

# Do the Origins of Biological General Intelligence Lie in an Adaptation of the Stress Response?

John Cummins

Cobblestones, Church Road, Dodleston, Chester CH4 9NG. UK  
deva\_research@lineone.net

**Abstract.** Research into natural and artificial intelligence can overlook that natural intelligence capable of anticipating the future has a potential cost, mediated by cognitive factors involved in the stress response, leading to high allostatic load ('wear and tear'). This theoretical paper suggests that nature may have partly resolved the problem by using the same mechanism—an adaptation of the stress response enabling a new and flexible balance of accurate and inaccurate assessments of the animal's control over the environment—to (a) generate flexible, high-level general intelligence in humans and (b) reduce allostatic load to within adaptive parameters. This new form of intelligence, probably appearing in early hominins, acts as a buffer between the animal and its environment. A tentative framework for information processing around the primate brain is proposed, showing where and how such 'buffer intelligence' could have arisen and been exploited in early hominins. This appears to be a development of a function undertaken in non-human primates by the neural correlates of consciousness, an area of the non-human primate brain where there is no, or very little, intelligence. There is a brief discussion of whether this principle might enable the spread of a capacity for intelligence throughout a complex adaptive system, with flexible linguistic syntax in humans as an example.

**Keywords:** Stress, Cost of intelligence, Allostatic load, Assessment of control, Chaos, Buffer intelligence, Consciousness, Language, Syntax.


## 1 Introduction

This theoretical paper offers a possible signpost towards understanding the origins and structure of flexible, high-level general intelligence in humans.

The author is indebted to many people for comments and advice, including the late Jeffrey Gray, Leslie Aiello, Igor Aleksander, Peter Dayan, Harvey Dearden, Wlodek Duch, Karl Friston, Robert Lowe, Bruce McEwen, Neil McNaughton, Robert Sapolsky, Colleen Schaffner, Murray Shanahan, Craig Smith, Derek Smith, Tom Ziemke and Pei Wang.

The ultimate goal of any living organism is to preserve its structural integrity long enough to promote its genetic material successfully into the environment. Processes aimed solely at achieving proximate goals such as survival, nutrition, reproduction and shelter, etc. occur within the overarching context of this ultimate evolutionary

goal. Proximate goals (including those pursued by ‘intelligence’), although contributing to the ultimate goal, are not necessarily identical to the ultimate goal. Achieving the ultimate goal involves factors that may affect how different proximate processes operate and are integrated with each other. The result may be a complex balance of costs and benefits. This paper examines the balance of costs and benefits in the relationship between intelligence, the stress response and allostatic load.



**Fig. 1. Potential Problem** In early hominins, the ‘weight’ of increased anticipatory intelligence could cause increased allostatic load that impacts adversely on the ultimate goal of reproductive success. The result is a non-adaptive imbalance between proximate and ultimate goals

Research into natural and artificial intelligence can overlook that natural intelligence capable of anticipating the future has a potential ‘downstream’ cost, mediated by the stress response, in the form of high allostatic load. In brief, allostatic load is the damage caused to the animal by over activation of the stress system. As intelligence increases, an animal lives less ‘in the moment’ and can prospect about future outcomes, leading to cognitively-mediated stress and high allostatic load. This can

impact adversely on health, longevity and reproductive fitness, leading in turn to an evolutionary bottleneck, as illustrated in Fig.1.

Nature may have partly resolved the problem by varying the primate stress system. The argument will be that a random genetic change raised a threshold in the brain of late –state activation of the physiological stress response. Feedback from this raised threshold created a flexible and adaptive balance of accurate and inaccurate assessments of the animal’s control of its environment. This same mechanism is (a) the origin of flexible, high-level human intelligence and (b) responsible for reducing the potential cost of stress response-mediated high allostatic load to within adaptive parameters. Human intelligence is not optimised to make accurate assessments of control in pursuit of proximate goals but to contribute towards the ultimate goal. This adaptive balance of accurate and inaccurate assessments of control acts as a buffer between the animal and its environment. Section 2 discusses the possibility of introducing purposeful flexibility into a brain, Section 3 expands on the evolutionary reasons why this might have occurred, and Section 4 proposes a framework for information processing around the primate brain, showing how such ‘buffer intelligence’ could have arisen in early hominins.

## **2 Accuracy of Control and the Generation of Purposeful Flexibility**

Accuracy of assessments of control in the primate stress system may offer a way to understand flexible high-level general intelligence, defined as an ability to respond adaptively to challenges, whether novel or similar to ones previously experienced, and flexible anticipation and modelling of a wide range of future scenarios of varying likelihood and the ability to choose (or where necessary create) and execute adaptive responses thereto. The key word is ‘adaptive’, which can only be defined by reference to the ultimate goal, not preconceived ideas of accuracy or ‘efficiency’ in performance of proximate tasks.

The cognitive architectures reviewed by Duch and colleagues [1], the large scale brain simulations discussed by de Garis and colleagues [2], and the biologically inspired cognitive architectures discussed by Goertzel and colleagues [3] assume that achieving accuracy of assessments of control over the external environment is a wholly beneficial goal. Within artificial systems, accuracy of assessment of control may have to be partly traded for other benefits such as speed or economy of resources, but it remains an ideal.

A problem in artificial general intelligence is how to introduce purposeful flexibility into a system. The limitations of an artificial system are apparent when it encounters a novel situation. If it encounters an uncertain situation, a learning-based, Bayesian-inspired system [4] may operate in accordance with Bayesian statistics to infer the conditional probability of an event P occurring, given that Q has occurred, therefore allowing the system to respond optimally under known uncertainty. This may suffice in an environment with ‘known unknowns’, but if a genuinely novel situation arises such a system will not be able to optimise its response or learning about the

environment and, indeed, *may not survive the initial encounter*. Whether it does survive will depend partly on luck. This is because Bayesian predictions based on a model cannot create new arbitrarily complex models out of smaller units [5]. A learning-based system without flexibility will therefore always tend to gravitate towards the most likely response based on what has occurred previously; it will have difficulty generating a functional response to a novel environmental challenge, or modelling novel scenarios and responses internally.

The problem and some possible solutions are outlined in Fig.2. This shows a memory / learning-based, probabilistic system operating especially at the level of sensory-motor coordination, anticipating the future based on the past. For this limited system to respond creatively to novel situations and model novel scenarios and possible responses internally, it must be integrated with a capacity to generate purposeful flexibility if it is to be described as having flexible, high-level general intelligence. Possible sources of purposeful flexibility include logic, randomness, ‘chaotic itinerancy’, and the focus of this paper, a variation of the stress response that harnesses accurate and inaccurate assessments of control to generate purposeful flexibility.




Fig. 2. Possible Sources of Purposeful Flexibility

Regarding option a in Fig. 2, there is interest in a synthesis of logical programming and Bayesian networks in the field of probabilistic programming [5], but it is unclear if this approach can generate sufficient flexibility or exploration if these are not inherent in Bayesian probability or logic.

Option b might achieve purposeful flexibility by introducing random combinations of existing models already in the system, and then sifting for utility as discussed by Calvin [6]. But brain tissue is expensive and the environment is dangerous; reliance solely on a stochastic exploration of model spaces approach would be expensive, very slow (because of the combinatorics) and high risk. A promising approach is Friston’s recent synthesis (option c) aimed at providing a theoretical framework for

understanding the brain [7]. In Friston's model, 'chaotic itinerancy' provides some flexibility, riding on top of Bayesian processes (equating to the lower box in Fig. 2), all within an overriding 'free energy' reduction framework. Put simply, the system expects to explore in a Bayes optimal fashion. There is insufficient space to discuss this here. Instead, the focus in what follows is on option d, which involves an adaptation of the stress response.


### 3 The Stress Response, Allostatic Load and the Cost of Intelligence

The primate stress response system comprises an integrated complex of central and peripheral neural and neuroendocrine processes that enable the animal to respond adaptively to threats ('stressors'), especially by vigorous muscular activity. For a review see Boyce & Ellis [8]. Ellis *et al* characterize its main features as follows [9]: Environmental events signaling threats produce responses within the neural circuitry of the brain and peripheral neuroendocrine pathways regulating metabolic, immunologic, and other physiological functions. This causes a shift to a state of biological and behavioral preparedness, involving increases in heart rate and blood pressure, metabolic mobilization of nutrients, preferential redirection of energy resources and blood to the brain and to the external musculature, and the induction of vigilance and fear. The neural basis for the organism's stress response comprises two anatomically distinct but functionally integrated circuits: the corticotrophin releasing hormone system and the locus coeruleus–noradrenaline system [10, 11, 12]. Co-activation of these two systems, along with their linkages to emotion regulatory brain regions such as the amygdala, anterior cingulate cortex and prefrontal cortex, produces the coordinated biobehavioral changes associated with the stress response in mammalian species [9]. The primary stress response axes, as well as their central and peripheral components, have been extensively conserved in the evolutionary history of vertebrate and mammalian species [9, 13,14].

In contrast to homeostatic systems, such as blood sugar, blood pH, and body temperature, which must be maintained within a narrow range, stress responses are allostatic [15]. Allostasis is the ability to achieve stability, or homeostasis, through change, as defined by Sterling and Eyer [16]. As an allostatic system, the stress response (or one of its elements) is adapted to turn on in response to a threat and then turn off again when the threat has passed. However, the stress response is *biphasic* [17]: It is highly adapted for dealing with acute stressors, but excessive, repeated or chronic stressors, or a failure to shut down the process appropriately, can cause high allostatic load [15, 18]. Allostatic load is the damage caused to the animal by chronic over activity, under activity or dysregulation of allostatic systems [15]. High allostatic load increases vulnerability to a variety of diseases. Cardiovascular, immune and hippocampal-mediated memory systems are especially vulnerable to high allostatic load (19). The stress system itself can also be damaged by stress, leading to a vicious circle of dysregulation and systemic damage [19]. In sum, chronic stress in humans can lead to illnesses such as heart disease, diabetes, depression, anxiety, substance abuse, ga-

strointestinal disorders, eating disorders, auto immune disorders, infections, and possibly even tumours [20].

In humans, intelligence, stress and the potential cost of high allostatic load intersect in the arena of appraisal and anticipation. Arnold [21] and Lazarus [22, 23, 24, 25, 26], have shown that in humans some stimulus- response pathways, including elements of the stress response system, are mediated by cognitive<sup>1</sup> appraisal, as illustrated in Fig. 3.



**Fig. 3.** Cognitive Appraisal (after Lazarus)

Allostatic processes, including stress responses, can be turned on not just in response to a present challenge, but also in anticipation of something that is likely to upset homeostasis [19, p.10]. Thus, anticipation and worry can contribute to allostatic load [15].

According to the field of ‘emotion appraisal theory’ (which includes stress), in humans there are two main appraisals. A primary appraisal checks the valence of a situation (positive or negative) and magnitude. A secondary appraisal identifies the causal agent (‘self’ or other) and any resources the person might draw on to meet the challenge. Secondary appraisal is an assessment of the controllability of the situation. In the formula of Lazarus [22], a stress response can result if a person assesses that an important negative challenge exceeds coping resources. An appraisal that a negative situation is uncontrollable can turn on a stress response, and an appraisal that it is controllable can inhibit or turn it off. Control does not have to be exercised; a belief that you have control is sufficient [27]. Crucially, the belief does not have to be correct; an incorrect belief that one has control can suffice to inhibit a stress response [19, p.261].

We can now begin to construct a hypothesis about the relationship between evolutionarily increasing intelligence and the cost of allostatic load. Intelligence increases the ability to anticipate threats, appraise controllability, and respond adaptively, initially increasing adaptive fitness. However, even if anticipation and assessment of controllability are accurate, and increased intelligence enables anticipation and avoidance or resolution of problems, this does not mean overall stress is reduced, because however many problems are accurately anticipated and solved there will always be a ‘sink’ of problems that are accurately anticipated and assessed as being uncontrollable. Thus, intelligence can increase cognitively -mediated stress and allostatic load. There is no limit to the number of problems that can be accurately anticipated but not solved, or to the ensuing physical damage. (Errors can also occur at any stage in this process.) In a living system such as a human, with an anticipatory self-preservation

<sup>1</sup> It should be noted that the term ‘cognitive’ in this context does not imply that the process necessarily occurs entirely at conscious level [24].

system driven in part by highly developed flexible intelligence, the quality and duration of information inside the system can, without any other input from the outside, cause the system to destroy itself physically from the inside out.

We now consider how nature may have partly resolved the problem.

### 4 Proposed Framework

Fig. 4 shows a conjectural model of information flow around a non-human primate brain




Fig. 4.

In this simplified functional framework there are three main types of module (specialised function). There are many M1 and M2 modules, organised around inputs and outputs, respectively, perhaps using memory-based probabilistic processes and heuristics, including Bayesian or Bayes-like processes. To the extent that these modules do use memory-based probabilistic processes, they incorporate a model (or part thereof) of the animal's changing relationship to the world as 'priors'. M1 and M2 modules generate nearly all the intelligence and behaviour of non-human animals, but their limited flexibility means they have difficulty generating novelty.

The third type of module is a hypothesized 'Late-Stage Comparator' ('LSC'). Primates (and perhaps all mammals) have an LSC that integrates highly processed information from many brain areas ---sensory, motor, somatic (including homeostatic

and allostatic), proprioceptive, nociceptive, affective and memory. As the LSC is situated within a larger feedback loop incorporating the animal and its external environment it can track the relationship longitudinally over time between the animal and its external environment, and assess the animal's control by comparing goal states and actual states. It creates a record of the level of control and the factors contributing thereto for the early stages of consolidation in long-term memory. Primarily, the LSC is part of the memory system.

The LSC is serial, as this is the most reliable-- perhaps only—way in which to capture the temporal order of association-based 'causality' and control<sup>2</sup>. Its output is a multi-modal serial record of effectiveness and context-- a 'narrative of control'--for later consolidation in long-term memory. ('Narrative' does not imply any lexical content, merely a serial record of events.) The LSC receives input from long-term memory (shown in Fig. 5) to create a more germane narrative of control, and somatic / affective input (not shown) to assist memory formation and retrieval.

The LSC has a secondary function. As it picks up late combinations of highly processed information it is well placed to register potential threats missed earlier by the stress response system. It can therefore be considered as a very late-stage part of the stress response system, instigating (or inhibiting) at a late stage a stress response. Even so, in non-human primates there is very little flexibility or intelligence here.

Fig. 5 shows the site of a proposed adaptation to the stress response system in early hominins. A chance genetic mutation slightly raised the threshold of late-stage activation by the LSC of the stress response, at point A. Feedback from A, in the form of a signal that the animal is in control, enters the LSC at point B. The initial effect may be tiny, but as it occurs at a sensitive point in the system the effects are wide-ranging and complex. Three points are relevant: First, perhaps such variations occur often in evolution, but the stress response is highly conserved as it aids survival. The majority of variations will fail. But the potential benefit of getting it 'just right' could be huge. Second, the variation is to a threshold of *late*-stage activation of the stress response. There are many earlier thresholds of activation. Most important is LeDoux's 'Quick and Dirty' system [29], that activates the primate (including human) stress system quickly in response to part-processed stimuli resembling a hardwired or fear-conditioned threat. This is not a part of the mechanism that generates high-level intelligence, but it has facilitated the appearance of this mechanism as it protects the animal from many dangers and allows evolution to vary other areas of the stress response system. Finally, for the variation to gain purchase in an individual and in the species, and not be 'washed out' by learning, the threshold should be relatively non-plastic. Perhaps it is used by other brain regions for calibration.

Generally, comparators and feedback loops are concerned with the past, and so offer an unlikely site for the origins of prospective intelligence. This is especially so if the LSC is primarily a part of the memory system. The pre-adaptation LSC tracks the relationship between animal and environment fairly accurately. However, as the LSC (pre and post-adaptation) is part of a feedback loop incorporating the animal and its

---

<sup>2</sup> For an alternative explanation of seriality at this point, see Baars [28].



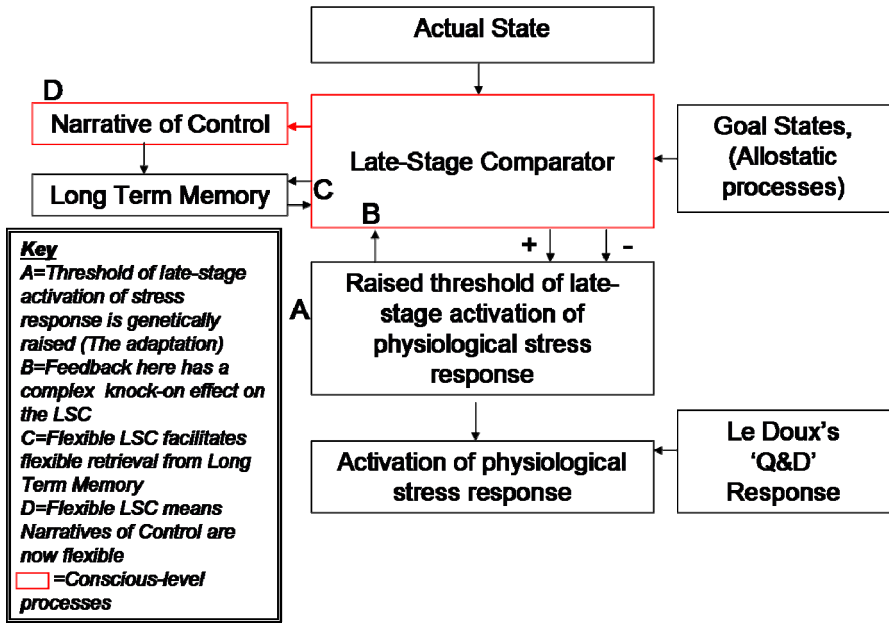


Fig. 5. Proposed Adaptation to the Primate Stress System

external environment it aims to bring actual and goal states together to solve proximate problems by sending error signals down to M2 modules. Pre-adaptation, there are just two elements, goal state and actual state, to be reconciled, so the phase space of possibilities is limited.

Fig.6 illustrates how such a process might generate purposeful flexibility.

Post-adaptation, feedback from the raised threshold enters the LSC at point B, indicating that the animal is in control of the situation. In some borderline situations this might cause a conflict, if the LSC otherwise assesses that the animal is not in control. Such borderline cases of conflict may not be important in themselves, but could become so if they open up a new landscape of possibilities. In the pre-adaptation LSC the relationship between the two elements of goal and actual states would be linear and able to express only a restricted range of possibilities. In the post-adaptation LSC there are three elements to be reconciled; the actual state, the goal state, and the message from feedback from the raised threshold indicating that the animal is in control. There is now the possibility of flexibility, as these three states can be reconciled in many ways. Goal states can be adjusted, actual states can be finessed, and pathways between them can become itinerant—guided by an imperative to seek ‘control’. Such itinerancy is shown by the blue lines in Fig. 6. The phase space has been expanded. Evolution may have generated and harnessed chaotic processes within a serial comparator, creating high-level intelligence in early humans.

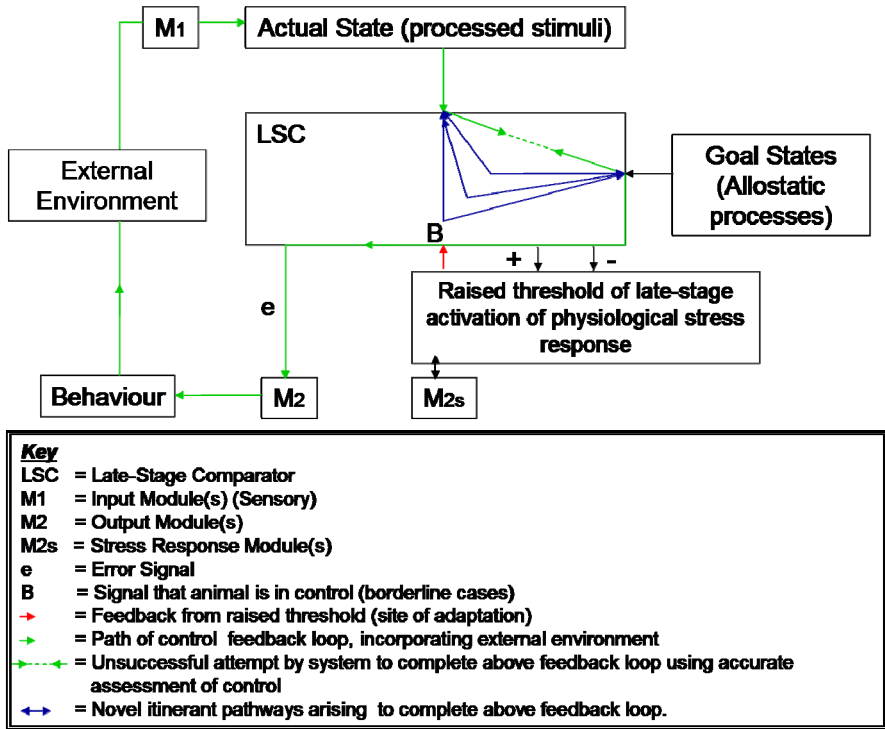


Fig. 6. Details of Post-adaptation Late-Stage Comparator (This does not show bi-directional pathways between LSC and Long-Term Memory.)

There are now two main types of negative feedback loop incorporating the LSC. The first is an ordinary negative feedback loop, with pathways in the LSC representing reasonably accurately the relationship between actual and goal states. Behaviour may be instigated to close the loop, and this is adaptive if it achieves a proximate goal. The second type may arise when the first type fails to achieve its goal. It involves itinerancy within the new landscape of possibilities, the expanded phase space within the LSC. This itinerancy, guided by an imperative to seek ‘control’, generates a flexible mixture of accurate and inaccurate representations of the relationship between actual and goal states. This may be adaptive in one of two ways.

First, guided itinerancy may generate novel pathways that eventually lead to a valid solution. As the LSC is freed from accurately tracking the present relationship between animal and environment it can now flexibly model new relationships, seeking control. This is facilitated by the LSC having bi-directional connections to long-term memory, as shown at point C in Fig. 5. The LSC can draw flexibly on information held in long-term memory, which includes elements of the animal’s model of its interactions with the world, to create new and anticipatory models in the LSC. Behaviour is then instigated by the LSC sending fairly crude error signals to M2 modules for detailed implementation. The LSC is now host to anticipatory, domain-general

intelligence. Second, the LSC may simply become ‘stuck’ in a closed loop incorporating inaccurate representation of actual and / or goal states and the relationships between them, causing the animal to assess *incorrectly* that it is in control. This latter possibility may be adaptive if it reduces allostatic load, and thereby contributes to the ultimate goal. The worst errors at this point will be winnowed out by evolution. If the problem is not resolved in any of the above ways a stress response may result, which may itself be adaptive.

Together, the above suite of options constitutes a form of flexible, high-level intelligence that acts as a buffer between the short and long term structural integrity of the animal and its environment. Nature has overcome the potential problem of an evolutionary bottleneck in an economical way. The same mechanism that creates flexibility in the LSC, guided itinerancy, can also make incorrect assessments of control, thereby reducing allostatic load. The brain is involved in a systemic balancing of proximate and ultimate goals over the lifetime of the person, in which buffer intelligence is optimised not for accuracy in pursuit of proximate goals but for achieving the ultimate goal. This balancing act is itself not intelligent; it is systemic.

This new form of intelligence has arisen in the hominin LSC at a point in the primate brain where there was previously very little flexibility. It is not an incremental development of other forms of intelligence, but rides on top of them, and is well-placed to exploit previously existing top-down pathways from the LSC to M2 modules by sending down intelligent instructions for execution. It is therefore a form of ‘top-down’ intelligence, complementing whatever ‘bottom-up’ intelligence earlier species may have. The new flexibility of the LSC might explain the origins of increased working memory, as a flexible LSC could reverberate both its own content and content received from long-term memory. A flexible LSC could share information between different M2 modules, increasing what Mithen [30] calls ‘cognitive fluidity’, and give rise to analogical thought. In principle, flexibility in the LSC could enable recursion of its own contents, and an ability to divide problems up into smaller constituent units and sub-goals, and flexibly manipulate potential solutions. It might also enable first order intentionality, with the animal becoming cognitively self-aware and self-monitoring at a conscious level. Lastly, as the LSC and its output, the narrative of control, may be the neural correlates of consciousness, we might hope for an eventual theoretical integration of high-level intelligence and consciousness that accounts for the relative lateness of conscious processes [31].

A capacity for increased intelligence could spread from the LSC to M2 modules in a way similar to that by which it originated in the LSC. A flexible LSC could send an increased range of signals to M2 output modules regarding controllability. This might in turn create the conditions for guided itinerancy to arise in an expanded phase space in any serial comparators in M2 modules. Intelligence could be ‘sucked down’ by M2 modules from the LSC, with different types and degrees of intelligent capacity arising in different M2 modules according to their structural capacities and ecological requirements. Such intelligence would generally operate automatically, below conscious awareness, with only highly processed abstracts of its output entering the LSC. An obvious candidate for this process would be M2 language module(s). The comparator in an M2 language module would initially be concerned to assess the

effectiveness of the communication process itself, but as in the LSC, flexibility in a comparator might also enable flexible organisation and recursion of its own contents, and give rise to flexible syntax.

## 5 Conclusion

It has been argued that the primate stress response system was a selective force and originating mechanism for increased hominin intelligence. The concept of ‘buffer intelligence’ may offer an evolutionarily plausible and economical framework for understanding the origins and development of several aspects of present-day human cognition, including flexible high-level general intelligence and increased creativity, expanded working memory, flexible retrieval from long-term memory, analogical thought, high-level monitoring of thought (cognitive self-awareness), and flexible linguistic syntax. A complex balance of costs and benefits has arisen from this stress-intelligence interaction over the hominin lineage. In present-day humans the main costs are quite frequent inaccurate appraisals of control and perhaps some rigidity of thought, and a residual but significant vulnerability to cognitively mediated stress and allostatic load and linked diseases. Outweighing benefits include a virtuous circle of increased intelligence, general health and longevity, and cross-generational transfer of the fruits of increased intelligence mediated by language.

Although for the sake of simplicity this paper presents the idea of a key genetic variation as a ‘one off’ event, most likely such variations occur frequently, with the majority failing (in line with their ‘high risk / high reward nature’) and those few that do gain an evolutionary purchase successively ramping up and refining the effects. Some version of the process may have occurred in different animal species including, perhaps, some that are now extinct.

Of the possible empirical ways of pursuing this question, comparative neuroanatomy, comparative genomics, and computer simulation to proof of principle of the concept of buffer intelligence, it is perhaps the last that may offer most interest to supporters of the INBIOSA project. Of particular interest is whether evolution may have generated and harnessed chaotic processes in a strategically placed serial comparator to expand phase space and create flexible high-level intelligence, whether the elements of this may be formulated mathematically, and whether the process can spread throughout a complex adaptive system.

## References

- [1] Duch, W., Oentaryo, R.J., Pasquier, M.: Cognitive architectures: where do we go from here? *Frontiers in Artificial Intelligence and Applications* 171, 122–136 (2008)
- [2] de Garis, H., Shuo, C., Goertzel, B., Ruiting, L.: A world survey of artificial brain projects, Part 1: Large –scale brain simulations. *Neurocomputing* 74, 3–29 (2010)
- [3] Goertzel, B., Ruiting, L., Arel, I., de Garis, H.: A world survey of artificial brain projects, Part 11: Biologically inspired cognitive architectures. *Neurocomputing* 74, 30–49 (2010)

- [4] Pearl, J.: Probabilistic reasoning in intelligent systems. Morgan Kaufmann, Los Altos (1988)
- [5] Ananthaswamy, A.: I, algorithm: A new dawn for artificial intelligence. *New Scientist*, 28–31 (January 31, 2011)
- [6] Calvin, W.H.: *How Brains Think*. Weidenfeld & Nicolson, London (1997)
- [7] Friston, K.: The free energy principle: a rough guide to the brain? *Trends in Cognitive Sciences* 12(7), 293–301 (2009)
- [8] Boyce, W.T., Ellis, B.J.: Biological sensitivity to context: 1 An evolutionary-developmental theory of the origins and function of stress reactivity. *Development and Psychopathology* 17, 271–301 (2005)
- [9] Ellis, B.J., Jackson, J.J., Boyce, W.T.: The stress response systems: Universality and adaptive individual differences. *Developmental Review* 26, 175–212 (2006)
- [10] Chrousos, G.P.: Stressors, stress, and neuroendocrine integration of the adaptive response. The 1997 Hans Selye Memorial Lecture. *Annals of the New York Academy of Sciences* 851(1 spec. no.), 311–335 (1998)
- [11] McEwen, B.S.: Protective and damaging effects of stress mediators. *New England Journal of Medicine* 338, 171–179 (1998)
- [12] Meaney, M.J.: Maternal care, gene expression, and the transmission of individual differences in stress reactivity across generations. *Annual Review of Neuroscience* 24, 1161–1192 (2001)
- [13] Bentley, P.J.: *Comparative vertebrate endocrinology*. Cambridge University Press, Cambridge (1998)
- [14] Nilsson, S., Holmgren, S. (eds.): *Comparative Physiology and Evolution of the Autonomic Nervous System*. The Autonomic Nervous System, vol. 4. Harwood Academic Publishers, London (1994)
- [15] McEwen, B.S.: Allostasis and allostatic load. In: Fink, G. (ed.) *Encyclopedia of Stress*, vol. 1, pp. 145–150. Academic Press, San Diego (2000)
- [16] Sterling, P., Eyer, L.: Allostasis: A new paradigm to explain arousal pathology. In: Fisher, S., Reason, J. (eds.) *Handbook of Life Stress, Cognition and Health*, pp. 629–649. Wiley, New York (1998)
- [17] McEwen, B.S.: Protective and damaging effects of stress mediators. *New England Journal of Medicine* 338, 171–179 (1998)
- [18] Sterling, P.: Principles of allostasis: Optimal design, predictive regulation, pathophysiology, and rational therapeutics. In: Schulkin, J. (ed.) *Allostasis, Homeostasis, and the Costs of Physiological Adaptation*, pp. 17–64. Cambridge University Press, Cambridge (2004)
- [19] Sapolsky, R.M.: *Why Zebras Don't Get Ulcers*. Henry Holt & Co, New York (2004)
- [20] McEwen, B.S.: *The End of Stress as We Know It*. Joseph Henry Press, Washington, DC (2002)
- [21] Arnold, M.B.: *Emotion and Personality*. Psychological Aspects, vol. 1. Columbia University Press, New York (1960)
- [22] Lazarus, R.S.: *Psychological Stress and the Coping Process*. McGraw-Hill Book Company, New York (1966)
- [23] Lazarus, R.S.: Thoughts on the relation between cognition and emotion. *American Psychologist* 37, 1019–1024 (1982)
- [24] Lazarus, R.S.: Vexing research problems inherent in cognitive –mediational theories of emotion—and some solutions. *Psychological Inquiry* 6, 183–196 (1995)
- [25] Lazarus, R.S.: *Stress and Emotion: A New Synthesis*. Free Association Books, London (1999)

- [26] Lazarus, R.S.: Relational meaning and discrete emotions. In: Scherer, K.R., Schorr, A., Johnstone, T. (eds.) *Appraisal Processes in Emotion: Theory, Methods, Research*, pp. 37–67. Oxford University Press, Oxford (2001)
- [27] Glass, D., Singer, J.: *Urban Stress: Experiments on Noise and Social Stressors*. Academic Press, New York (1972)
- [28] Baars, B.J.: Conscious contents provide the nervous system with coherent, global information. In: Davidson, R., Schwartz, G. (eds.) *Consciousness and Self-regulation*, pp. 45–76. Plenum, New York (1983)
- [29] LeDoux, J.: *The Emotional Brain*, pp. 161–174, Phoenix, London (1999)
- [30] Mithen, S.: *The Prehistory of the Mind*, p. 77. Thames & Hudson, London (1996)
- [31] Libet, B.: Neural processes in the production of conscious experience. In: Velmans, M. (ed.) *The Science of Consciousness*, pp. 96–117. Routledge, London (1996)