



UNIVERSITY OF LEEDS

This is a repository copy of *Maximum utility unitary coherent perception vs. the Bayesian brain*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/108617/>

Version: Accepted Version

Proceedings Paper:

Fox, C orcid.org/0000-0002-6695-8081 and Stafford, T (2012) Maximum utility unitary coherent perception vs. the Bayesian brain. In: Proceedings of the Annual Meeting of the Cognitive Science Society. Annual Meeting of the Cognitive Science Society, 01-04 Aug 2012, Sapporo, Japan. .

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Maximum utility unitary coherent perception vs. the Bayesian brain

Charles Fox (charles.fox@sheffield.ac.uk)

Department of Computer Science, University of Sheffield, UK

Tom Stafford (t.stafford@sheffield.ac.uk)

Department of Psychology, University of Sheffield, UK

Abstract

Our subjective experience of the world is ‘unitary coherent’ (UC). ‘Unitary’ means we only perceive one interpretation at a time rather than a blur of multiple possible worlds. ‘Coherent’ means that we almost always perceive scenes that do not contain contradictory parts. While this form of first-person perceptual experience may seem obvious, it is in opposition to the requirements of optimal decision making, and to some forms of the ‘Bayesian brain’ hypothesis. We hypothesise that there are at least three types of ‘Bayesian’ action selection occurring in cognition, including a ‘maximum utility (MU) percept’ strategy that makes use of UC percepts. We give evidence from a video game experiment that is compatible with MU/UC perception and action selection, and is incompatible with optimal actions. Furthermore, it is compatible with the presence of utility bias in MU/UC perception: by changing the available actions we may be able to manipulate the subject’s percept of a fixed ambiguous stimulus.

Keywords: Bayesian; psychophysics; utility; bias; perception

A paradox about perception

Our subjective experience of the world is ‘unitary coherent’. *Unitary* means we only perceive one interpretation at a time (e.g. either a face *or* a vase in the Rubin Vase illusion) rather than a blur of multiple possible interpretations (never the face *and* vase together). *Coherent* means that we almost always perceive scenes that do not contain contradictory parts. (e.g. we do not see part face and part vase). While the UC nature of perception may seem obvious from subjective experience, it is in opposition to the requirements of optimal decision making, which require consideration of *all* possible interpretations of sensory data (Bernardo & Smith, 2000). In particular, the ‘Bayesian brain’ hypothesis (Doya, Ishii, Pouget, & Rao, 2007) views perception as computing probabilities of many interpretations, and optimal actions would be found by integrating out the utility of all actions under all percepts. If both Bayesian brain research and optimal action theory (Körding & Wolpert, 2004) suggest that perception should operate using a distribution, or ‘Bayesian blur’ of possible percepts, why then is our subjective experience limited to a unitary coherent percept instead? And which unitary coherent percept do we perceive: the most probable one or the most useful one? This study argues that as full Bayesian perception and action selection is computationally hard, an approximation which we call ‘maximum utility (MU)’ perception is a useful surrogate. It then presents evidence in support of the maximum utility perception hypothesis using a video game style experiment.

Types of perception and action selection

When a unitary coherent percept is required in machine perception, such as the output of a machine vision (Felzenszwalb & Huttenlocher, 2005) or speech recognition system (Young et al., 2006), the maximum a posteriori (MAP) state is often used by system engineers,

$$s_{MAP} = \arg \max P(s|d), \quad (1)$$

where s are world states and d is the available data. However real-world agents are often required to make actions as well as – or instead of – reporting percepts. In these cases, perceiving the MAP state does not necessarily lead to the best action if the following naive action-selection rule is used as a separate stage following MAP perception,

$$a_{naive} = \arg_a \max U(a, s_{MAP}), \quad (2)$$

where U is utility, a are actions. Instead, optimal actions are obtained (Bernardo & Smith, 2000) by maximising *expected* utility (MEU), which requires integrating over the ‘Bayesian blur’ of possible worlds,

$$a_{MEU} = \arg_a \max \int_s U(a, s) P(s|d) ds. \quad (3)$$

MEU action selection has no role for unitary coherent percepts. Instead it must consider *every* interpretation s . Computational approximations to this integral might ignore some improbable interpretations (Spiegelhalter, Thomas, Best, & Gilks, 2008), but still sum over a *set* of possible world states s rather than privileging any particular unitary state.

In many cases humans have been claimed to make optimal actions (Griffiths & Tenenbaum, 2006). This may occur for low-level, rapid stimulus-response type actions, and for high-level cognitive decisions such as business and financial decisions. Much recent work in ‘Bayesian Cognitive Science’ proceeds by *assuming* an MEU framework, then reasoning backwards from observed actions to report human priors on various stimuli (Stone, Kerrigan, & Porrill, 2009).

So why then do we bother to perceive UC percepts? Do they have some functional significance as well as being correlates of subjective experience? If they do have a function, this would suggest that Bayesian Cognitive Science’s assumption of optimal action is flawed, and could potentially invalidate some of its reported human priors.

Bayesian inference and hence MEU decision making is generally an NP-hard problem (Cooper, 1990) so is impractical for all but the most constrained percepts and actions.

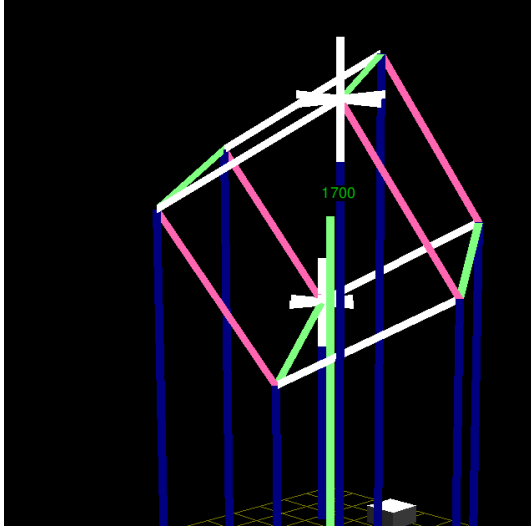


Figure 1: 3D environment used in the training phase of the game. A joystick moves the missile launcher around a 2D (x, z) plane on the ground. Pressing and holding the joystick button fires a missile (not shown) vertically upwards (y axis). Releasing the button detonates the missile. Points are scored for detonation close to the target(s) shown by white crosses. The training phase shown here includes colour, overlap, perspective and support cues to make the cube's configuration unambiguous. These cues are removed in the test phase to leave an ambiguous Necker cube, with ambiguous 3D target positions. Figure is best viewed in colour.

It has been suggested (Gigerenzer & Todd, 1999; Goldstein & Gigerenzer, 2002) that making actions based on a single 'best' percept (such as the 'take the best' heuristic and 'less if more; effect) could be a useful heuristic to speed up the decision making process at the expense of optimality. However the 'percepts' in these cases are high level logical states of the word rather than actual perceptual objects in three dimensional space.

We propose an alternative form of perception and action selection to MAP perception and MEU action selection, which we call *maximum utility perception (MU)*. In MU we choose a UC state and action together,

$$(s_{MU}, a_{MU}) = \arg_{s,a} \max U(a, s) P(s|d) \quad (4)$$

which yields the best possible action assuming that only a single world state can be considered.

The MU Hypothesis

We hypothesise that humans have at least *three* kinds of decision making behaviour, moving from fast and simple to slow and accurate:

1. *Immediate stimulus-response (S-R)*. A fast association from input data directly to an action. Such mappings do not need to build a UC percept. The could be implemented neurally at the sub-cortical level, such as direct links from supe-

rior colliculus to basal ganglia (Redgrave, Prescott, & Gurney, 1999). For simple mappings, it is possible that simple computational structures such as small neural networks could learn to perform near-optimal action selection, as has been demonstrated by computational experiments (Ramamoorthy & Verguts, 2012). Near-optimal performance in fast, low-level tasks such as reaching and pointing *quickly* at spheres having different locations and utilities (Kording & Wolpert, 2007); and 'simple heuristics' (Kahneman, Slovic, & Tversky, 1982; Gigerenzer & Todd, 1999) would be candidates for this mechanism. For simple tasks this method would give MEU-like results but without explicitly performing the MEU integration.

2. *Fast MU/UC percept-response (P-R)* which achieves suboptimal MU action in reasonable computation time. UC perception could be implemented cortically, with high-level perceptual areas computing a single most *useful* percept of the world, jointly with action selection under utility bias. Evidence for UC perception is found in binocular rivalry experiments (Srinivasan & Nunez, 2006), and in computational models (Riesenhuber & Poggio, 1999) as well as in everyday subjective experience. This paper gives evidence for MU/UC perception.

3. *Full MEU* action selection, via conscious sequential consideration of many possible percepts and responses. This slow type of decision making would occur for example when making a business decision, where several minutes (or even hours or days) are set aside to consciously perceive one possible world at a time, and the effects of many possible actions in them are simulated, and the resulting utilities averaged over. Humans are well-known to be poor at this kind of computation (Kahneman, 2003), and real-life action selection of this type is often performed in the business world by specialised operations researchers making use of computers to calculate the expected utilities (Pourret, Naïm, & Marcot, 2008), rather than relying on their own cognitive faculties.

If multiple decision making systems exist, it seems likely that the basal ganglia system is used to switch between them, for example taking account of time pressures for the type of decision to be made (Lengyel & Dayan, 2008; Redgrave et al., 1999; Daw, Niv, & Dayan, 2005). Strong support for the existence of at least two systems comes from the Ebbinghouse illusion, which produces different perceptual reports in verbal and stimulus-response type actions. It has been shown (Goodale & Milner, 1992) that the motor actions are consistent with optimal MEU-like decisions in the *same* subjects that make incorrect verbal reports.

While research on near-Bayesian optimal decisions of the S-R and Full MEU types abounds, there has been comparatively little work on the role of unitary-coherent perception in decision making. While our subjective experience tells us very clearly that *something* in the brain is computing a UC percept (which is incidentally presented to our conscious experience), and researchers have modelled how MAP percepts could be computed in this way (Riesenhuber & Poggio, 1999)

there has been little study of how this type of perception could be used in action selection as a replacement for S-R and MEU behaviour. Our hypothesis is that MU perception and action is in fact the dominant mode of everyday, aware, perception and action – the type of cognition that occurs consciously but not deliberately.

It is difficult to design experiments to isolate this middle, MU, level of perception, because as soon as subjects know their performance is being monitored they tend to start deliberating as in Full MEU, rather than performing ‘everyday’ perception and action selection. Conversely, if tasks are too low-level and fast-paced, they will use rapid S-R behaviour. Perhaps that is why few experiments have noticed MU effects before. To this end, we have carefully designed a simple 3D perception task, and examine two hypotheses:

Hypothesis H1 is that there are examples of human behaviour that are consistent with UC perception and inconsistent with both Full MEU (deliberative) and approximate MEU (S-R). A positive result here would stimulate further research into delimiting the circumstances in which the different behaviour types are employed.

Hypothesis H2 is that the particular kind of UC used in human perception is the MU percept. To find evidence for this stronger hypothesis, we will examine if it is possible to bias the percept from equally a priori probably percepts by altering the available action set, as predicted by MU.

Methods

A video game – loosely based on “space invaders” – was designed and implemented¹, having optimal MEU actions that require consideration of multiple scene interpretations, and having MU actions giving suboptimal rewards. If human behaviour in this (or any other) game could be shown to deviate from MEU behaviour and be consistent with MU, then evidence is provided for H1. Further, if the human behaviour is consistent with predictions made by MU selection, then evidence is provided for H2. An overview of the phases of the game is given here, followed by details of each phase.

In phase one of the game, shown in fig. 1, subjects were trained in several rounds to fire missiles from a launcher in a 2D plane, in an unambiguous simulated 3D environment. They received rewards according to how close to aerial targets (shown as white crosses in the figure) they get. After demonstrating that they understand the utility function and controls by passing a second, ‘examination’ phase, they are then tested (phase three) in an ambiguous bi-stable environment. A true MEU strategy would consider both interpretations of this environment, whereas a UC based strategy would use only one and lead to a different action. Phase one consists of several rounds which teach the subject about the game. The choice of tasks here is fairly arbitrary, as the logic of the experiment is that *if* subjects can pass the phase two exam *then* they have demonstrated an understanding of the rules sufficient to play the real test game in phase three. Phases one and two may

¹in Python, source code available on request.

be repeated in any order as many times as the subject desires, until the exam is passed successfully.

Phase 1: exploratory training round

The game environment consists of a visible 2D horizontal plane on which a missile base can move around, a wire-frame cube in the 3D space above the plane, and one or two targets located at vertices of the cube. The environment is drawn using very strong perspective², and the vertices of the cube are connected vertically to the plane by lines to make their 3D locations unambiguous. In addition, edges of the cube are drawn with thick lines of different colours, producing additional disambiguation cues where one lines is seen to cross in front of another.

Subjects control the (x, z) position of a missile base using an analogue joystick (Logitech Extreme 3D Pro) and fire a missile by pressing and holding the joystick trigger. Once fired, the missile moves upwards (the y direction) until the trigger is released. The missile then explodes at position $\mathbf{m} = (m_x, m_y, m_z)$. $N \in 1, 2$ targets are present in the environment at positions $\mathbf{t}^i = (t_x^i, t_y^i, t_z^i), i = 1 : N$, and a Gaussian reward R is received and displayed centre-screen after the explosion,

$$R = \sum_{i=1}^N r_i, \quad (5)$$

$$r_i = 100 \times \exp - \frac{(\mathbf{m} - \mathbf{t})(\mathbf{m} - \mathbf{t})^T}{2\sigma^2}. \quad (6)$$

The spread, σ was fixed at a large enough value ($\sigma = 2\sqrt{3}/\sqrt{2\ln 2}$) so that if two targets are present, the score is always highest when firing at the point between them than when firing directly at one of the targets.

In the exploratory round, single targets are presented at different vertices of a fixed cube. Subjects have unlimited time to position the missile base and fire. They are then represented with a visual display of the reward, then the next target is presented. They are given 50 such targets to practice with, no cumulative score display, and are encouraged to experiment to learn about the rewards available at different distances from the target by an introductory message.

Phase 1: utility training round

In the exploratory round, subjects obtained high scores by firing as close to the target as possible. To help them learn about the shape of the Gaussian utility function, a series of rounds takes place in which fixed cubes are shown and the subject is asked to deliberately score *only* 50, 70 or 90 points. Thus they are encouraged to try aiming at locations at different distances from the target.

Phase 1: double target training

This round is similar to the first exploratory round, but uses two targets presented together at each trial. By construction

²OpenGL: gluPerspective(45, 1.0*width/height, 0.1, 100.0); gluLookAt(7,0.05,7, 0,-1.25,0, 0,1,0). Cube faces are 2*2 units.

(choice of σ), the optimal action is now always to aim at the point midway between the targets.

Phase 1: blackout training rounds

Two further training rounds take place. In the first, the lower half ($x+z > 0$) of the missile-launching area is ‘blacked out’. It is coloured red, and the joystick is unable to move into the red area. In the second round, the situation is reversed at the top half of the grid ($x+z < 0$) is blacked out. Subjects learn that that optimal strategy when faced with a target in the blacked out region is to fire from a position as close to it as possible that is on the centre line.

Phase 2: examination round

The purpose of the examination round is to demonstrate that the subject has learned the optimal actions for single and double targets, as well as possessing sufficient motor skills to control the game using the joystick. 20 trails are presented in rapid succession (one every 5s) and a cumulative score is maintained. If subjects fail to score 2170 points or more, they are sent back to the exploratory phases then made to repeat the examination (or allowed to leave the experiment). The qualifying score was chosen such that it can only be obtained by using the optimal strategy of aiming closer to the centre of each pair of double targets than to either individual target. Thus by passing the examination phase, subjects demonstrate knowledge of this strategy.

Phase 3: Ambiguous test round

In trials within this round, a bi-stable ambiguous (Necker) cube is presented to the subject very quickly, at random orientations. In some (80% of) trials there is one target at a random vertex. In others there are two targets, which may be at opposite (10% of trails) or non-opposite (10% of trials) vertices. The percept is made ambiguous by switching the projection from perspective to orthographic, dropping the vertex-to-plane cues, and drawing all edges in white to remove overlap cues. To motivate subjects, they are told that their cumulative score from all rounds of phase three will be their reported result, and that the subject with the highest reported result will receive twenty UK pounds in cash. (Subjects were undergraduate psychology students and were not paid otherwise; but received credit towards their degree for participating.)

Phase 3: Blackout test rounds

The ambiguous test round is repeated twice more, with blacked out near and far regions as in the learning phase.

Debriefing

It is crucially important that subjects do not become aware of the ambiguity, because this could allow high-level (Full MEU) reasoning to aim in the centre, and destroy any UC-revealing behavioural effects. For this reason, after phase 3, subjects were told about Necker cubes and asked if they were aware of the Necker ambiguity.




Figure 2: Example of a single-target trail, (x, z) plane. The viewer’s position is in top right corner. Blue crosses show the two ambiguous target positions resulting from a single target vertex on a Necker cube. The red cross shows the firing position. Black lines show the centre line and the two classification boundaries, dividing the launching area into near (top-right), centre and far (bottom-left) firing regions.

Processing

25 subjects were tested. Of these, 20 completed the exam and proceeded to generate data in the test phases. In debriefing, no subjects reported awareness of the ambiguity in the Necker cubes. For each trial, the 3D positions of both ambiguous locations of the target or targets were computed. This was achieved by transforming the x, z joystick co-ordinates into a horizontal and depth pairs, (h, d) , then flipping the depth coordinate,

$$\begin{bmatrix} h \\ d \end{bmatrix} = H \begin{bmatrix} x \\ z \end{bmatrix}, \quad (7)$$

where H is the Hadamard matrix,

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (8)$$

The complementary ambiguous location is thus

$$\begin{bmatrix} x' \\ z' \end{bmatrix} = H^{-1} \begin{bmatrix} h \\ -d \end{bmatrix}. \quad (9)$$

Furthermore, the height coordinate was transformed by $y' = y - cd$, where c is a constant ($c = 0.9$) which compensates for the choice of viewing angle in the projection images.

All shots were classified into three regions (fig 2), according to whether their (x, z) firing locations were closest to the near ambiguous target location, the far ambiguous target location, or the centre line.

Results

Non-blackout trials with single targets

In these trials, the MEU action by construction (i.e. choice of σ as described in phase 1) is to fire at the point on the centre line between the two possible ambiguous locations. (fig. 2). In contrast, the MU action is to fire directly at a

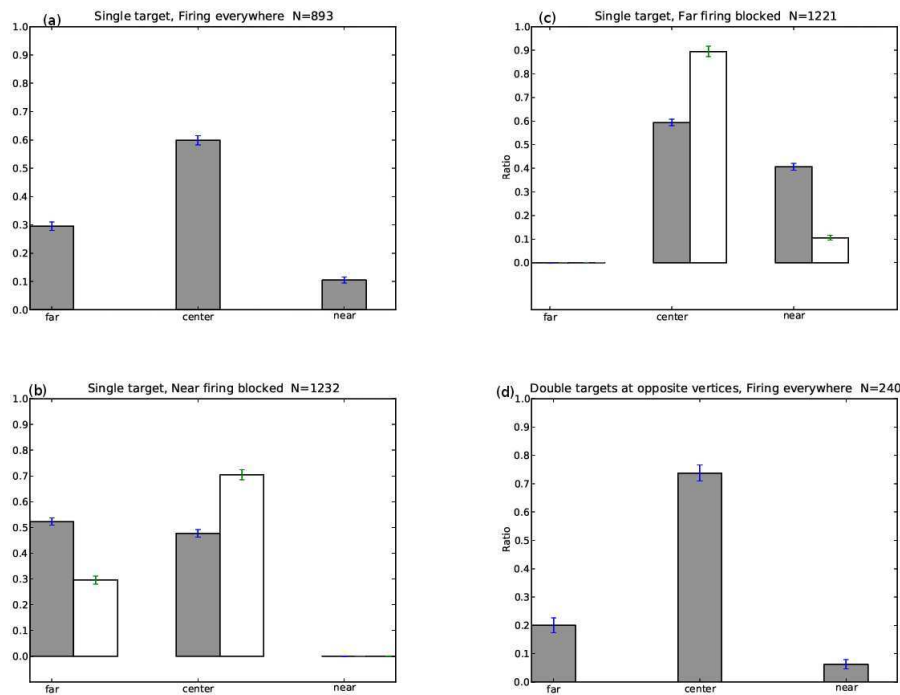


Figure 3: Results. Grey bars: observed frequency ratios, with Beta posterior, one-standard-deviation confidence intervals. White: predictions for blocked firings under the null hypothesis. The null hypothesis is that the unblocked area is unchanged from the 'single target, firing everywhere' case, while the centre is the sum of the centre and blocked firings from that case. All error bars assume IID observations and ignore which observation came from which subject.

randomly-chosen single one of those locations. This is because the MEU action averages over the two possible states of the world, which gives the same calculation as choosing where to fire in an unambiguous double target case; whereas the MU action picks just one interpretation of the target location, then fires directly at it.

Fig. 3a shows the distribution into region classes, over all subjects and trials of this type. Treating each action by each subject as an independent observation (i.e. ignoring subject-specific effects), and beginning with a flat Beta prior over the ratio of shots in each region, we infer posterior ratios, along with uncertainties. The figure shows the mean and one standard deviation error bars inferred about the population ratio of shots fired of each type. Signal detection theory can be used to obtain the p values, but broadly two ratios are significantly different if pairs of error bars do not overlap.

The results of these trials are surprising but inconclusive. Although target locations are perfectly ambiguous between near and far positions, subjects show a preference for the far target over the near one. That is, they are already interpreting the Necker cubes percept in a biased way, to favour interpretations with the target at the back of the scene.

If the MEU strategy was followed perfectly, we would see all shots fired in the centre and none in the near or far regions. If the MU strategy was followed perfectly, we would see all shots in the near and far regions and none in the centre. Unfortunately, we see shots in all three regions. Whilst this

is incompatible with a pure MEU strategy, it can weakly be explained from a MU perspective: A large number of shots are fired from the centre line, which may be due to subjects losing all depth perception (i.e. not perceiving the cube as 3D at all) and hedging by firing in the centre; they may also be due to limited depth perception resulting in a stable cube percept but an inaccurate joystick placement. (Some subjects commented on the lack of training in the absence of the vertical supports, and consequent loss of skill at pointing to the depth of the targets.) The near and far shots would be correct MU actions, and the centre shots due to a problem with the experiment, requiring a better communication of the depth to the subjects in future versions whilst retaining the ambiguity.

Blackout trials with single targets

In these trials, the pure MEU action is *still* to fire at the point on the centre line between the two possible ambiguous locations. Points on the centre line are still available during a blackout, so the optimal strategy is unchanged.

Fig. 3b shows the results when the near-side is blacked out. The majority of shots are now fired in the far region. This is consistent with the MU strategy: actually *perceiving* and acting on the Necker cube in the configuration which enables the target to be reached; an optimism bias. If we assume UC perception and action, these new results then show MU-like bias occurring within in. For comparison, we show in white the prediction of a null hypothesis. This is obtained taking

the Single target histogram of fig. 3a and moving its near mass to its centre mass, as would occur if UC percepts were unchanged by the utility bias, and near shots were substituted by firing on the centre line as close to blacked-out near targets as possible. MEU theory is unable to explain why the observed frequencies are so different from this null hypothesis. MU theory explains it easily: there is no utility in *perceiving* near targets; but if they are re-perceived as far targets, then an increased utility can be obtained by firing at them.

Similarly, fig. 3c shows the results when the far-side is blacked out. Again compared to a null hypothesis (white bars) which moves the mass from the far region to the centre region from fig. 3a, we see a significant difference, again suggesting the MU-like change in both percept and action towards the obtainable (non-blacked-out) target position.

Fig. 3d is shown as a control. It is the distribution from the non-blacked-out trials having two targets at opposite cube vertices. In these cases, the MEU and MU strategies are the same – fire in the dead centre of the grid, and the results show a significantly increased rate of firing in the centre region over that of fig. 3a. This again supports the presence of MU over MEU, because MEU would give identical results in 3a and 3d but MU would give an increase in the centre region in 3d over 3a, which we do see here.

Discussion

In informal discussions, we often argue that “perception is obviously UC from subjective experience”. Operationalists challenge this statement and would prefer us to cite an experiment to demonstrate the claim in the third person. While MEU actions are known to occur at both low-level psychophysical tasks and at high-level cognitive reasoning tasks, we have here presented evidence for the existence of a largely unexplored middle ground in which action selection is consistent with MU, and inconsistent with both Full deliberative and S-R approximate MEU, as in Hypothesis H1.

MU, but not MEU, can explain the deviations from the null hypotheses seen in figs. 3b and 3c, and also the difference between figs. 3a and 3d. The match of the data to MU behaviour in these cases gives some support for Hypothesis H2.

It was disappointing for the MU theory that fig. 3a did not present conclusive evidence by itself of MU over MEU, as pure MU would predict all shots to be fired in near and far regions, not in the centre. One way to explain the data away here is that the stimuli used were insufficiently informative to give subjects a sense of space, so they fire in the centre by default in the absence of any meaningful percept. Future work should try to refine the experiment to see if such hypothetical null percepts can be replaced by true percepts, for example by using different input and display systems while retaining the ambiguity in the Necker cube itself. Finally, the assumption that all shots by all subjects are mutually independent is strong, and future work should employ more subjects so that the assumption of independence of shots belonging to each subject can be dropped in the analysis.

References

- Bernardo, J., & Smith, A. (2000). *Bayesian Theory*. Wiley.
- Cooper, G. F. (1990). The computational complexity of probabilistic inference. *Artif. Intell.*, *42*, 393–405.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci*, *8*(12), 1704–1711.
- Doya, K., Ishii, S., Pouget, A., & Rao, R. P. N. (2007). *Bayesian brain*. MIT.
- Felzenszwalb, P. F., & Huttenlocher, D. P. (2005). Pictorial structures for object recognition. *International Journal of Computer Vision*, *61*(1), 55–79.
- Gigerenzer, G., & Todd, P. (1999). *Simple heuristics that make us smart*. OUP.
- Goldstein, D., & Gigerenzer, G. (2002). Models of ecological rationality. *Psych. review*, *109*(1), 75.
- Goodale, M., & Milner, A. (1992). Separate visual pathways for perception and action. *TINS*, *15*(1), 20–25.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psych. Science*, *17*, 767–773.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, *93*(5), 1449–1475.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge.
- Körding, K., & Wolpert, D. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.
- Kording, K., & Wolpert, D. (2007). *Bayesian statistics and utility functions in sensorimotor control*, in Doya et al.
- Lengyel, M., & Dayan, P. (2008). Hippocampal contributions to control: the third way. In *NIPS*. MIT.
- Pourret, O., Naim, P., & Marcot, B. (2008). *Bayesian networks: a practical guide to applications*. Wiley.
- Ramamoorthy, A., & Verguts, T. (2012). A computational model of instruction following. *Brain Research, pre-press*, doi:10.1016/j.brainres.2011.12.025.
- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*(11), 1019–1025.
- Spiegelhalter, D., Thomas, A., Best, N., & Gilks, W. (2008). *Bugs: Bayesian inference using Gibbs sampling*. Available from <http://www.mrc-bsu.cam.ac.uk/bugs/>
- Srinivasan, R., & Nunez, P. (2006). *Brain networks with distinct spatial and temporal structure*. Abstr. in Assoc. Scientific Study of Consciousness, Oxford.
- Stone, J., Kerrigan, I., & Porrill, J. (2009). Where is the light? Bayesian perceptual priors for lighting direction. *Proc. Royal Soc. B*, *276*(1663), 1797–1804.
- Young, S., Evermann, G., Gales, M., Hain, T., Kershaw, D., Liu, X., et al. (2006). *The HTK book*.