Modeling Human Vision for Heuristics

Sean McGerty, Frank Moisiadis

University of Notre Dame Australia

sean.mcgerty@gmail.com

frank.moisiadis@nd.edu.au

Abstract

Human vision is a highly complex system that has evolved to enable interactions within and on our environment in an expedient and resource efficient way. Applying black box testing principles gives us insight to functional parameterizations within the brain. Neuroscience further helps us understand the localization of functions within the brain. By modeling these specializations we gain a taxonomy for interactions between focused and persistent attention modes. Using this taxonomy we break down the interactions within evolutionary heuristic showing possibilities for granular hybrid behaviors. Finally we correlate this approach in a byplay between strategic and tactical concerns in military simulation.

Keywords: vision, attention, artificial intelligence, evolution, satisficing, particle swarm, genetic algorithm, simulated annealing, mutation.

1. Modeling human vision

Human vision acts like an accessible real-time processing system with modes of operation that have evolved [1] by a process of natural selection over millions of years. The modes of operation such as perception and identification allow us to understand, interact with and react to our environment.

Black box testing can help give us an understanding of the operating modes of a system without needing to know the detail of the implementation. Boundary testing helps us understand the acceptable ranges for inputs when correlated against outputs. Observation allows us to identify the operating modes. We can separately verify and validate these modes as though we were unit testing.

By careful performance testing we can introspect the functional modes of the brain. For example, we could separately test the performance for vision being able to differentiate the apparent motion of dots against a field of noise [2]. We could then separately test for the performance for the same person being able to similarly differentiate the apparent movement of a simulated sound. Then we can similarly test for both at the same time where the movement is synchronized between visual and audio prompts. If there were few shared components then the expected performance of the combined test would most likely be in the range of the first two tests. However if responses were quicker in the combined test then this may indicate a shared component that was correlating the signals into an aggregated result improving efficiency. In this case the signals and their apparent motion may be codified into a taxonomy that is shared between the visual and audio modes, and it is the combined inputs from the two modes that is leading to a faster response.

Therefore these approaches approximate black box testing and suit our need to parameterize visual inputs as we rarely have reason to open up a healthy brain. We are able to perform imaging and detect electrical impulses within the brain [3]. However there can be limitations with these approaches. Some methods lack clear resolution. Note too that the better we learn a task, we need less effort to perform it. This impacts methods that can measure blood flow within the brain, as we are less likely to resolve areas that have learnt a task well.

Much of the information we have is taken from patients who present with localized damage then re-present with reduced function. In this way we can extrapolate the same localization of function in a healthy brain. For example there are remarkable differences between the abilities of patients with right or left hemisphere damage to recall faces, or even recognize themselves [4][5].

Our aim is to produce a picture of shared subsystems, which have evolved in response to challenges and tasks in our environments. The possibilities for shared taxonomies, boundary interfaces and holistic approaches are intriguing. We will show that these behaviors and taxonomies are already significant for evolutionary software and related systems.

2. Anatomical Visual Pathways

Within each eye we find the macula at the center has the highest concentration of color sensing cones, while further from the center we find more rods, which are balanced towards black/white and motion detection [6][7]. So effective is the by-play between the detailed solution in the

macula, and awareness in the rest of the eye, that the macula may only occupy 10 degrees of the curvature [8]. Below our level of awareness we maintain a detailed focus in three dimensional space, while our awareness notifies us of external changes so that we can adapt our focus to a new item of interest if needed [9].
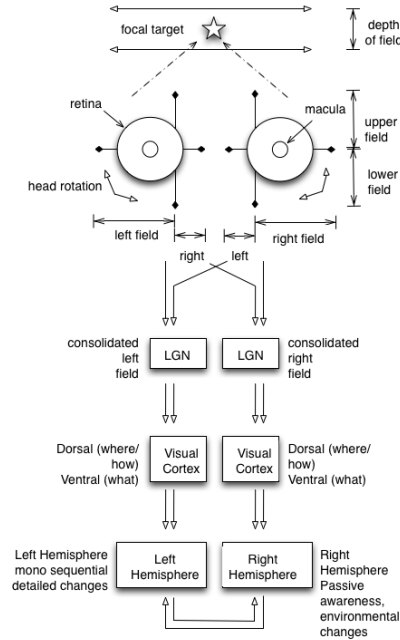

Fig. 1: Vision Processing Path

The Lateral Geniculate Nucleus (LGN) takes signaling from the eye, differentiating the left field of vision from both eyes towards the right hemisphere and the right field towards the left hemisphere [10]. In general terms the LGN continues the allocation of resources seen in the eyes. Of particular interest we find that there are differentiated cells in the LGN matching the breakdown of cones (EA) and rods (SB) in the eye [9][11]. Each LGN is broken down into layers [12]. The first two layers interface with the rods, one per eye. These are biased towards movement, depth perception and small variations in brightness. The next

four layers are also split between the cones in each eye. These are biased towards color, form and fine detail.

The Visual Cortex (VC) sits in one of the oldest areas of the brain, at the back near the brain stem [13][14]. Signals arrive at the V1 area, which are then orchestrated with neighboring dorsal and ventral regions. The ventral region assists with recognizing objects and surroundings. The dorsal region assists with the location of objects, having a role in feedback loops for the body interacting with its surroundings in three-dimensional space. Once again, we find discrete resources are differentiated between central/peripheral, and even upper and lower fields.

3. Vertical symmetry

When we look at the components along the visual pathways we find functionalities that are consistently replicated for each hemisphere and each eye along a vertical symmetry. In general terms this allows us to separately test a hemisphere by signaling into the visual field managed by the corresponding eye. As discussed, the signaling from both eyes is reconciled in the LGN into left and right fields of vision. If we hope to isolate the hemisphere we need to both obscure the other eye and ensure that the visual queue does not represent in both fields. This is achieved by asking the subject to watch a focal point while the item of interest is shown off to the side being tested.

The greater subset of functionality is duplicated on both sides of the visual pathway as we would expect, given the life experiences of people who have completely lost vision in one eye or the other. For example elementary abilities to detect motion in two dimensions relative to the subject, and the ability to detect changes in light intensity, have been demonstrated in the symmetrical visual pathways of fruit flies [15]. Different neural pathways fire when movement is perceived in the up, down, left or right directions. These studies also show differentiation between detecting edges via intensity changes and detecting movement of the edges. For expediency we would not expect to require a correlation between the hemispheres at this stage of processing, and so these functions are completely duplicated on each side.

This duplication of function continues to the visual cortex where the ability to recognize objects, and the ability to place them in our visual field represent equally on both sides. For example if we are driving a car our attention is largely forward and outside the car. However each hand is largely interacting with our environment in separate visual fields. Each cortex is assisting the actions within its visual field as the motion of a hand to a surface or object occurs.

4. Functional modeling

4.1. Qualitative vs quantitative

The usual experience of a person is that they have an awareness of their environment or the room that they are sitting in. However we know that only a small fraction of the surface of the back of the eye has the receptors capable of high definition or strong color recognition. The reconciliation of this impoverished view is achieved through orchestrated efforts to approximate detail and passively shift focus. Your visual system will focus on an object of attention while passive systems watch your surroundings for signs of change. As well when we are not focusing on a task our gaze passively moves about sampling our environment such that the remembered detail of an object is still

available once it moves into our peripheral vision or out of sight.

These paired ideas that we can see within the eye and represented back into our perception are very like qualitative and quantitative modes. When we hold focus on an object we are attempting to quantify it as accurately as possible. We register the relative movement and intensity of objects in our peripheral view, supplementing these with memories of detail acquired earlier. We can argue that we have a qualitative view on these objects before they become of interest again.

The optic nerve head is a good example of how the visual system compensates for lack of fine detail [16]. The area where the optic nerve connects to the back of the eye has no photoreceptor cones or rods, and yet we hardly ever notice. Our eyes seem constantly in motion sampling our environment and so the region obscured by the optic nerve has recently been seen. With an absolute lack of current detail a qualitative impression persists such that most people are not aware there is any deficiency.

4.2. Handedness

There are differences in the mix of focused and peripheral processing given the handedness of an individual [3]. A right-handed person spends more time performing tool manipulation with their right hand, which implies that the left hemisphere spends more time doing granular and detailed tasks. The right hemisphere therefore complements this activity by maintaining more of a perceptual view of the environment. When not performing a task, both hemispheres participate to help create the awareness of the environment and so both hemispheres have qualitative perception modes. However the right eye and left hemisphere are preferred for de-

tailed work, and so solution mode quantitative processing is preferentially performed on that side. These statements lead us to a view that quantitative solution based processing is a mode that largely appears on in the left hemisphere above qualitative processing modes in both hemispheres.

Split-brain theory correlates this idea by attributing specific mono-procedural solution modes to the left hemisphere [3]. Similarly the right hemisphere is better at qualitative broadly aware correlated tasks. The left hemisphere concentrates on tasks while the right hemisphere is more aware of other tasks and the passage of time.

5. Functional aspects

5.1. Satisficing vs. optimizing

What we have appearing therefore is a satisficing behavior [17], where our passive awareness is maintained until a threshold is reached and we then apply focused attention on the subject. 'Satisficing' is the idea of combining the concepts of satisfying and sufficing rather than optimality. A person will often need to make choices on partial information, which is termed a 'bounded reality' [18][19]. People intentionally make suboptimal choices, which they have rationalized to themselves as preferential. As flawed optimizers people are sometimes seen to choose an option from a range of solutions, which meet a fitness threshold. A satisficing behavior is a way of reaching a good enough solution without spending the resources to find the best possible. Satisficing behaviors suit qualitative processing mechanisms.

The quantitative solution processing of the left hemisphere is working for an optimal level of accuracy for the task at

hand. When a person moves to improve their view of a task or their surroundings they are repurposing their focus, and the focus can almost entirely be devoted to a task with the preferred hand.

However, optimality comes with a cost and it is the perception of the environment and the passing of time that sometimes lead us to change an approach to a task or stop it altogether. Therefore we balance the optimizing behaviors of the left hemisphere with the awareness and satisficing behaviors of the right.

There are indications that our ability to differentiate these abilities has grown as the ratio of the size of the corpus callosum to the rest of the brain has fallen. Our abilities to concentrate and devote ourselves to left hemisphere processing may also be related to neurons in the corpus callosum which inhibit rather than enable signaling from the right hemisphere to the left [20].

5.2. Heuristics

We can differentiate these functions if we look at the byplay within an evolutionary algorithm (EA). EAs are artificial intelligence (AI) heuristics with a lifecycle whereby they apply random mutations (RMs) to a population of candidate solutions. They then use a fitness function (FF) to assess which solutions are more successful, and they attempt to share these attributes with the weaker candidates [21][22][23].

Evolutionary algorithms are popular for producing good solutions to large and difficult problems such as transport networks, parameter optimization and operations research. Given a definition of a problem and related constraints the algorithms improve a population of solutions, usually until incremental improvement slows. The algorithms achieve this goal

without needing to check all possible solution combinations.

With vision, when we move the focus of our attention away from an object we retain a memory of it and may hold it in our peripheral vision in case something changes. If we follow the suggested EA lifecycle we can notice that iterative checks of the fitness function may be combined with random mutations so as to flag when a candidate has met a solution threshold. In effect for both cases we want to be aware of unexpected changes, which then lead us back to full attention. The satisficing threshold is the appropriate level for action.

The remainder of the EA is the ability to correlate attributes from successful solutions. The process of verification and validation of calculated changes requires comparisons with higher granularity. Therefore the processing of the left hemisphere seems ideally suited to this need.
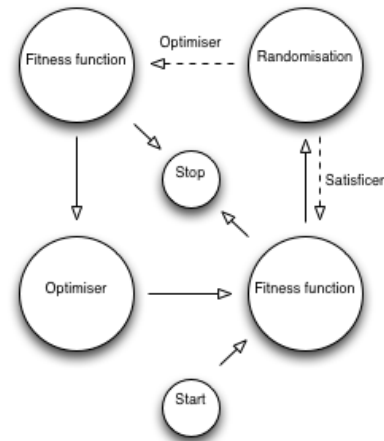

Fig. 2: Heuristic Lifecycle

If we represent these relationships in the heuristic graphically, we see that the qualitative process can represent as a subset of the quantitative process. This bears

direct correlation to the function we expect with vision in the left hemisphere.

We have also accomplished a secondary consideration by localizing the solution mode of the EA into a qualitative process. We can now have multiple instances of an EA with separate parameterizations sharing the same qualitative process.
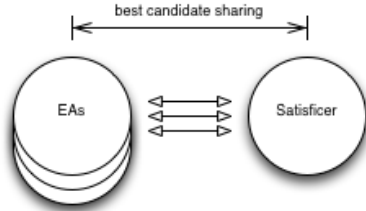

Fig. 3: Satisficer Sharing

We can even have multiple different EAs sharing the same qualitative process. In this case the satisficing threshold could be set to include the best candidates so far and these could be represented back towards each EA in turn. This would adaptively allow the best EA - for example particle swarm optimization, simulated annealing and genetic agorithms - to lead the others in response to success in their problem space.

6. Strategic vs tactical byplay

Consider a strategic simulation where a collection of units with varying abilities and operating modes require orchestration to reach an objective.

It may not make sense to attempt to find an optimal combination of units and deployment tasks given the bounded reality of the combat situation. We may not have a complete view of the battlefield and enemy dispositions, and even if we did, we may not correctly predict their movements or objectives. In this case it is far better to take a satisficing approach where deploying complementary units leads to more reliable results. The right hemisphere with its expansive view and abilities to correlate multiple inputs is ideally placed to react to contact with the enemy.

In this situation however each unit is tasked with achieving its objectives with singular focus unless told otherwise. That task may include support of correlated units, but the efficiency of the unit is related to singular focus. Resources spent on unexpected tasks deplete the unit and extend their arrival times at objectives.

In response to updates from the tactical units the satisficing behavior is in a better position to respond to a tactical unit that has stalled in pursuit of an objective. If progress has generally been good then the satisficer may have the unit hold their opponent in place while other units make progress. If progress has been low then the fitness threshold for the satisficer will be lower, leading to a recommendation that the unit detach and attempt to redeploy elsewhere.

7. Conclusions

The evolution of localized functions along the visual pathway gives clues to how expediency and resource efficiency has been achieved in response to and with our environment. Satisficing behaviors reappear many times as self-limiting decision making abilities.

We have shown that it is possible to model this satisficing byplay within the lifecycle of an evolutionary heuristic. Taking this view of the evolutionary heuristic exposes a separation of concerns similar to that seen between the hemispheres of the brain.

These concerns give a taxonomy for a discussion between strategic and tactical actors on a fictional battlefield.

7.1. Future work

It would be of great value to implement evolutionary heuristics within a component framework allowing a similar separation of concerns around a satisficing behavior.

A component framework implementation supporting multiple evolutionary algorithms would allow a highly granular hybrid model. Such a model holds the promise of unifying behaviors between diverse solution agents. There may also be new approaches for better mapping a namespace, which would better enable recognition and navigation around local maxima.

You could argue for example that the fitness function has more to do with the representation of the namespace than the heuristic. In this case a component framework would allow the sharing of the same optimized fitness function between each heuristic. This would also reduce inconsistencies between heuristics through the satisficing cache.

8. Acknowledgements

9. References

[1] Charles Darwin. On the origin of the species by natural selection. 1859.

[2] Georg F Meyer, Sophie M Wuerger, Florian Röhrbein, and Christoph Zetzsche. Low-level integration of auditory and visual motion signals requires spatial co-localisation. Experimental Brain Research, 166(3-4):538–547, 2005.

[3] Iain McGilchrist. The master and his emissary: The divided brain and the making of the western world. 2009.

[4] David J Turk, Todd F Heatherton, William M Kelley, Margaret G Funnell, Michael S Gazzaniga, and C Neil Macrae. Mike or me? Self-recognition in a split-brain patient. Nature neuroscience, 5(9):841–842, 2002.

[5] Michael S Gazzaniga. Forty Five years of split-brain research and still going strong. Nature Reviews Neuroscience, 6(8):653–659, 2005.

[6] Austin Roorda and David R Williams. The arrangement of the three cone classes in the living human eye. Nature, 397(6719):520–522, 1999.

[7] Robert W Rodieck and Robert W Rodieck. The ?rst steps in seeing. 15, 1998.

[8] David A Atchison, George Smith, and George Smith. Optics of the human eye. 2000.

[9] Eric L Schwartz. Spatial mapping in the primate sensory projection: analytic structure and relevance to perception. Biological cybernetics, 25(4):181–194, 1977.

[10] Soon Keen Cheong, Chris Tailby, Samuel G Solomon, and Paul R Martin. Cortical-Like Receptive Fields in the Lateral Geniculate Nucleus of Marmoset Monkeys. The Journal of Neuroscience, 33(16):6864–6876, 2013.

[11] Alumit Ishai, Leslie G Ungerleider, Alex Martin, Jennifer L

Schouten, and James V Haxby. Distributed representation of objects in the human ventral visual pathway. Proceedings of the National Academy of Sciences, 96(16):9379–9384, 1999.

[12] Margaret Livingstone and David Hubel. Segregation of form, color, movement, and depth: anatomy, physiology, and perception. science, 240(4853):740–749, 1988.

[13] Aina Puce, Truett Allison, Maryam Asgari, John C Gore, and Gregory McCarthy. Differential sensitivity of human visual cortex to faces, letterstrings, and textures: a functional magnetic resonance imaging study. The Journal of Neuroscience, 16(16):5205–5215, 1996.

[14] Kalanit Grill-Spector and Rafael Malach. The human visual cortex. Annu. Rev. Neurosci., 27:649–677, 2004.

[15] Matthew S Maisak, Juergen Haag, Georg Ammer, Etienne Serbe, Matthias Meier, Aljoscha Leonhardt, Tabea Schilling, Armin Bahl, Gerald M Rubin, and Aljoscha Nern. A directional tuning map of Drosophila elementary motion detectors. Nature, 500(7461):212–216, 2013.

[16] Lelia A Paunescu, Joel S Schuman, Lori Lyn Price, Paul C Stark, Siobahn Beaton, Hiroshi Ishikawa, Gadi Wollstein, and James G Fujimoto. Reproducibility of nerve fiber thickness, macular thickness, and optic nerve head measurements using StratusOCT. Investigative ophthalmology & visual science, 45(6):1716–1724, 2004.

[17] Herbert A Simon. Motivational and emotional controls of cognition. Psychological review, 74(1):29–39, 1967.

[18] Herbert A Simon. Theories of bounded rationality. Decision and organization, 1:161–176, 1972.

[19] Herbert A Simon. Rationality as Process and as Product of Thought. The American Economic Review, 68(2):1–16, 1978.

[20] Axel Kammerer, Álvaro TejeroCantero, and Christian Leibold. Inhibition enhances memory capacity: optimal feedback, transient replay and oscillations. Journal of Computational Neuroscience, 34(1), February 2013.

[21] Scott Kirkpatrick, D Gelatt Jr, and Mario P Vecchi. Optimization by simulated annealing. science, 220(4598):671–680, 1983.

[22] Carlos M Fonseca and Peter J Fleming. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. 423:416–423, 1993.

[23] J Kennedy and R Eberhart. Particle swarm optimization. . . . 1995 Proceedings, 1995.