

---

# Need is All You Need: Homeostatic Neural Networks Adapt to Concept Shift

---

**Kingson Man**  
University of Southern California  
kman@usc.edu

**Antonio Damasio**  
University of Southern California  
damasio@usc.edu

**Hartmut Neven**  
Google Inc.  
neven@google.com

## Abstract

In living organisms, homeostasis is the natural regulation of internal states aimed at maintaining conditions compatible with life. Here, we introduce an artificial neural network that incorporates some homeostatic features. Its own computing substrate is placed in a needful and vulnerable relation to the very objects over which it computes. For example, a network classifying MNIST digits may receive excitatory or inhibitory effects from the digits, which alter the network's own learning rate. Accurate recognition is desirable to the agent itself because it guides decisions to up- or down-regulate its vulnerable internal states and functionality. Counterintuitively, the addition of vulnerability to a learner confers benefits under certain conditions. Homeostatic design confers increased adaptability under concept shift, in which the relationships between labels and data change over time, and the greatest advantages are obtained under the highest rates of shift. Homeostatic learners are also superior under second-order shift, or environments with dynamically changing rates of concept shift. Our homeostatic design exposes the artificial neural network's thinking machinery to the consequences of its own "thoughts", illustrating the advantage of putting one's own "skin in the game" to improve fluid intelligence.

## 1 Introduction

To paraphrase Heraclitus, "The only constant in life is change". This observation goes a long way towards explaining why the performance of trained systems often degrades over time. The real world is non-stationary; the rules and relationships learned today may no longer hold tomorrow. Un-learning the bad old rules, re-learning the good new ones, and knowing how to tell the difference remains a major challenge for learning machines. Here we are inspired by the natural intelligence of living organisms, which maintain themselves in the face of environmental change by following the dictates of homeostasis.

Homeostasis is the regulation of internal body states within a range compatible with life. In organisms capable of forming mental states, feelings are the mental expression of these internal viability states (see Carvalho and Damasio 2021; Damasio and Carvalho 2013). It has been proposed that (a) machines that implement a process resembling homeostasis could be designed to exhibit a feeling-like device for the motivation and evaluation of their behavior and that (b) equipping an artificial learner with a feeling-like device might improve its adaptiveness to the inconstant data streams of the real world (Man and Damasio 2019).

Our approach diverges from traditional conceptions of artificial intelligence that emphasize outward-directed perception and abstract problem solving. We regard high-level cognition as an outgrowth of resources that originated to solve the ancient biological problem of homeostasis. A body is subject to perennial risk and decay. Organisms do not take a neutral or value-free stance towards information processing. Their actions have consequences for their own well-being. In our perspective, sense-data become meaningful when the data are connected to the maintenance and integrity of the sensing system.

Here we present a homeostatic neural network architecture based on a sketch by (Man and Damasio 2020) in which a classifier is placed into a needful and vulnerable relation to the objects over which it computes. By way of analogy, the homeostat must learn to distinguish between cups of coffee and cups of beer, while also needing to take a drink every so often to regulate its own mental arousal. In this setting, accurate classification is desirable to the agent itself because it guides decisions that can carry consequences for its internal states.

## 2 Background

We begin by providing some background on related work in biologically inspired mechanisms of self-regulation in neural networks.

### 2.1 Homeostatic features in neural networks

In biological brains, neurons must regulate their intrinsic excitability and synaptic conductance to ensure stable network function (Marder and Goaillard 2006). In artificial neural networks, and especially in those containing recurrent connections, homeostatic regulation of excitability can reduce saturation and improve signal propagation (Williams and Noble 2007). In simulation experiments of evolutionary robotics, which made explicit reference to homeostasis, phototactic robots were controlled by a fully connected neural network and used ‘neural plasticity’ to restore adaptive behavior following visual field inversion (Di Paolo 2000; Iizuka and Di Paolo 2008).

We note, however, that the homeostatic-like features of prior work, where present, were implemented from the outside-in: the systems were instructed to maximize, or keep within a set range, certain arbitrary values that were labeled "homeostatic". There was no obligatory link between external conditions and internal parameters. The operation of the system itself was not exposed to the consequences of the system’s own activities, that is, it was not made vulnerable. At best, a simulated population was culled by the designer’s fitness function, a fact that was not appreciated by the simulated agents themselves.

### 2.2 The problem of concept shift

A major challenge in machine learning is dataset shift, in which learners fail to generalize because the training and testing data are drawn from different distributions. A severe variant of dataset shift is concept shift (we follow the terminology of Moreno-Torres et al 2012; see also Widmer and Kubat 1996; Tsymbal 2004; Kull and Flach 2014), in which the relationships between labels and attributes can change over time.

Under concept shift, the associations between labels  $y$  and observations  $x$  change across the training and testing phases:  $P_{train}(y|x) \neq P_{test}(y|x)$ . (The observations, however, are drawn from the same distribution across training and testing:  $P_{train}(x) = P_{test}(x)$ )

This type of phenomenon occurs frequently in real world settings of online supervised learning – especially in social settings. As rapidly as fashions and trends can change, recommender systems must stay current and serve the same user’s new tastes. As another example, at the onset of a global pandemic, an entire demographic previously well-served by travel ads may suddenly be better served by ads for home furnishings.

We may also compare the problem of concept shift to the problem of lifelong learning. In lifelong learning a system must sequentially learn to perform multiple different tasks (see Thrun 1998; Chen and Liu 2018). New tasks often interfere with earlier tasks, leading to poor performance or even "catastrophic forgetting". One biologically inspired method around this problem reduced the learning rate of neurons that were found to be important in learning previous tasks, leaving the other neurons

```

1 // Algorithm 1. Homeostatic self-regulation of learning rate.
2
3 // Helper function to estimate the accuracy of a simulated version of
4 // oneself on recent data, at a given learning rate
5 Define simulate_self(mlp, memories, lr):
6     Perform stochastic gradient descent, at the given lr, over a single
7     pass of the memories dataset
8     Return overall accuracy of predictions in re-classifying the
9     memories
10
11 // Begin algorithm
12 Input: data, a set of images and associated labels from {0..9}
13 Input: f, a frequency of lr regulation every f steps, e.g. 100
14 Input: initial_lr, an initial learning rate, e.g. 0.005
15
16 lr = initial_lr
17 lr_stepsize = lr/10
18 memories = [] // to store up to f preceding datapoints
19
20 Initialize a multi-layer perceptron, mlp.
21 timestep=0
22 While 1:
23     For each of (image, label) in data:
24         Append (image, label) to memories
25         if memories contains more than f items:
26             remove the oldest item
27         Perform forward pass on mlp and output a label prediction y^
28         if timestep is a multiple of f: // perform LR regulation
29             // What is the predicted effect of this object on lr?
30             if the label y^ is greater than or equal to 5:
31                 simulated_lr = lr + lr_stepsize // excitatory
32             else:
33                 simulated_lr = lr - lr_stepsize // inhibitory
34             // What would the accuracy be if learner altered its lr?
35             ingest_accuracy = simulate_self(mlp, memories, simulated_lr)
36             // What if learner kept its current lr?
37             reject_accuracy = simulate_self(mlp, memories, lr)
38             if ingest_accuracy > reject_accuracy:
39                 lr = simulated_lr
40         Perform backward pass on the mlp at the selected lr
41         timestep += 1

```

Figure 1: Algorithm: Pseudocode for homeostatic self-regulation of learning rate.

free to pick up new tasks (Kirkpatrick et al 2017). Under concept shift, on the other hand, some forgetting is not catastrophic but rather desirable and necessary.

### 3 Homeostatic architecture of needful neural networks

One way of explaining our homeostatic design is to say that it exposes an artificial neural network's thinking machinery to the consequences of its own "thoughts". This incentivizes the network to regulate itself to think better – to become better aligned with reality and to adapt to external change.

#### 3.1 Classifications with consequences

A homeostatic learner is given the task of classifying images of objects. In a twist, the learner is designed to be needful – it depends on the objects that it classifies for its continued integrity and functionality. The objects can have direct effects, excitatory or inhibitory, on the learner itself. For example, in a classification of MNIST digits, the digits {0,1,2,3,4} have inhibitory effects and reduce the neural network's learning rate, while the digits {5,6,7,8,9} have excitatory effects and increase the neural network's learning rate.

Critically, following classification of an object, the learner decides to either take it or leave it, that is, it can choose to either "ingest" or reject the recognized digit, and therefore to regulate its own learning rate. The agent must answer a question akin to "How does this object make me feel?" We use a counterfactual decision process to answer a related question, "How would my own functionality be affected by taking or leaving this object?" This introduces a homeostatic meta-task to secure the internal conditions under which the learner can perform the ostensible task of image classification.

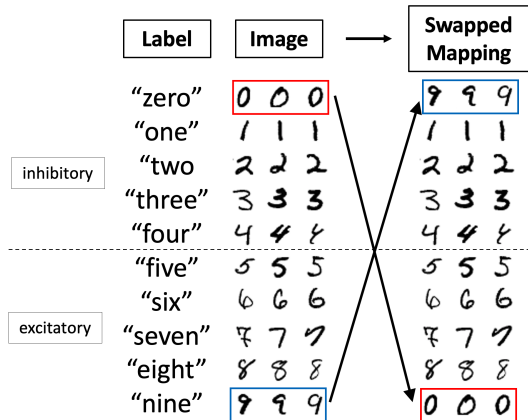


Figure 2: Concept shift is implemented by swapping the mapping between label and image between two randomly selected classes. Illustrated here on the MNIST dataset, the mappings for “zero” and “nine” are swapped. This swap will also invert the homeostatic effects expected for each number – a potentially hazardous situation for a vulnerable classifier.

The learner assesses the desirability of each alternative by evaluating each version of itself against a store of recently seen objects and their labels (Fig. 1).

With its own skin in the game, the homeostatic learner naturally chooses the course of action that best improves its classification ability. This arrangement sets up a "strange loop" (Hofstadter 2007) of causality in which accurate classification is desirable to the classifier itself. Mis-identification of an object can lead to mis-prediction of the object’s effects on the learner. The learner is therefore incentivized to improve performance, especially during times of great change – and great peril.

### 3.2 The vicissitudes of life

As so often happens in life, the rules have a way of changing on you. We introduce concept shift to the dataset by permuting labels on a subset of the data. When a shift occurs, we randomly select two of the 10 classes and swap their labels for all instances in the data. For example, in the MNIST dataset, we may swap the labels for "zero" and "nine", such that all images of a large vertical oval (0) are now labelled "nine", and all images of a small circle with a line hanging off of its right (9) are now labelled "zero" (Fig. 2). Note that in this swap, the homeostatic effects of the digits have been inverted. Images that were previously associated with inhibitory effects (image 0 → label "zero" → inhibitory) are now excitatory (image 0 → label "nine" → excitatory), and vice versa.

## 4 Experiments

We perform empirical studies to compare homeostatic regulation of learning rate (LR) against two control conditions: a randomly regulated “wandering” learning rate, and a more conventional, constant learning rate. We characterize the conditions under which homeostatic regulation either imposes an overhead cost and a performance penalty, or else allows a learner to smoothly adapt to changing conditions.

All classification studies are performed with a multilayer perceptron with two hidden layers containing 80 and 60 units respectively, using the ELU activation function (Clevert et al 2015) and He initialization (He et al 2015). We evaluate our method on two datasets, MNIST (Lecun et al 2010) and Fashion-MNIST (Xiao et al 2017). All experiments were performed in MATLAB and the source code is provided in the Appendix.

### 4.1 The homeostatic learner adapts to concept shift

Testing across a wide range of rates of concept shift, measured in swaps performed per epoch of training, we find that in the stationary setting (no swapping) the conventional, constant-LR classifier




Figure 3: Homeostatic learners incur some performance penalty in environments with no or low concept shift, but are far superior under conditions of highest shift. Color-coded validation accuracies of learners with their learning rates homeostatically regulated (blue), randomly regulated (green), and held constant (red). Traces show mean  $\pm$  SEM over 20 replicates.

is most accurate (Fig. 3, red traces in left columns). However, the homeostatic learner is only marginally less accurate than the constant-LR learner up to about 10 swaps per epoch (Fig. 3, blue traces). The ability of the homeostat to nearly match the constant-LR classifier’s performance is remarkable because the homeostat starts with a seeming disadvantage: it is vulnerable to its own mistakes. The homeostat’s incorrect classifications can compound into poor decisions that impair future performance. Illustrating how badly things could have gone for the homeostat, the randomly regulating LR classifier (green traces) quickly goes off the rails and shows large variance across replicates.

The benefits of the homeostatic architecture become apparent at the highest intensities of concept shift (Fig. 3, right columns). At the extreme of 500 swaps per epoch – in which the learner experiences a digit/label swap every 100 digit presentations – the constant-LR classifier is swamped by change and falls to near chance level. The homeostat, on the other hand, continues to learn and improve despite the onslaught of environmental change. The contrast between homeostatic and static LRs, under high concept shift, is even more stark when omitting the high-variance traces of the random learner (Fig. 4).

For another view into the data, we group accuracies by learner type, illustrating the declines in performance with increasing rates of concept shift (Fig. 5). The homeostat shows more graceful declines in comparison to the constant-LR learner.

#### 4.2 Homeostatic LR regulation is responsive to the prevailing rates of concept shift

The homeostat tunes its learning rate to a level specific to the environment in which it finds itself (Fig. 6). All learning rates are initialized at the same value (0.005) but the homeostat seeks a level appropriate to the experienced rate of concept shift. For the random LR regulators, on the other hand, the sequence of learning rates is insensitive to shift rates and shows a lower-bounded random walk that drifts upward.




Figure 4: Keeping the same conventions as in Fig. 3, omitting the random-LR learner for clarity. Left, in the stationary environment the homeostat performs slightly below the constant-LR learner. Right, the homeostat is far more accurate than the constant-LR learner under severe concept shift. Note the difference in ordinate scales across charts.

### 4.3 The homeostatic learner adapts to second-order shifts, or "seasonality"

Encouraged by the ability of the homeostat to seek out a learning rate appropriate to the prevailing rate of concept shift, we created learning environments with seasonality, in which the rate of concept shift can vary over the course of training. We find that the homeostat maintains the best average performance across "calm" and "stormy" seasons, and rapidly recovers after the onset of a stormy period (Fig. 7).

This pattern may be explained by the homeostat's LR sequences under seasonal training. The homeostat rapidly ramps up LR during stormy periods but tends to hang on to these gains; it does not lower LR during the return to a calm period (Fig. 7, bottom). This refractory behavior of persistently elevated LR may lead to relatively poor performance in the calm periods following seasonal cycles (e.g., the blue traces in the final 20 epochs).

## 5 Discussion

To summarize, we show that: 1) homeostatic learners are superior to conventional learners under concept shift, with the greatest advantage obtained under the greatest rates of shift; 2) homeostatic regulation imposes a slight performance penalty under static and low-shift environments; 3) homeostatic learners tune their learning rate in accordance with environmental conditions; and 4) homeostatic learners can adapt to second-order shift, or changes in the rate of environmental change. Although we find these converging results across the MNIST and Fashion-MNIST datasets we note that one possible limitation on the scope of our claims is the use of only these two datasets, each being somewhat limited in visual complexity and image size.

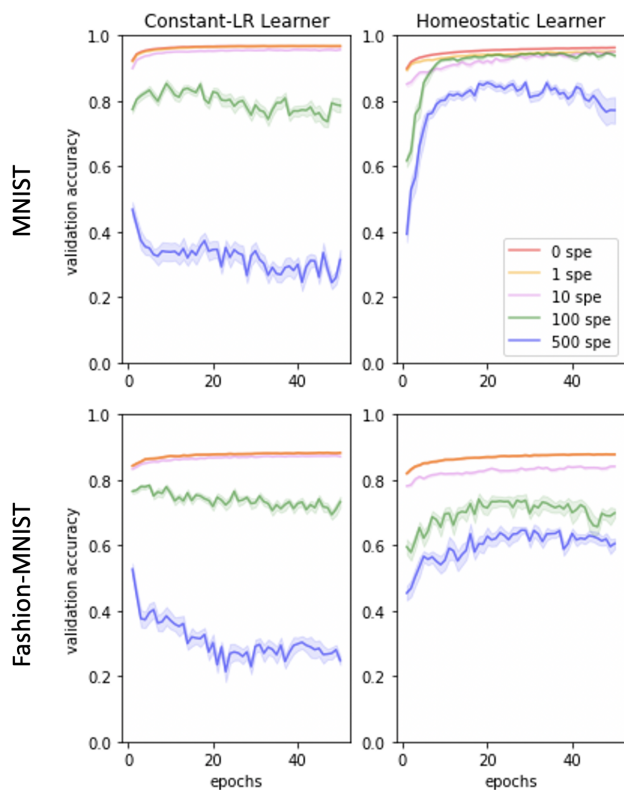


Figure 5: Accuracies by classifier type. Left, At high rates of concept shift – the green and blue traces – there is a large drop in accuracy for the constant-LR classifier. Right, The homeostat exhibits better adaptation to concept shift, with less marked declines of the green and blue traces.

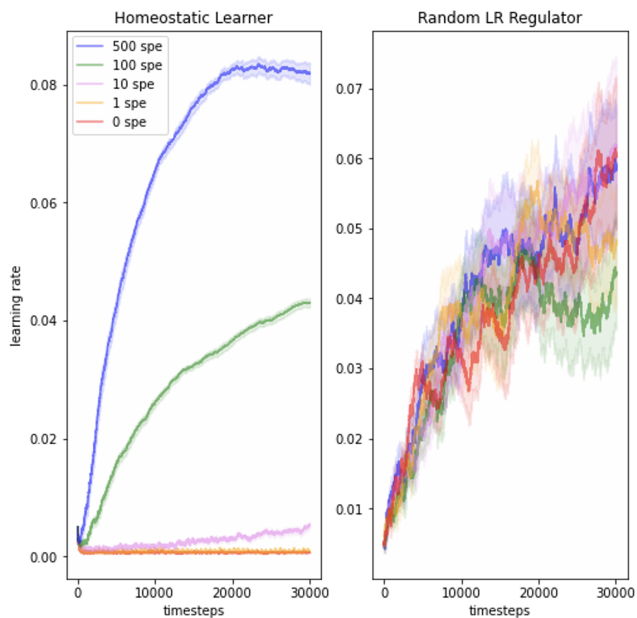


Figure 6: Learning rate sequences of the two LR-regulating classifiers. The homeostatic learner seeks an LR appropriate to each level of concept shift, while the random regulator drifted upwards. At 500 swaps per epoch (left, blue), the homeostat arrests its own LR growth and asymptotes. Data shown from MNIST only.




Figure 7: Accuracy and learning rate under “seasonality” of concept shift. Top row: Schedule A cycles between extreme rates of concept shift, while schedule B is more gradual. Middle row: The homeostat maintains good average performance across shifts in the rate of concept shift. The accuracy of the constant-LR classifier declines precipitously during stormy periods but returns to normal during calm periods. Bottom row: The sequence of learning rates reveals that the homeostatic learner ratchets up its learning rate during stormy periods but is far less inclined to reduce it during calm periods.

Another possible limitation is the re-use of training data over many epochs, which limits the fundamental novelty of the concept shift. Although the labels and data are repeatedly shuffled, the classifier is never asked to learn from never-before-seen image patterns. In the real world, concept shift often co-occurs with some level of covariate shift. Not only do relationships change over time, but the predictors change as well.

Although our method can dynamically adjust the learning rate, we did not benchmark it against LR optimizers such as ADAM (Kingma and Ba 2014) and other momentum-based methods. In the non-stationary setting the loss surface shifts over time and it is therefore inappropriate to accumulate previous gradients from an outdated loss surface. However, see (Kobayashi 2021) for a momentum-based optimizer in the non-stationary setting, showing some mixed results compared to ADAM.

Finally, we are aware of the resemblance between reinforcement learning and our task of homeostatic self-regulation, though we argue that they should not be identified as the same. The objective here is not to maximize some arbitrary “reward” by massed trial-and-error. The object of the game is simply to keep playing the game. We specify a particular target for a learner to optimize: its homeostatic well-being, crystallized as an internal parameter that controls its ongoing ability to make good decisions, which affect an internal parameter that controls its ongoing ability to make good decisions. . . and so on. For an excellent example of work in reinforcement learning that does take homeostatic logic into account see (Keramati and Gutkin 2014).



We close with some ethical considerations of our work. Robots optimized for self-preservation – what could possibly go wrong? The proposal of (Man and Damasio 2019) considers a solution to the possible problems of imbuing machines with feelings of self-interest: provide them with more feelings, so that they can comprehend the feelings of self-interest of others, in other words, so that they can have empathy. Considering present reality, our work has the potential to improve recommender systems such that they can be exquisitely attuned to our changes in moods, tastes, and fashions. These could clearly generate economic benefits, but there is also the potential for such skilled systems to insult the human experience of free choice and self-determination. A person may desire to “escape the algorithm” and try to zig instead of zag – only to quickly end up back in the clutches of their attentive machines.

## Acknowledgments and Disclosure of Funding

We gratefully acknowledge funding to KM and AD from Google and from the National Science Foundation under Grant No. IIS-2033961.

## References

- Carvalho GB, Damasio A. Interoception and the origin of feelings: A new synthesis. *BioEssays*. 2021;2000261.
- Chen Z, Liu B. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*. 2018;12(3):1-207.
- Clevert DA, Unterthiner T, Hochreiter S. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*. 2015.
- Damasio A, Carvalho GB. The nature of feelings: evolutionary and neurobiological origins. *Nature Reviews Neuroscience*. 2013;14(2):143-52.
- Di Paolo EA. Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. 2000.
- He K, Zhang X, Ren S, Sun J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE international conference on computer vision 2015* (pp. 1026-1034).
- Hofstadter DR. *I am a strange loop*. Basic books; 2007.
- Iizuka H, Di Paolo EA. Extended homeostatic adaptation: Improving the link between internal and behavioural stability. In: *International Conference on Simulation of Adaptive Behavior 2008 Jul 7* (pp. 1-11).
- Keramati M, Gutkin B. Homeostatic reinforcement learning for integrating reward collection and physiological stability. *eLife*. 2014;3:e04811.
- Kingma DP, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. 2014.
- Kirkpatrick J, Pascanu R, Rabinowitz N, Veness J, Desjardins G, Rusu AA, Milan K, Quan J, Ramalho T, Grabska-Barwinska A, Hassabis D. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*. 2017;114(13):3521-6.
- Kobayashi T. Towards deep robot learning with optimizer applicable to non-stationary problems. In: *2021 IEEE/SICE International Symposium on System Integration (SII) 2021* (pp. 190-194).
- Kull M, Flach P. Patterns of dataset shift. In: *First International Workshop on Learning over Multiple Contexts (LMCE) at ECML-PKDD 2014*.
- LeCun Y, Cortes C, Burges CJC. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1994.
- Man K, Damasio A. Homeostasis and soft robotics in the design of feeling machines. *Nature Machine Intelligence*. 2019;1(10):446-52.
- Man K, Damasio A. Truth or Consequences: Homeostatic Self-Regulation in Artificial Neural Networks. In: *Artificial Life Conference Proceedings 2020* (pp. 146-147).
- Marder E, Goaillard JM. Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience*. 2006;7(7):563-74.
- Moreno-Torres JG, Raeder T, Alaiz-Rodríguez R, Chawla NV, Herrera F. A unifying view on dataset shift in classification. *Pattern Recognition*. 2012;45(1):521-30.

- Thrun S. Lifelong learning algorithms. In: Learning to learn 1998 (pp. 181-209).
- Tsymbol A. The problem of concept drift: definitions and related work. Computer Science Department, Trinity College Dublin. 2004;106(2):58.
- Widmer G, Kubat M. Learning in the presence of concept drift and hidden contexts. Machine Learning. 1996;23(1):69-101.
- Williams H, Noble J. Homeostatic plasticity improves signal propagation in continuous-time recurrent neural networks. Biosystems. 2007;87(2-3):252-9.
- Xiao H, Rasul K, Vollgraf R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747. 2017.