

Absolute pitch accessible to everyone by turning off part of the brain?

TERRY BOSSOMAIER^{†‡} and ALLAN SNYDER[‡]

[†]School of Information Technology, Charles Sturt University and [‡]Centre for the Mind, Joint Venture, Australian National University

and University of Sydney, Australia

E-mail: tbossomaier@csu.edu.au, allan@centreforthemind.com

Absolute (or perfect) pitch exists in fewer than 1/10,000 of the adult population and many claim that it can not be taught. On the other hand, research suggests that the mechanisms for absolute pitch exist in us all but access is inhibited during early maturation. We here argue that this inhibition can be switched off by artificially turning off part of the brain, allowing everyone access to absolute pitch. This possibility has profound implications for understanding the strategies adopted by the complex networks of the mind. We describe agent-based modelling techniques to understand the computational rationale for these inhibitory processes and to develop new artificial music recognition and synthesis techniques.

1. INTRODUCTION

Absolute, or perfect pitch (AP),¹ would not appear to confer a survival advantage on an individual, either now or at any time during our evolutionary history. On the contrary, there is a ‘cost’ for those who possess it (Zatorre 2003). Fewer than 1/10,000 of the population have absolute pitch (Profita and Biddes 1998) and that it is seemingly very difficult if not impossible to teach (Takeuchi and Hulse 1992). But, the fact remains that certain rare individuals do possess perfect pitch. Why is this so? Who are they and why do they have privileged access and not the rest of us to this ability?

To begin with, our mechanism for hearing consists of discrete frequency analysers which afford every individual with the raw apparatus necessary to confer absolute pitch. In other words, the information for absolute pitch is available to us all. Furthermore, it is likely that all infants have absolute pitch but that this ability is inhibited in the first year with the onset of language acquisition (Saffran and Griepentrog 2001, Saffran 2003). In other words, language is what is of ultimate importance to an individual and not the raw details (the frequency components) that make up the sounds (Snyder and Mitchell 1999). So,

¹We follow Deutsch’s (2002) approach where absolute pitch is essentially immediate and non-conscious amongst those who have it.

it would appear advantageous to suppress or inhibit the awareness of such details in favour of the meaning of sounds as suggested by Snyder and Mitchell (1999).

This observation about hearing has its analogy with vision. For example, our brain performs the calculations necessary to label a three-dimensional object from subtle shading across its surface. Yet, we are unable to consciously access such raw details, otherwise we could draw naturalistically without training (Snyder and Thomas 1997). Clearly such information is inhibited from conscious awareness (Snyder and Barlow 1986). In other words, it is the object label itself that is of ultimate importance and not the attributes or raw details processed by the brains to formulate the label. As with absolute pitch, it is even possible that newborn infants are aware of subtle shading prior to learning how to see. Some people, particularly artists, may perhaps acquire the skill to switch off such processing in a so-called ‘right-brain shift’.

But, why do some people have absolute pitch? First we note that all musical (autistic) savants have absolute pitch (Miller 1989). Furthermore, the incidence² of absolute pitch is below that of the incidence of adult autism. It remains to be investigated whether or not some of those reported to have absolute pitch fall within the high-functioning upper limit of the autistic distribution.³

On the other hand, absolute pitch seems to be preserved from infancy in those who were exposed to musical instruments at a very young age (Deutsch 2002, Schaug, Jänke, Huang and Steinmetz 1995). This, in our opinion, could be an instance where the brain is “tricked” into believing that the harmonically related tones of musical instruments constitute a language, resulting in special (left-hemisphere) brain structures for processing such tones.⁴

²The inheritability of absolute pitch is uncertain but Drayma, Manichaikul, de Lange, Snieder and Spector (2001) show that performance on the distorted tunes test correlates very strongly with identical compared to fraternal twins.

³Although autism is often associated with low IQ, it is not always the case. Average or higher IQ autistics are referred to as high-function (viz. Dustin Hoffman in *Rain Man*).

⁴The location of absolute pitch in the brain has recently been found

The fact that certain individuals maintain their ability to access absolute pitch from infancy proves that the process of inhibition can be delayed or frustrated, or that additional brain structures are created. We argue here that the inhibition can even be dynamically inhibited by artificial means (Snyder, Bossomaier and Mitchell 2004). And, importantly, recent experimental findings are consistent with the possibility that savant-like skills are accessible (can be switched on) by turning off (disinhibiting) the part of the brain that inhibits access to such skills (Snyder, Mulcahy, Taylor, Mitchell, Sachdev and Gandevia 2003). In particular, fifteen minutes of transcranial magnetic stimulation (TMS) applied to the left frontotemporal lobe enabled normal individuals to exhibit savant-like skills in drawing and in proofreading. Such magnetic stimulation is known to inhibit the normal functioning of localised regions of the brain (Pascual-Leone, Bartres-Fox and Keeman 1999). We here suggest that the application of TMS could be used in an analogous fashion in order to access absolute pitch.

Possibly related to the loss of absolute pitch is the way the ability to acquire an accent-free second language disappears around puberty. Henry Kissinger, Joseph Konrad and Vladimir Nabokov, for example, all developed excellent capabilities of expression in English but never lost very strong foreign accents. On the other hand Kissinger's brother, just a few years younger, had no such accent (Pinker 1994). We conjecture that the loss of the ability to hear the raw sensory data, underlying the loss of absolute pitch, may also underlie the loss of the ability to accurately hear the phonemes of a foreign language and thus develop a good accent. Thus disinhibition may also facilitate language learning in adults (Bossomaier 2003), which we discuss in more detail in Snyder, Bossomaier and Mitchell (2004).

The human brain is highly malleable. Blind people recruit the visual cortex; for speech, deaf people can recruit the auditory cortex. Yet, even when it might be advantageous, some things abilities are not accessible. Thus there must be adaptive computational strategies, either for speed or efficiency, underlying these restrictions. Computer simulation is a necessary adjunct to brain experimentation to understanding these strategies and for reusing them in artificial systems. The agent model we describe in this paper is intended to elucidate the trade-offs implicit in such computational strategies.

Before examining the implications of absolute pitch to inhibitory processes, it is first worth examining at what point in development it occurs; in human beings, whether it is present in children; and whether it is present (or even meaningful) in animals. Setting such a context will help

us in understanding the computational issues.

1.1. Childhood development

Saffran and Griepentrog (2001) and Saffran (2003) present evidence that in the first few months of life infants preferentially respond to absolute rather than relative pitch differences. Deutsch (2002) argues that the sensitivity to absolute pitch in early life may be captured by early musical training (rare) or by the exposure to tonal languages such as Mandarin or Vietnamese. Thus there is a high degree of accuracy of absolute pitch generation in speakers of these languages for the tones of the language. Deutsch conjectures that such populations will also display a much higher incidence of absolute pitch in music. However, speakers of tonal languages need to be able to generate only a few tones and thus have them explicitly labelled. In fact it is arguable that absolute pitch would interfere with language perception, since individuals, children versus adults, male versus female, will differ remarkably in the base frequency of their voices. It may be important, though, for an individual to have a consistent set of tones to be easily understood. Thus the number of labels required is much smaller than the 70 or so required for absolute pitch in music. Again, once the critical period has passed, and the left hemisphere architecture consolidated its inhibitory networks, it will be impossible, unassisted without disinhibition (e.g. by TMS), to learn all these additional labels for music.

The closing of the doors of perception does not only occur with sound. Pascalis, de Haan and Nelson (2002) show that infants 6 months old showed interest in novel faces, whether they were human or monkey. By 9 months of age, the behaviour is like adults, where the interest and discrimination is only focused on human faces. Human adults are worse at discriminating between monkey faces than human faces on forced-choice comparisons.

1.2. Animals and language

In trying to unravel strategies of the human brain, it is always advantageous to see what features are present in other species, and thus may have evolved for different purposes. We have argued that absolute pitch serves no use to humans and is actually disadvantageous. What happens in animals, particularly those which have musical behaviours?

Although no other species approaches humans in terms of language sophistication, there are language and music precursors in many other species. Thus contrary to the linking of music to language processing in humans, are the strong analogues of music in other species of the animal kingdom. Gray, Krause, Atema, Payne, Krumhansl and Baptista (2001) highlight strong similarities between animal species and humans in their construc-

to be the planum temporale (PT). Schaug, Jäncke, Huang and Steinmetz (1995) found that the left PT is enlarged in musicians without absolute pitch.

tion and learning of musical songs. Whales are noted for their songs, and humpbacks in particular show strong structural similarities to human music in rhythm, intervals, phrase length and larger scale form.

Three bird orders are known to have special brain structures active in bird song and talking and to be capable of vocal learning (parrots, songbirds and humming birds) (Holden 2002). Gray et al. (2001) describe how such birds also show strong similarities in rhythm and musical intervals. The hermit thrush uses a pentatonic scale, the canyon wren all 12 semitones. Are these behaviours imitation and manipulation of complex (absolute) spectra, or do they embody any sense of key, or, are animals more like human infants? The answer seems to be the latter.

Hauser and McDermott (2003) suggest that transposition is rare and at best minimal amongst other animals. Rhesus monkeys seem to have octave generalisation i.e. they judge tonal melodies to be the same if they are separated by one or two octaves, but they fail on even transposition by a perfect fifth. Songbirds do not have any transposition skill, implying that the capacity evolved after the divergence of birds and mammals.

Yet Tchernichovski, Mitra, Lints and Nottebohm (2001) found some octave knowledge in zebra finches. As they learn to mimic a song, they start either too high or too low in pitch. If they start too high, they continually increase the pitch until they reach an octave higher, at which point they simply halve the frequency.

Thus it seems that the loss of absolute pitch is associated with more advanced conceptual abilities, which seem to be related to right-left asymmetry. Great apes are now known to exhibit distinct asymmetry in the analogue of Broca's (language) area (Cantalupo and Hopkins, 2001), but there is a paucity of data to judge to what extent they have lost transpositional recognition capabilities. It remains a moot point as to whether the loss of absolute pitch is driven specifically by human language or by advanced conceptual skills.

1.3. The penalty of absolute pitch

If inhibition subserves a computational strategy, then, all other things being equal, we should expect to find a penalty for introducing absolute pitch (low-level detail) into the conceptual areas of the left hemisphere. There is evidence that these penalties do exist. Miyazaki and Rakowski (2002) examine the performance of subjects with and without AP in the recognition of notated melodies. They found that AP interferes with the recognition of melodies presented aurally transposed relative to the manuscript notation. Although extremely interesting these results depend critically on a small number of sub-

jects and further experiments are eagerly awaited.⁵

A similar interference occurs with language. Peña, Bonatti, Nespor and Mehler (2002) examine the ability of subjects to learn a grammar composed of artificial words made up of French phonemes. The boundaries between words have clear statistical properties, and, just given a stream of phonemes, they can determine the word boundaries quite successfully. But they are not able to ascertain the grammar even after prolonged exposure. The interesting result, however, is that if the boundaries between the words have a slight, subliminal gap of 25ms, inserted between them, the grammar is easily learned. Thus, *if we take away focus on the details, we can learn higher level structures.*

There are several possible reasons for these interference effects. There may be simply a shortage of neurons, but we doubt this is the case, since many neurons are lost shortly after birth. Perhaps there is too much competition at higher levels if the lower level representations are allowed to compete with higher level concepts. This is consistent with suggestions that experts manipulate fewer components in coming to a decision than novices (section 4). Alternatively absolute pitch may promote recognition "units" which are locked to a particular key, thus blocking generalised key-independent recognition.

The importance of fluidity of representation is apparent in recent work attempting to map the location of tonal relationships in the brain. Psychophysical and EEG measurements support a torus model for the tonality and keys of Western music. The major and minor keys each form a loop which wraps around three times, with each major flanked by its relative and parallel minors (Zatorre and Krumhansl, 2002). But However Janata, Birk, van Horn, Leman, Tillman and Bharucha (2002) find that the tonality map is distributed throughout the medial prefrontal cortex. But what is ; and what is particularly interesting is that it varies from session to session for a given listener. There is no fixed population corresponding to parts of the tonality surface. Thus the tonal representation of music is itself highly dynamic and even further dissociated from the absolute frequencies of the incoming sound.

In fact, too much low-level data may be the computational difficulty experienced in autism. There are indications that the brain grows too fast in the early years in autistic children. Thus they get swamped by detail and can never successfully learn the concepts they need for normal development. This notion of starting with a small brain and letting it grow at a controlled rate in line with experience is supported by computer simulations by El-

⁵The editor (David Worrall) noted that this is anecdotally well known by teachers of music students who possess AP. Such students seem to learn to readily transpose on-the-fly as do those that play transposing instruments. But out-of-tune instruments seem to produce similar interferences and cause annoyance at the necessity for them to invoke conscious attention to the activity.

man and others (see Elman, Bates, Johnshon, Karmiloff-Smith, Parisi and Plunkett 1997, p340ff) Although investigative neuroscience will reveal progressively more of how the brain works, we contend that understanding these computational trade-offs requires computer simulation as we discuss below.

2. RESTORING ABSOLUTE PITCH

Snyder and Mitchell (1999) predicted that savant skills like perfect pitch were latent within us all, and that under some circumstances they might be released. Since that time a good deal of supporting evidence has come to light. Snyder, Bossomaier and Mitchell (2004) discuss a range of ways in which such skills may be released in ordinary people. All of these essentially involve destroying or blocking a higher-level conceptual process which is itself blocking access to lower-level sensory data, or just turning off the blocking (inhibiting) mechanisms. Brain injury, especially strokes, various forms of mental illness, such as schizophrenia, and drugs may all have this disinhibitory effect, while some exceptionally creative people may have developed skills to bring this blocking under direct control (Snyder and Mitchell 1999, Snyder et al. 2003, Snyder, Bossomaier and Mitchell 2004). But the most exciting method at the moment however is magnetic brain stimulation or transcranial magnetic stimulation (TMS).

The application of a magnetic field at about one pulse-per-second can inhibit localised regions of the brain (Pascual-Leone, Bartres-Fox and Keeman 1999). The effects normally persist for only minutes after the field is turned off. Snyder et al. (2003) show that TMS inhibition of the left frontotemporal lobe enables otherwise artistically-challenged individuals to draw convincing dogs and to enhance proofreading skills, where a focus on the lexical details must dominate reading for understanding.

Magnetic fields do not penetrate more than a few centimetres below the surface of the skull, and thus not all brain areas are accessible. The sheer range of different brain areas, sometimes quite small, necessitates much experimentation in finding the optimal stimulus location. Thus to get the most from TMS, and in particular to use it to induce absolute pitch, requires both a detailed understanding of neuroanatomy, obtained from imaging techniques such as fMRI, and a modelling framework for understanding the way the mind processes music.

3. MODELLING DYNAMIC INHIBITION

Our central premise is that absolute pitch, like all other savant skills, can be switched on by turning off part of the brain. In other words, the mechanism for absolute




Figure 1. The concept takes raw data and feeds processed information to the executive. It blocks transmission of raw data to the executive process. (Adapted from Snyder, Bossomaier and Mitchell 2004)

pitch resides equally within us all but only certain individual (mainly savants) have privileged access (Snyder and Mitchell 1999, Snyder, Bossomaier and Mitchell 2004). To appreciate the process of artificial access, we designed an agent-based computational model which encapsulates the dynamic inhibition of absolute pitch. We say dynamic because it allows for the inhibition to be turned off and on. Figure 1 shows the canonical dynamic inhibition model: raw data feeds concepts which block its access to executive processes.

The agent model has three phases and goals. The first is to deepen our understanding of why absolute pitch is suppressed by determining the computational trade-offs in speed, efficiency, number and complexity of agents and flexibility. By making *teams* of agents compete in recognition of melodies, independent of key, we can identify the characteristics of the best performers. The second stage is to investigate how the agent model will generalise to recognition of style, i.e. higher level patterns, or large scale grouping of notes. Transposition of fragments is crucial to being able to recognize such higher level patterns, and, again, locking in to fixed frequencies may be very disadvantageous. The third and final stage is the exchange of agents within a team to promote creativity, analogous to the replacement creativity model implemented by Goldenberg, Mazursky and Solomon (1999).

Neural modelling has reached a very high degree of sophistication from the membrane/synapse level up to thousands of nerve cells. Simulation at such a low level can show us that our model of neural dynamics will lead to brain-like behaviour. Yet at the same time such simulations do not tell us about the dynamics of the higher-level emergent phenomena, such as, say, AP. Thus although inhibition is itself a low-level neural process, understanding its cognitive implications, particularly as revealed by TMS, requires modelling at a coarser-grained (higher) level, where the building blocks may consist of

many neurons. In line with current software terminology, we refer to these building blocks as agents (Weiss 1999). We saw above in section 1.3 that music perception and generation activates many brain areas, again indicating this agent-based approach.

In the present article, we do not consider the mapping between agents and nerve cells, other than to observe the following. There is a mapping between an agent and a collection of nerve cells at any given time. But each neuron may simultaneously contribute to more than one agent. Conversely, as an agent's behaviour changes with time, its neural makeup may change. Agents are an emergent property of neural dynamics. As in any complex system, (Bossomaier and Green 2000), emergent properties cannot in general be predicted from the lower level building blocks.

3.1. Collective agent intelligence

Just as in human organisations, teams work better than individuals. Many people have suggested that mind is a collection of agents, such as Marvin Minsky's Society of Mind (1985). Now, with the advent of brain imaging we are learning more and more about the intricate structure of the brain; in fact as Colheart (2001) asserts, the whole field of neuropsychology rests upon the assumptions of functional modularity. The challenge is in understanding what the modules are and why things are divided up the way they are. For example, the facial fusiform area has long been thought to be the centre of face recognition, but recent studies have shown that the situation is much more complicated (Haxby 2001).

Thus there is considerable interest in building teams of software agents to carry out complex tasks, especially if the agents can self-organise into effective teams. By studying the way agents evolve in competition and the way they interact, one can gain deeper insight into the computational constraints on agent organisation.

Collective agent intelligence requires ways in which agents can learn and evolve to perform their individual roles better. At the same time agent teams must avoid the "tragedy of the commons" (Harbin 1968), wherein a free, shared resource is destroyed through overuse. Two approaches have achieved some success in autonomous agent team building.

Baum and Durdanovic (2000) propose the notion of an agent economy, wherein each agent can bid for ownership of the whole world, on which it then performs its own individual actions and reaps appropriate rewards. Since the agents benefit directly from the growth of the world economy, the tragedy of the commons is avoided. Baum's approach performs well on tasks such as the Towers of Hanoi or Rubik's Cube. However, agents operate as individuals at all times, although they effectively cooperate to carry out sub-tasks. So, Baum's systems are rather like a

symbiotic ecology rather than a single organism.

Wolpert and Tumar (2002) provide a mathematical framework in which agent fitness is defined partially with respect to the good of the whole population. Any increase in agent fitness also increases the global system performance. Agents may fight and gain at the expense of each other, but the net system gain has to be positive. Wolpert and Tumar (2002) apply their methodology successfully to optimisation of computer networking and traffic.

3.2. Collective agents for word/melody recognition

As a precursor to studying the recognition and creation of melodies, we have studied the process of agents recognizing long words (15 letters or more, as a written (not spoken) string) from a dictionary (Cornforth, Kirley and Bossomaier 2004). Agents begin with ownership of single letters, which they are then able to trade via a tendering process. The number of letters they can hold is not limited after initialisation. Agents have a set of parameters which governs their willingness to trade, their price aggressiveness and so on. These parameters evolve as agents meet with more or less success in the market. The performance of the system as a whole requires a diverse set of parameters at the beginning, in other words a diverse set of agent skills.

Words are analogous to melodies, stripped of rhythm, metre and timbre. However, the adjacency statistics are quite different and repetitive structures are less common and less systematic. The other key difference is there is no analogue of transposition, a key element to our music agent system. Thus agents have to recognize melodies, independent of key. Following Baum et al. (2000), any agent can bid to recognize part of a given melody and receives a reward related to the length of melody it is able to match. The relationship between reward and length is a key element to collaboration, as we discuss in the next section.

3.3. Building teams of agents

For an agent to become part of a team rather than to operate by itself, it needs to obtain a greater reward as a team member. In the melody context we achieve this by providing a nonlinear reward for increasing melody recognition, such as the square of the melody length. It then becomes possible to feed back a reward to agents which recognize melody components a reward greater than they could receive for their component alone. Each agent thus contracts its skills to a higher level agent, no longer competes as an entity in its own right and loses the option to "bid" for the melody. Figure 2 shows the overall agent framework. Destruction or suspension of an agent mimics the disinhibition process, by TMS or injury as we discussed

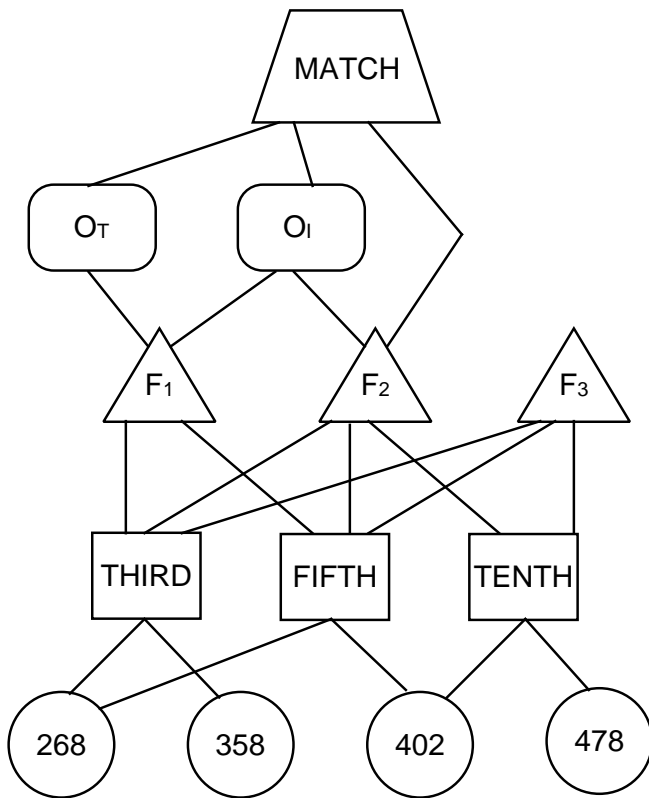


Figure 2. Raw frequencies are interpreted in terms of the intervals between them. F_1 - F_3 represent melody fragments, while O_T and O_I indicate such operations as transposition and inversion.

above. The lower level agents are now free to interact with other agents.

In the case of word recognition, word length is an adequate measure of complexity. But with melodies, there are repetitive structures which can be captured with simpler descriptions than an enumeration of every note. Thus agents need to be able to execute various sorts of procedures such as repetition, transposition etc. Thus we move from agents as simple holders of patterns which they can sell to the highest bidders to LISP functions which can implement learn a variety of musical operations.

These more intelligent agents can now receive a reward for a different sort of partial recognition. Rather than exact sequences of notes (absolute or key-referred), they can replicate higher level structural patterns and thus generate and recognize stylistic patterns.

3.4. Towards creativity

The final phase of the agent system moves from recognition to generation of new music. Once we have evolved agents capable of augmenting and transforming an existing melody, by swapping agents around, or changing the order in which they operate, entirely new musical patterns can be created.

Goldenberg *et al.*(1999) demonstrate computer creativity for the generation of advertisements by substitution of one component in an assembly by another, the so-called replacement route to creativity. Thus agents which add blocks of notes, or carry out operations on existing note sequences, may be replaced by other agents from the current pool of agents. The evolutionary approach we adopt goes further in allowing the components to develop independent of human control. At this stage, objective assessments of agent team performance have to give way to more subjective evaluation. But this is a common practice in evolutionary design, beginning with Richard Dawkins biomorphs (Dawkins 1986) and now a fully fledged design practice (Saunders and Gero 2001; Bentley 2002).

3.5. Other machine learning methods

State-of-the-art methods in machine learning use statistical inference techniques to determine likely musical sequences. Dubnov, Assayag, Larlillot and Berjerano (2003) use Incremental Parsing (IP) and Prediction Suffix Trees (PST) in their OpenMusic system, which has generated concert-level compositions. IP follows the ideas of Ziv Lempel coding, used extensively in text compression and underlying image processing formats such as GIF (Witten, Moffatt and Bell 1994). IP builds a dictionary of patterns and phrases which it uses to generate likely successors at any given point in time. PST behaves in a similar way, but emphasizes common patterns and trades robustness for accuracy.

Such methods are very effective at a computational level, but they often conceal the nature of what is happening in a black-box manner. Statistical approximation has proved valuable and frustrating in other areas of neuroscience and psychology, for example Julesz early work on preattentive texture discrimination (Julesz 1975) characterised textures by their statistical properties. But as the experimental evidence and theoretical models became more precise, the structure of image tokens became the defining properties. Since we want our models to generate insight into human musical processing, we avoid such information theoretic approaches.

3.6. Inhibition, encapsulation and object-oriented programming

The dynamic inhibition framework has a very strong conceptual similarity to the major new paradigm in software engineering of the last decade object-oriented programming. The idea of objects is that they encapsulate data, providing access to it through functions (methods) which process and abstract it. No other computational process is allowed access. The virtues of object-oriented programming lie in safety (data cannot be corrupted by some other process) and reusability. Objects may be plugged

into other software, or be grouped into larger assemblies known as components or frameworks. Encapsulation also facilitates updating: providing the interface the object presents to the world outside it remains unchanged, its internal working can be modified at will (Jacobson, Griss, and Jonsson 1997).

4. DISCUSSION

We have argued here that AP is within everyone but that it is normally inhibited from conscious awareness. Our central premise is that AP can be switched on by turning off parts of the brain responsible for the inhibition. To appreciate the mechanism and strategy underlying dynamic inhibition we constructed an agent-based computational model.

We have argued that a grand strategy at work in the human brain is one of creating higher level concepts which discard low level detail and reduce the number of items one has to manipulate. Expertise itself is often characterised by simplification and the ability to work with fewer variables.

But this loss of access to detail is not necessary within vertebrate brain structure. Some animals have impressive memories for detail. Birds can bury nuts and retrieve them months later when the terrain is snow-covered. Every child learns that elephants never forget; that their excellent memory is essential for survival. They need to remember where there might be water holes many days walk away in times of drought. Navigational errors would bring a thirsty death (McComb, Moss, Duant, Baker and Sayialel 2001). The older elephants have travelled the furthest and experienced the greatest range of climactic conditions. Pachyderms value age!

Even sheep can do better on detail than man. Recent studies show that capacity for recognition of the faces of other sheep is extremely good. It is also 3D, in that they can recognize a sheep from the front when having seen only the side view (Kendrick, da Costa, Leigh, Hinton and Pierce, 2001).

The brain seems to be able to manipulate only a limited number of pieces of information at executive level, the famous Miller doctrine of 5 plus or minus 2. More recent studies suggest that working memory is similarly limited, as is the level of embedding possible in language. Studies of cognitive processes in design, comparing experts to novices, show that experts paradoxically operate with fewer variables than do novices. They have developed higher-level concepts, embracing and suppressing lower-level detail.

But the evolutionary need for dynamic inhibition is not simply a matter of shielding the executive from lower level detail. We can see this by looking at the performance of supervised, feed-forward neural networks. Representations

are formed by hidden units which transform and manipulate the input in complicated ways, before passing processed information upwards. What such artificial networks lack, however, is the ability to adapt quickly to changing requirements, a hallmark property of the prefrontal cortex. The interacting agent model can adapt quickly by just combining agents together in different ways.

The need for flexible templates is apparent throughout the history of western music. As the 19th century progressed, the intervals and harmonies which had stood for the previous two hundred years gave way to increasing chromaticism, reaching its apotheosis in Richard Wagner. But the 20th century was far worse. The whole notion of tonality started to collapse with new compositional methods such as serialism. Pierre Boulez (1971), one of its strictest composers and theorists, writes:

These actual octaves create a weakening, or hole, in the succession of sound relationships by way of provisionally reinstating a principle of identity denied by the other sounds, so that they are at variance with the principal of structural organisation of the world in which? they appear: actual octaves must be completely avoided at the risk of structural nonsense. (Boulez 1971).

Thus octaves act like sensory detail: they lock us into a particular perceptual mode and obstruct our grasp of higher-level structure.

Following the loss of harmonic traditions, we then get, more or less at the same time, the loss of determinism (addition of random elements), the production of unusual sounds from traditional instruments and the use of unconventional instruments such as concrete mixers. Thus as the musical paradigm shifts we have to break up and reconfigure agents at progressively deeper levels, first the intervals and finally the harmonic structure of the sounds themselves. Nevertheless, in Western musical cultures tonal and metric structures still remain prevalent.

We have argued in this paper that the loss of both absolute pitch and the sensitivity to phonemes of foreign languages results from the need for computational efficiency. But prevailing evidence suggests that this raw frequency information is not rewired into internal circuitry, forever lost to consciousness and executive processing. Rather it is dynamically blocked, like other savant skills, and can therefore be released by disinhibiting techniques such as TMS. We have shown that absolute pitch does entail trade-offs and many excellent musicians and composers have managed perfectly well without it. Restoring sensitivity to the phonemes of a foreign language has potentially major benefits to a wide range of people.

The reasons for dynamic inhibition are not yet established. The building of agent models for processing musical signals enables us to see the pressures on communication between agents for optimal speed, efficiency and creativity. Not only do such models help us to understand how and why the brain operates in these surprising ways, but they hold out the potential for new machine algorithms in music recognition, reconstruction and creation.

The hazards of absolute pitch emerge as a proliferation of agents owning or generating musical fragments locked to particularly keys. The most efficient solutions are where the building blocks are key independent at the lowest levels. This is analogous to the idea of constancies in sensory information processing – for example, absolute light level, which is generally irrelevant to object recognition, is removed at the retina of the eye and not transmitted to the brain.

The advantages of the hierarchical structure are apparent especially at the creativity level, where coarse grained “chunks” can be swapped, akin to whole subtrees of a computer program, analogous to the importance of crossover for search space exploration in evolutionary programming. Future developments will allow such chunks (agents) to be tagged with connotations outside the musical form itself – emotional labels, leit-motifs with specific narrative meaning and other application driven requirements.

Computer games now exceed the movie industry with revenue over *US\$20 billion* per annum. As in film, music is a crucial dimension, but can benefit from contextual variability. By modelling adaptive real-time synthesised sonic frameworks, our model, whilst helping to understand the explosion of new information about how the human brain operates, promises to supply insights into how to model flexible compositional tools such as those necessary for the development of music for real-time interactive gaming environments.

ACKNOWLEDGEMENTS

We thank the referees for many helpful comments in improving the clarity of the manuscript. We would also like to thank Toby Hawker for his assistance.

REFERENCES

- Baron-Cohen, S. 2003. *The Essential Difference. Men, Women, and the Extreme Male Brain*. Allen Lane.
- Baum, E. B. and Durdanovic, I. 2000. *An evolutionary post production system*. Submitted, available from <http://www.neci.nec.com/homepages/eric/eric.html>.
- Bentley, P. 2002. *Digital Biology*. Hodder Headline.
- Bossomaier, T.R.J. 2003. Complexity: metrics and modules. *Advances in Complex Systems* **6**: 313–29.
- Bossomaier, T.R.J. and Green, D.G. (eds). 2000. *Complex Systems*. Cambridge University Press.
- Boulez, P. 1971. *Boulez on Music Today*. Translated by Susan Bradshaw and Richard Rodney Bennett, London: Faber and Faber.
- Cantalupo, C. and Hopkins, W.D. 2001. Asymmetric Broca’s area in great apes. *Nature* **414**: 505.
- Colheart, M. 2001. *Assumptions and methods in cognitive neuropsychology*. In *Handbook of Cognitive Neuropsychology*, ed. B. Rapp. Taylor and Francis.
- Cornforth, D., Kirley, M. and Bossomaier, T.R.J. 2004. Agent heterogeneity and coalition formation: Investigating the effects of diversity in a multi agent system. *Third International Joint Conference on Autonomous Agents and Multi Agent Systems*. New York.
- Dawkins, R. 1986. *The Blind Watchmaker*. Harlow: Longman.
- Deutsch, D. 2002. The puzzle of absolute pitch. *Current Directions in Psychological Science* **11**: 2002–4.
- Drayma, D. and Manichaikul, A. and Lange, M. de, and Snieder, H. and Spector, T. 2001. Genetic correlates of musical pitch recognition in humans. *Science* **291**: 1969–72.
- Dubnov, S., Assayag, G., Lartillot, O. and Bejerano, G. 2003. Using Machine Learning Methods for Musical Style Modelling. *IEEE Computer* **36**: 73–80.
- Elman, J.L., Bates, E.A., Johnson, M.H., Karmiloff-Smith, A., Parisi, D. and Plunkett, K. 1997. *Rethinking Innateness*. MIT Press.
- Goldenberg J., Mazursky D., and Solomon S. 1999. Creative Sparks. *Science* **285**: 1495–6
- Gray, P.M., Krause, B., Atema, J., Payne, J., Krumhansl, C. and Baptista, L. 2001. The music of nature and the nature of music. *Science* **291**: 52–4.
- Harbin, G. 1968. The tragedy of the commons. *Science* **162**: 1243.
- Hauser, M. and McDermott, J. 2003. The evolution of the music faculty: a comparative perspective. *Nature Neuroscience* **6**: 663–8.
- Haxby, J.V. Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L. and Pietrini, P. 2001. Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science* **293**: 2425–30.
- Holden, C. 2001. How the brain understands music. *Science* **292**: 623.
- Holden, C. 2001a. Wordless Memories. *Science* **296**: 1233.
- Holden, C. 2002. Big bucks for bird brain studies. *Science* **296**: 1017.
- Holden, C. (ed). 2003. Mozart effect revisited. *Science* **301**: 914–15.
- Jacobson, I., Griss, M., and Jonsson, P. 1997. *Software Reuse*. ACM Press.
- Janata, P., Birk, J.L., van Horn, J.D., Leman, M., Tillman,

- B. and Bharucha, J.J. 2002. The cortical topography of structures underlying Western tonal music. *Science* **298**: 2167–70.
- Julesz B. 1975. Experiments in the visual perception of texture. *Scientific American* **232**: 34–43.
- Kendrick, K.N. da Costa, A.P., Leigh, A.E. Hinton, M.R. Pierce, J.W. 2001. Sheep don't forget a face. *Nature* **441**: 165–6.
- McComb, K. and Moss, C. and Durant, S.M. and Baker, L. and Sayialel, S. 2001. Matriarchs as Repositories of Social Knowledge. *Science* **292**: 491–4.
- Miller L. 1989. *Musical Savants: exceptional skills in the mentally retarded*. Hillsdale NJ: Laurence Erlbaum.
- Minsky, M. 1985. *The Society of Mind*. London: Heinemann.
- Miyazaki, K. and Rakowski, A. 2002. Recognition of notated melodies by possessors and nonpossessors of absolute pitch. *Perception and Psychophysics* **64**: 1337–45.
- Pascalis, O. and de Haan, M. and Nelson, C.A. 2002. Is Face Processing Species-Specific During the First Year of Life? *Science* **296**: 1321–3.
- Pascual-Leone A, Bartres-Fox D and Keeman J P. 1999. Transcranial magnetic stimulation: Studying the brain-behaviour relationship by induction of “virtual lesions”, *Philos Trans, R Soc London, B Biol Sci* **354**: 1229–38
- Peña, M. Bonatti, L.L, Nespore, M. and Mehler, J. 2002. Signal-driven computations in speech processing. *Science* **298**: 604–607.
- Pinker, S. 1994. *The Language Instinct*. Allen Lane, Penguin Press.
- Profita J. and Biddes, I. G. 1988. Perfect Pitch. *Am J Med Genet* **29**: 763–71.
- Saffran, J. R. 2003. Absolute pitch in infancy and adulthood: the role of tonal structure. *Developmental Science* **6**: 35–47.
- Saffran, J.R. and Griepentrog, G.J. 2001. Absolute pitch in infant auditory learning. Evidence for developmental reorganisation. *Developmental Psychology* **17**: 74–85.
- Saunders, R. and Gero, J.S. 2001. The digital clockwork muse: A computational model of aesthetic evolution, in *Proceedings of the AISB01 Symposium on AI and Creativity in Arts and Science*, Wiggins G., ed.
- Schaug, G. Jäncke, L., Huang, Y. and Steinmetz, H. 1995. In vivo evidence of structural brain asymmetry in musicians. *Science* **267**: 699–701.
- Snyder, A.W., Bossomaier, T. and Mitchell, D. J. 2004. Concept formation: ‘object’ attributes dynamically inhibited from conscious awareness. *J Integrative Neuroscience* **3**: 19–34.
- Snyder, A.W. and Mitchell, D.J. 1999. Is integer arithmetic fundamental to mental processing? The minds secret arithmetic, *Proc Roy Soc Lond. B Biol Sci* **266**: 587–92
- Snyder, A.W. Mulcahy, E., Taylor, J.L., Mitchell, D.J., Sachdev P. and Gandevia S.C. 2003. Savant-like Skills Exposed in Normal People by Suppressing the Left Fronto-temporal Lobe. *J Integrative Neuroscience* **2**: 149–58
- Snyder, A.W. and Barlow, H.B. 1986. Revealing the artists touch, *Nature* **331**: 117–18.
- Snyder, A.W. and Thomas, M. 1997. Autistic artists give clues to cognition, *Perception* **26**: 93–96.
- Takeuchi, A. H. and Hulse, S. H., 1992. Absolute pitch. *Psychol Bull* **113**: 345–61.
- Tchernichovski, O., Mitra, P.P., Lints, T. and Nottebohm, F. 2001. Dynamics of the Vocal Imitation Process: How a Zebra Finch learns its song. *Science* **291**: 2564–72.
- Trano, M.J. 2001. Music of the hemispheres. *Science* **291**: 54–6.
- Trehub S.E. 2003. The developmental origins of musicality. *Nature Neuroscience* **6**: 669–72.
- Weiss, G.(ed). 1999. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press.
- Witten, I.H. and Moffatt, A. and Bell, T.C. 1994. *Managing Gigabytes*. Van Nostrand Reinhold New York.
- Wolpert, D. and Tumar K. 2002. Collective intelligence, data routing and Braess’ paradox. *Journal of Artificial Intelligence Research* **16**: 359–87.
- Zatorre, R.J. and Krumhansl, C.L. 2002. Mental models and musical minds. *Science* **298**: 2138–9.
- Zatorre R J. 2003. Absolute pitch: a model for understanding the influence of genes and development on neural and cognitive function. *Nature Neuroscience* **6**: 692–5.