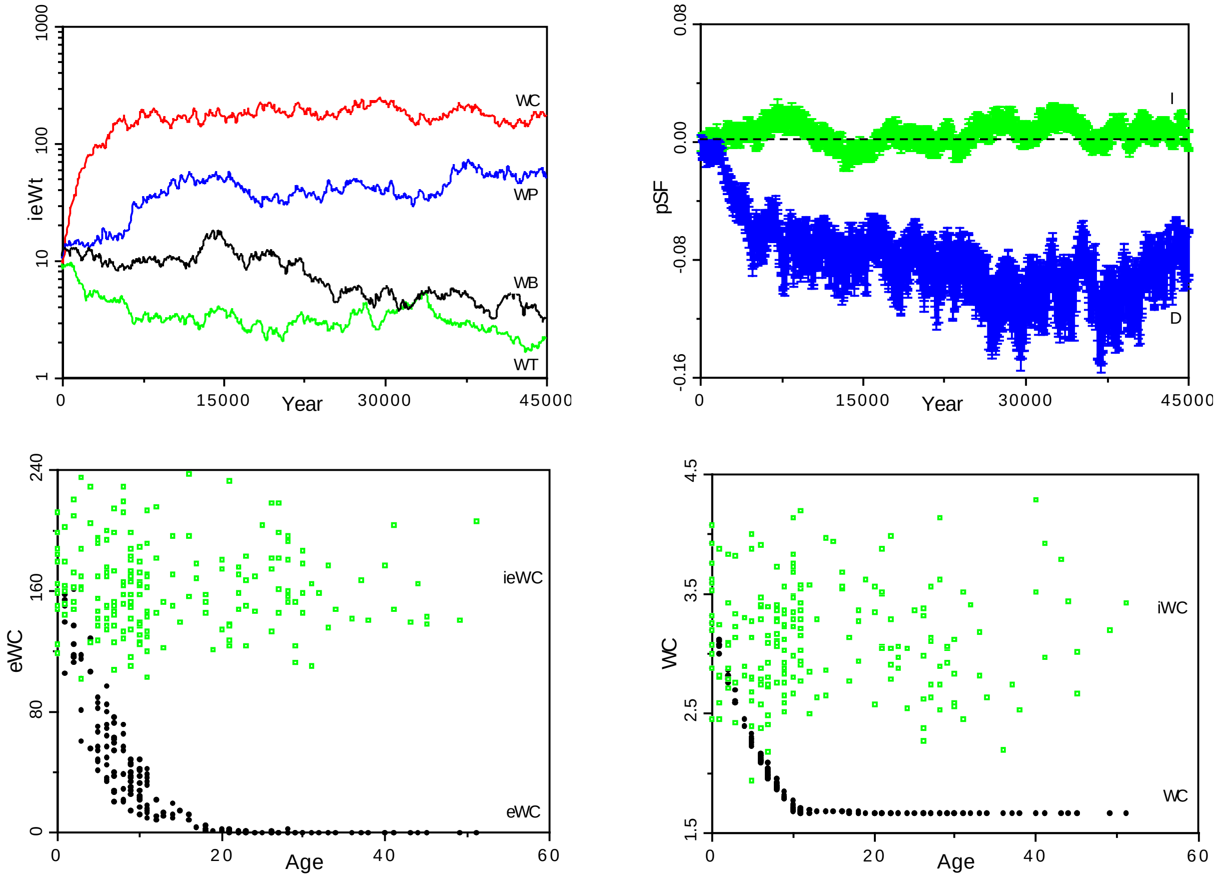

Figure 3: Evolution of the basic system with adaptable plasticity variation.

performance after a noisy procreation process and/or being born into an unpredictable environment. In many cases it will also need plasticity to compensate for the changes that naturally take place during its maturation period. For this study, such a maturation process was simulated by a simple output scale factor that varies linearly from 0.5 to 1.0 over the first ten years. In humans this might correspond to changes in inter pupil distance for oculomotor control, or changes in arm length for reaching or pointing. The important consequence is that the appropriate innate/newborn weights will not be the same as the adult values. The plasticities that evolve will allow the system to learn most efficiently how to change the weights throughout its life.

Unfortunately, limited computational resources allowed only a rather coarse simulation of the evolutionary process, but it proved sufficient for our purposes to have a fixed population size of only 200, with around 12 deaths per year due to competition, and 4 individuals over 30 years old dying each year due to old age. The procreation and mutation parameters were chosen to speed the evolution as much as possible without introducing too much noise into the process. All these details were kept constant across all the simulations.

Figure 2 shows the evolution of the basic model with deterministic plasticity variation. As I have also found elsewhere [3], the system quickly evolves appropriate initial weights (*iWt*) and initial learning rates (*ieWt*), despite there being no direct inheritance of learned behaviours. The scatter plot of the weights *WC* against *Age* for the evolved population shows how the plasticity allows it to change from the innate newborn values *iWC* to the appropriate adult values. The big question for this study is: how does the evolved population vary its plasticity with age? The scatter plot of the learning rates *eWC* against *Age* shows a clear decline in plasticity with age. The plot of the initial learning rates *ieWC* against *Age* indicates that this is not simply a matter of the low *ieWC* individuals living longer. The graphs of the plasticity variation parameters *pAGE* and *pSF* against *Year* confirm that the system does evolve an exponential decay of plasticity that starts from an early age (1 or 2 years old) with a scale factor 1+*pSF* of around 0.96.

Figure 3 shows that a remarkably similar population evolves for our adaptable plasticity variation version. The increasing plasticity scale factor 1+*pSFI* applied after consistent weight changes differs little from 1, whilst the decreasing

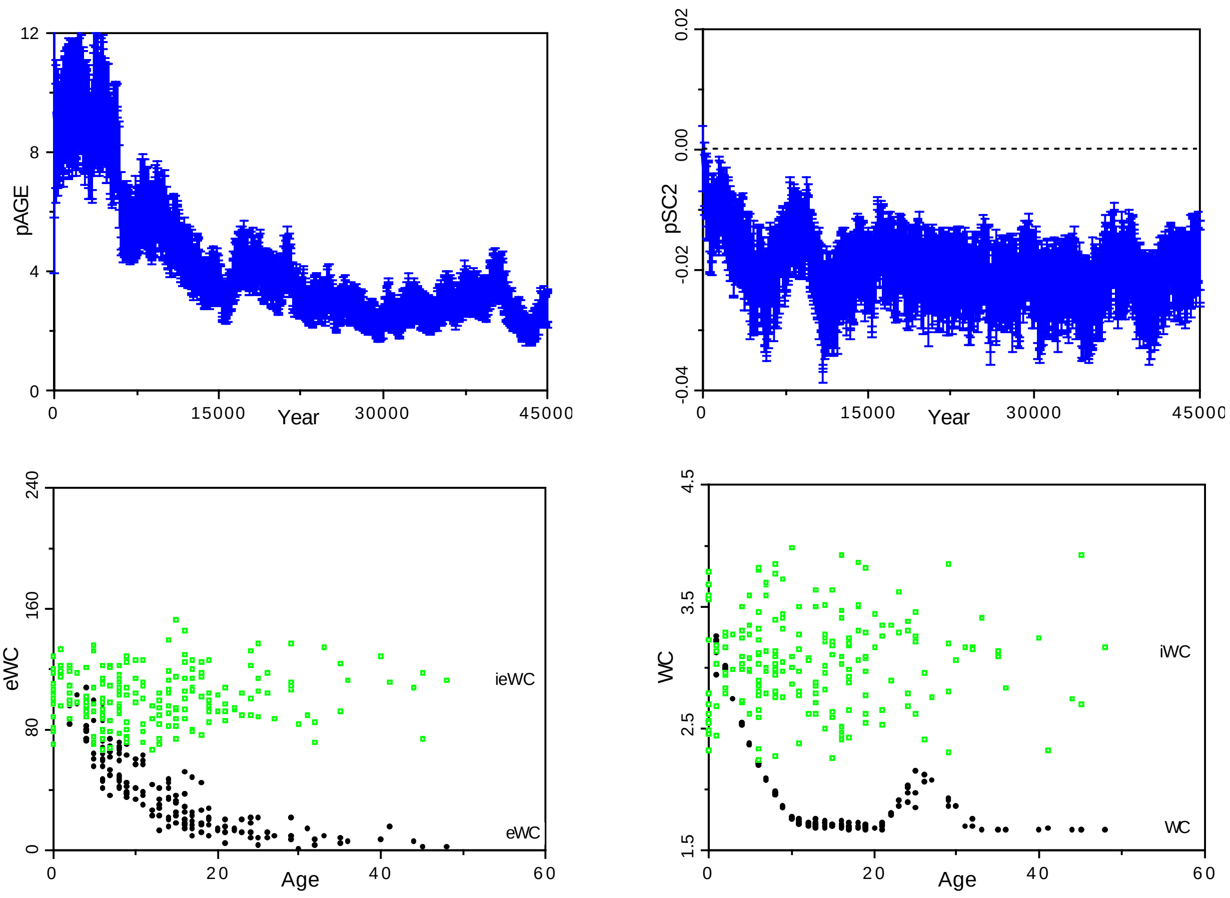

Figure 4: Evolution of a system requiring later life adaptation with deterministic plasticity variation.

scale factor 1+*pSFD* applied after inconsistent weight changes is around 0.92. The plots of *eWC* and *WC* against *Age* indicate that these scale factors have a similar effect to the more frequently applied deterministic scale factor observed above.

The obvious question arises as to what happens if an individual needs to adapt/learn later in life, after the standard maturational period is over. It seems unlikely that evolution will allow the plasticities to decay away to very small values in this case. To test this we make our output scale factor decrease linearly from 1.0 to 0.75 between the ages of 20 and 25 and return linearly back to 1.0 between 25 and 30. There is no need to specify whether this corresponds to an internal factor (e.g. compensation for system damage or deterioration) or an external factor (e.g. adaptation to changes in the operating environment), as they will have the same effect. Obviously, real late life adaptation will rarely be so predictable, but the consequences for our model will be similar, and the simplification makes it easier to interpret the results.

Figure 4 shows how this need for additional adaptability affects the deterministic plasticity variation model. We see that the initial learning rates are lower, the decrease in plasticity is slower, and the decrease starts at a later age. The net effect of all this is that there is enough plasticity remaining to bring about the necessary weight adaptations between the ages of 20 and 30, but not too much at any other age.

Figure 5 shows how the adaptable plasticity variation case is affected by the need for later life adaptation. Here there appears to be little reduction in the initial learning rates, but the plasticity varying scale factors are modified to result in appropriate plasticity during the late life adaptation period.

The important point to note from both these cases is that, even when late life adaptation is required, the systems still evolve plasticities that have large innate values that then decrease during the maturation period.

5 Conclusions

By simulating evolving populations of simple adaptable neural network control systems we have seen that there is a natural propensity for the evolution of plasticities that decrease with age, quite independently of any physical overheads of the plasticity. This is consistent with the well known "critical periods" of brain development [4].

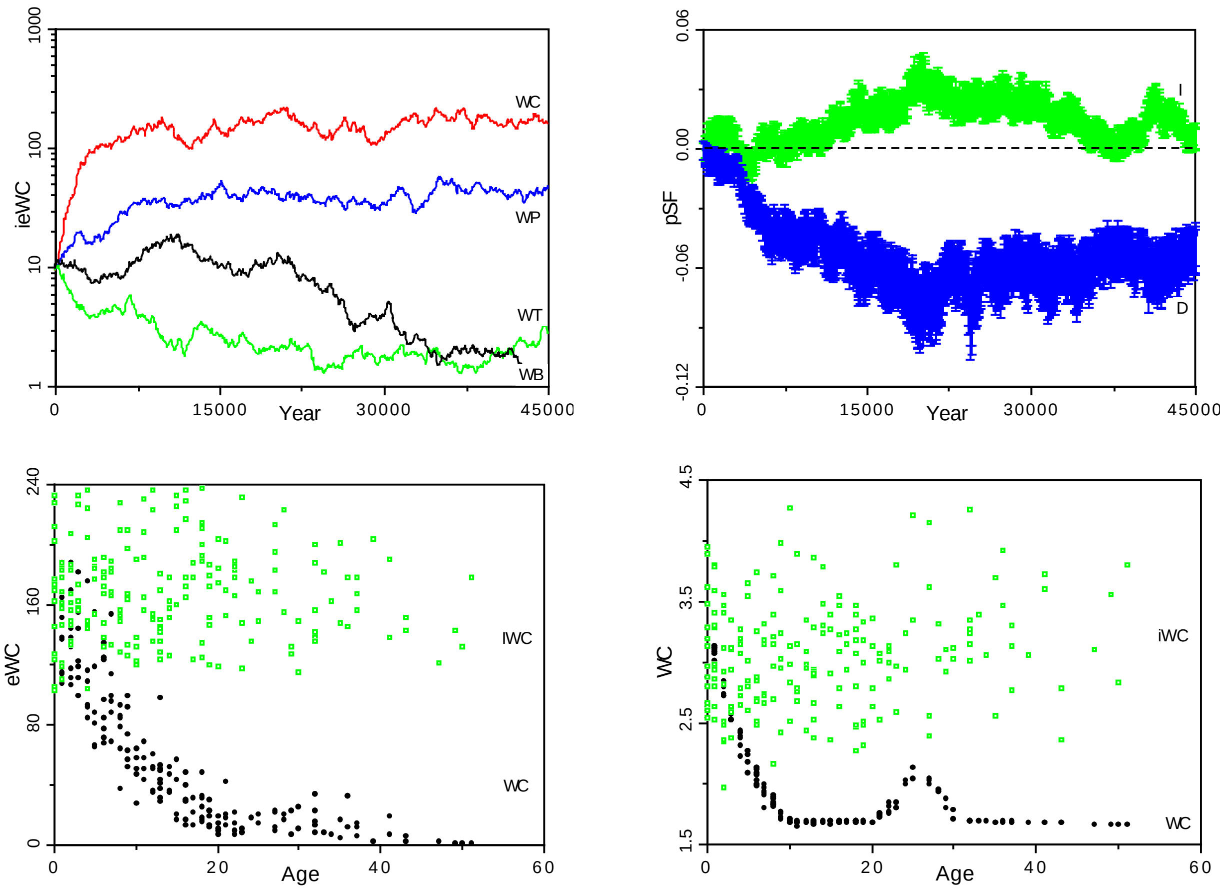

Figure 5: Evolution of a system requiring later life adaptation with adaptable plasticity variation.

There are two competing effects at play. In order to survive in competition with fitter adults and/or a hostile environment, a newborn needs to adapt as quickly as possible to its environment. It also needs to adapt efficiently to its own maturation. Large plasticities will be beneficial for both. In adults, however, large plasticities can lead to an unstable learning system, in which unusual/extreme experiences can potentially result in a large shift of the systems' parameters and a serious reduction in overall fitness. Lower learning rates here allow smoother optimal parameter estimation and more consistently good responses in a varied environment. In this paper it has been demonstrated how a process of evolution by survival of the fittest can result in a population of individual systems that deal with these conflicting requirements by having plasticities that decrease appropriately with age. Moreover, the process will provide a pattern of decrease that is appropriate for its environment. For example, a system that requires adaptation later in life will evolve differently to one that does not.

In complex systems, such as the human brain, we can expect each of the various sub-systems to evolve appropriately for its own requirements, so there may well be no single global behaviour. It is also likely that they will be able to evolve more complicated plasticity variability than was possible in our simple two parameter simulations. The next stage of this work will be to develop and test more realistic simulations of specific human sub-systems, and to explore how these ideas could be applied to the formulation of efficient artificial adaptable systems for real world engineering applications.

References:

[1] Baldwin, J.M., A New Factor in Evolution. *The American Naturalist*, **30**, 1896, 441-451.

[2] Belew, R.K. & Mitchell, M. (Eds), *Adaptive Individuals in Evolving Populations*. Addison-Wesley, 1996.

[3] Bullinaria, J.A., Exploring the Baldwin Effect in Evolving Adaptable Control Systems. *Proceedings of the Sixth Neural Computation and Psychology Workshop*. Springer, 2000.

[4] Greenough, W.T., Black, J.E. & Wallace, C.S., Experience and Brain Development. *Child Development*, **58**, 1987, 539-559.

[5] Jacobs, R.A., Increased Rates Of Convergence Through Learning Rate Adaptation. *Neural Networks*, **1**, 1988, 295-307.