UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Prediction and learning under unsignalled changing contexts

Permalink
https://escholarship.org/uc/item/2790m865

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Wall, Laura
Gronau, Quentin
Cooper, Gavin
et al.

Publication Date
2023

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Prediction and learning under unsignalled changing contexts.

Laura Wall [1], Quentin F. Gronau[1], Gavin Cooper[1], Guy Hawkins[1], Scott D. Brown[1], and Juanita Todd[1]

School of Psychological Science, University of Newcastle, Australia
Callaghan, NSW 2308 AUS[1]

Abstract

Predictive inference and error-driven learning are critical to optimal performance across many different contexts. However, the specific context determines the informativeness of errors in updating predictions. In this study, participants experienced two changing, unsignalled contexts with opposite optimal responses to errors; the change-point context, where errors were informative, and the oddball context, where they were not. The changes to the context occurred under two task structures: 1) a fixed task structure, with consecutive blocks of each context, and 2) a random task structure, with the context randomly selected for each new block. We modelled participants' performance using a Hierarchical Gaussian Filter (HGF) model. We found that performance was greater in the oddball than change-point context, with more accurate and precise estimates. The estimates from the fixed task structure were also more precise than those in the random task structure. We showed that consistency in context can improve precision.

Keywords: predictive inference; learning; context; Hierarchical Gaussian Filter

Introduction

Predictive inference, the prediction of future observations based on past observations, underpins much of everyday learning and behaviour and is critical to optimal performance across a variety of different contexts. Further, many systems are set up under the assumption that past observations are good predictors of future ones. For example, in criminal sentencing, the judge predicts the likelihood of future criminal activity by past criminal activity; in real estate, house prices are set based on previous sale prices in the area. However, all of these observations occur with a certain level of noise or variability; e.g. some houses sell for more than expected and some less. This variability introduces *uncertainty* into the estimation of the underlying "truth" – the actual probability that the perpetrator will commit a future crime, or the real value of a house in a given area.

However, sometimes, the uncertainty is amplified by outliers – one-off, extreme observations beyond the usual variability, that do not reflect the underlying truth. For example the sale of a mansion or run-down shack amongst suburban two bedroom homes. With enough samples and prediction errors from the environment, outliers are detected for what they are, the uncertainty is reduced, and the estimate of the underlying truth is improved. However, in most contexts, this underlying truth is not always stable and can change – a previously "good citizen' gets in the with "wrong crowd" and starts regularly committing crimes, or a noisy nightclub gets built in the area and house prices drop. These changes reflect *volatility*. Optimal performance of predictive inference under usual uncertainty, outliers, and volatility requires quick and accurate determination of whether a particular observation, e.g. a low house price, reflects the natural variability in observations, a surprising but uninformative outlier, or a change in the underlying value.

There has been considerable theoretical, experimental and computational development on this topic, improving the understanding of the mechanisms underpinning optimal learning and predictive inference under uncertainty and volatility (Nassar, Wilson, Heasly, & Gold, 2010; Nassar, Waltz, Albrecht, Gold, & Frank, 2021; Bruckner, Heekeren, & Nassar, 2022; Brown & Steyvers, 2009; d'Acremont & Bossaerts, 2016; Marković & Kiebel, 2016; Mathys et al., 2014). The predictive inference tasks used across these papers share many common features. Typically, participants are required to predict the next item in a series of samples which are drawn from a Gaussian distribution with a mean reflecting the underlying state and the standard deviation creating some estimation uncertainty. In these tasks, under the standard context, the mean is stable for most trials, but occasionally changes (referred to as a change-point) at a probabilistic rate (referred to as the hazard rate). Some studies also employ an oddball context (Nassar, Bruckner, & Frank, 2019; Nassar et al., 2021; d'Acremont & Bossaerts, 2016), where the mean moves very slightly in accordance with a random walk on each trial. Occasionally an outlier is sampled outside of the distribution, at the same probabilistic rate as change-points. In both contexts the change-points and outliers lead to surprising events, with large prediction errors. In the changing context, these errors help the participants to register the change and *update their predictions* accordingly. In the oddball context, these large errors help participants to identify the outliers and *ignore* them.

Optimal performance differs between the two contexts, and thus to perform well participants are required to know what context they are in, and how to respond to surprising events. In previous studies, these contexts have occurred, and been modelled, in distinct blocks where participants were made aware of a context change through instruction. However, in reality people are often required to infer what context they are in, and if it has changed, by the behaviour of the samples. This requires longer-term learning where information from multiple surprising and unsurprising observations is integrated into a hierarchical model which informs not only the variability and mean of the current underlying distribution,

[1]Corresponding author: please address any enquiries to laura.wall@newcastle.edu.au

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*. ©2023 The Author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

but also the current context. Such longer-term learning has only just begun to be examined in this field, despite clear implications for interpreting learning deficits across different groups (Wall, Cooper, Hawkins, Brown, & Todd, 2023).

If the explicit instruction separating these contexts were removed, participants would only be able to apply a context-appropriate response to surprising events if they had *learnt* what context they were in. In this study we will include multiple short blocks of each context, but without explicit information about which context is occurring in each block, or when a change in context will occur. This will provide an opportunity to examine the extent to which participants learn over a longer time-frame. However, this continued learning is likely contingent on key aspects of the environment being consistent across these longer time-frames. To further investigate the conditions of longer-term learning, we will manipulate the presentation of the multiple short blocks of each context to create two task versions; one where all blocks of each context are presented consecutively and one where the context randomly changes from one block to the next. Theoretically, participants in the consecutive contexts version should have better estimates of the underlying distribution and context, than those in the randomised version, as they have more opportunity to refine their model of the task environment and its predictions before the context changes. If one were to analyse performance in this task with a computational model, these improved estimates of the underlying generative states should translate into less variable and more accurate parameter estimates of the mean.

Although a variety of computational models have been used to understand learning in this style of predictive inference task, the two main theoretical models are Nassar et al's. (2010) normative model based on reduced Bayesian belief updating and delta rule models, and the Hierarchical Gaussian Filter model (HGF; Mathys et al., 2014). Both models estimate aspects of uncertain and volatile environments, such as the mean and variance of the generative distribution, and the volatility of the changes in the mean. For a comparison and description of the differences between these two main models see Marković and Kiebel (2016) and Bruckner et al. (2022).

In our previous work (Wall et al., 2023), (author et al, 2022)[2], we utilised this multi-context experimental design and implemented with a modified version of the normative model developed by Nassar et al. (2010, 2021) to understand participant learning over time and contexts. However, there was no obvious differentiation in the parameter estimates across the two task versions (consecutive and random). Whilst we could continue to modify this model to better accommodate longer-term context learning, we propose that implementing the HGF, which is more flexible and task agnostic, has the potential to provide better estimation in this multi-context design. Further, given that longer-term context learning would naturally lead to a hierarchical internal model of the task environment for participants, the hierarchical structure of the HGF should better capture this process. Therefore, in this paper, we will use the data from Wall et al. (2023) but utilise the HGF model to determine if the estimates of the mean of the distribution are less variable and more accurate for participants in the fixed task version than those in the random task version.

Methods

The detailed methods can be found in Wall et al.(2023), however, they are noted briefly below.

Participants

There were 101 (51 female) participants in total. All participants were recruited from an online marketplace, Prolific, in two separate samples. The first were reimbursed at £8.16/hr (approx 15 AUD/hr) while the second were reimbursed at £7.52/hr (approx 14 AUD/hr) due to an update to the estimated time to complete the task. The mean age of participants was 25.54 (median= 22, $SD = 7.9$; range = 19 - 68). The sample was generally experienced with online studies and tasks but there was substantial variability.

Design

This study was a 2 (context: change-point v oddball) x 2 (task version: fixed, consecutive contexts v randomised contexts) mixed subject design, with all participants receiving both contexts, but the task version randomised across participants.

Task

The predictive inference task asked participants to catch bags of money (for points) dropped from a helicopter, which was hidden behind clouds. The bag locations were drawn from a Gaussian distribution with the helicopter's current location as the mean. The goal was to guess the location of the next bag drop and move a bucket to the predicted location, through key presses. An example task sequence can be seen in Figure 1.

The task had two contexts which produced surprising locations of bag drops, as seen in Figures 1 and 2. In both contexts these surprising drops (change-points and outliers) occurred on 12.5% of trials, with a minimum of three trials between each surprising drop. In the change-point context, change-point trials were defined by a randomly selected change to the mean (helicopter) position. In this context the mean was stable for all non change-point trials. In the oddball context, the mean drifted in accordance with a random walk of mean 0, SD 1, on each trial, but surprising outlier trials were occasionally randomly drawn from a uniform distribution the width of the screen. The outlier trials had no influence on the random walk motion of the helicopter. These contexts were explained to participants in the instructions as two different levels of wind. In the oddball context, it is windy, so it is hard to keep the helicopter stable (explaining the mean drift),

[2]This study is an extension of our work on this paper - to ensure blind reviewing but representative reporting, we have noted that this previous work is our own but redacted our names. All references to author et al, 2022 refer to the same previous paper

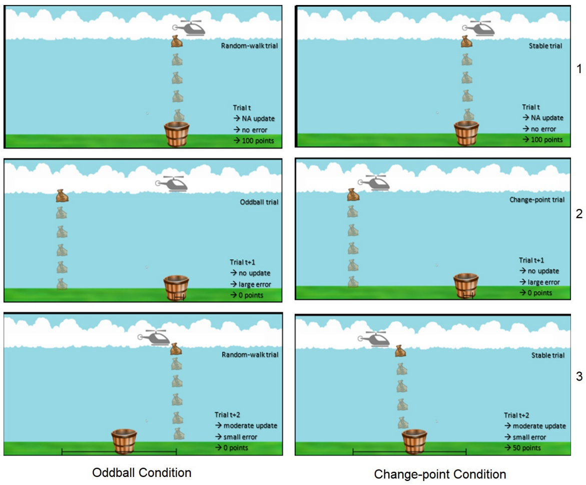

Figure 1: An illustration of an example 3 trial sequence for the oddball context (left panel) and change-point context (right panel). NOTE: The background image, bucket, non transparent bag and line at the bottom are all as seen by the participant. The text, helicopter and transparent bags are added for illustrative purposes.

it is dangerous to fly the helicopter (explaining the lack of change-points) and occasionally a big gust of wind blows a bag further away than expected (explaining the outlier). In the change-point context, it is calmer, so the pilot can control the helicopter better and keep it stable (explaining the lack of drift) and fly it to different locations (explaining the change-point), and there are no big gusts of wind (explaining no outlier drops). It was made clear to participants that in both contexts a drop would occur that was unexpected, but that it meant something different depending on the context: if windy, ignore it, it's an outlier; if calm, move your bucket, it's a change-point.

Procedure

The experimental session took approximately 20 minutes and involved eight blocks of 50 trials each (four for each context), as seen in Figure 2. The order of the blocks was manipulated across subjects, such that approximately 50% of subjects (N = 52) were allocated to a "consecutive context blocks" version of the task, and the remainder allocated to the "random context blocks" version. As seen in the top row of Figure 2 in the consecutive version, the task was essentially split into two halves of four blocks each, with one context as the first half and the other as the second. The presentation of these 'halves' was randomised across participants such that each order occurred around 50% of the time. As seen in the bottom row of Figure 2, in the random version of the task, the context randomly changes from one block to the next. All order combinations of the blocks were equally possible.

Modelling

We used the Hierarchical Gaussian Filter model (HGF; Mathys et al., 2014) with two levels as described in Marković and Kiebel (2016). The HGF specifies the following generative model for the observed bag location on trial t, denoted by o_t:[3]

$$\begin{aligned} x_t^{(2)} \mid x_{t-1}^{(2)} &\sim \mathcal{N}(x_{t-1}^{(2)}, \eta) \\ x_t^{(1)} \mid x_{t-1}^{(1)}, x_t^{(2)} &\sim \mathcal{N}(x_{t-1}^{(1)}, \exp\{x_t^{(2)}\}) \\ o_t \mid x_t^{(1)} &\sim \mathcal{N}(x_t^{(1)}, s). \end{aligned} \tag{1}$$

The hidden state $x_t^{(1)}$ corresponds to the true helicopter location on trial t and the hidden state $x_t^{(2)}$ to the volatility of the helicopter location on trial t. Since participants do not know the true states of the world, they need to invert this generative model to obtain a *perceptual model*. Mathys et al. (2014) described how this can be achieved using Bayesian Variational Inference. This way, participants can obtain an estimate of the posterior mean, $\mu_t^{(1)}$, and posterior variance, $\sigma_t^{(1)}$, for the hidden state $x_t^{(1)}$ on each trial. Similarly, they can obtain an estimate of the posterior mean, $\mu_t^{(2)}$, and posterior variance, $\sigma_t^{(2)}$, for the hidden state $x_t^{(2)}$ on each trial. We assume that participants' bucket placement on trial t, denoted by y_t, is then generated as follows:[4]

$$y_t \sim \mathcal{N}(\mu_t^{(1)}, \sigma_t). \tag{2}$$

To capture both the accuracy (absolute discrepancy) and certainty (precision) of participant's estimates of the mean, we calculated a new variable as an indicator of performance. We denote this variable as π, but refer to it as the relative discrepancy. To calculate this variable, for each trial, the mean discrepancy – difference between the mean $x_t^{(1)}$ (the actual mean of the generative distribution) and $\mu_t^{(1)}$ (participants estimate of that mean) – was divided by the precision of that estimate – $1/\sigma_t^{(1)}$. This results in a variable where a *smaller* value - a smaller relative discrepancy - is indicative of *better* performance in both the absolute discrepancy (smaller mean difference) and the certainty (more precise/smaller variability). The formula penalises participants who are either very certain but wrong, or very accurate but uncertain.

$$\pi = \frac{|x_t^{(1)} - \mu_t^{(1)}|}{1/\sigma_t^{(1)}} \tag{3}$$

The model was fit to the participant data. The model's fit to the data was assessed through visual inspection. The model code, data, and plots of each participant's actual (data) and

[3]Note that the HGF described in Marković and Kiebel (2016) is a special case of the original HGF described in Mathys et al. (2014). Specifically, they implicitly assume that the intercept in the term describing the variance of $x_t^{(1)}$ as a function of $x_t^{(2)}$ is zero and the corresponding slope is 1.

[4]This part is often referred to as the *response model*.

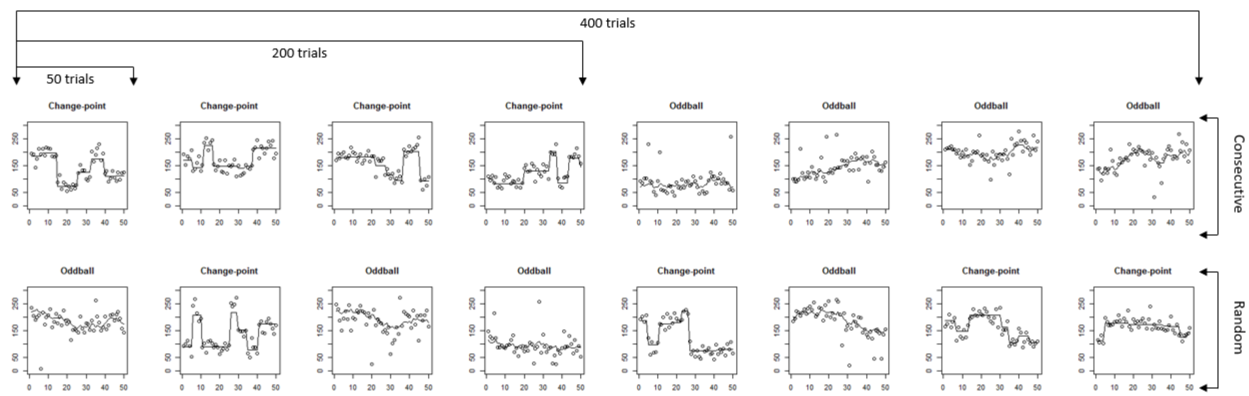

Figure 2: An example block sequence for the consecutive task version (top row - participant 92) and the randomised task version (bottom row - participant 93). The eight squares from left to right represents the experimental blocks 1-8, each of 50 trials, with the dots showing the position of the bags, and the lines showing the position of the helicopter, on the 0-300 unit screen width.

predicted (model generated) bucket placements can be found at `https://osf.io/mf5ne/` One of these plots for an example participant and block for each context is shown in Figure 3.

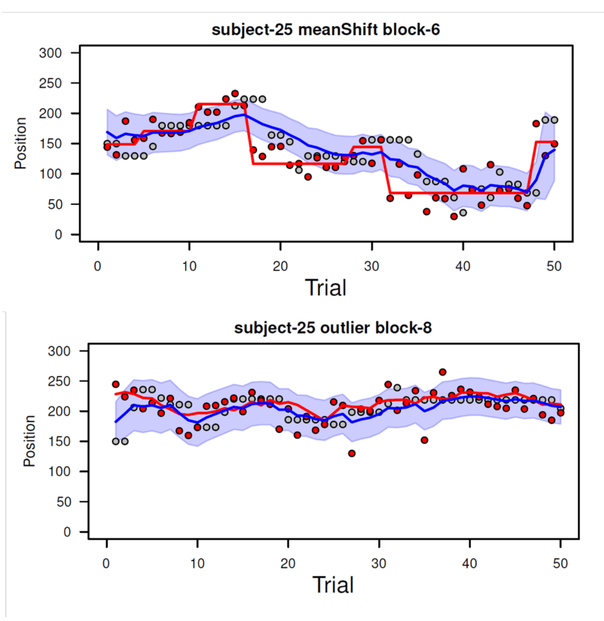

Figure 3: An example block of each context (block 6, change-point and block 8, outlier) for an example participant (25). The red lines indicate the true mean ($x_t^{(1)}$), the blue lines indicate the participant's estimate of the mean ($\mu_t^{(1)}$) and the shaded blue area is a 95% credible interval. The grey dots are the participant's bucket placements and the red dots are the bag drops.

Results

Extensive model free analysis was reported in Wall et al. (2023) and so in this paper, we will provide a brief summary of those results for context, but only report on results from our HGF analysis.

In summary, we found in Wall et al. (2023) that participants generally responded to the task as expected, where larger updates followed larger errors across trials, and bigger proportional updates (update / error) followed change-point trials compared to standard trials, while smaller proportional updates followed outlier trials. These differences, between responses following change-points and outliers, increased across trials within a block, and the mean of these responses generally increased across blocks, but only for participants in the fixed task version. We suggested that the refinement to one's internal model across a longer period of time was likely dependent on the broader environmental stability.

The means of our HGF model-estimated individual, trial by trial, computed variables reflecting the absolute discrepancy ($x_t^{(1)} - \mu_t^{(1)}$), precision ($1/\sigma_t^{(1)}$) and relative discrepancy (π) for those in the fixed and random task version are seen in columns 2 & 3 and 4 & 5 of Table 1 respectively. These variables, summarised at the participant level across trials, can be further seen in Figure 4.

As seen in Figure 4 and column 2 of Table 2, two sided Bayesian t-tests found that there was decisive evidence of differences in all variables between the two contexts, such that there was a smaller difference in the estimated and actual mean location (absolute discrepancy), greater precision, and also better overall performance due to smaller relative discrepancy for the oddball than change-point context.

As seen in column 3 of Table 2, Bayesian one-sided t-tests found *no* evidence of differences between the two task versions for absolute or relative discrepancy. However, despite the small size of the difference in precision, the Bayesian t-tests showed decisive evidence of higher precision for the fixed, consecutive task structure, than the randomised task structure. Similarly, Bayesian ANOVAs of the effect of context and task version on absolute and relative discrepancy both found the model with the strongest evidence as the one with an effect of context only (winning model $BF_{Acc} = 3.39 \times 10^{123}$ and $BF_{Per} = 3.43 \times 10^{64}$ respectively). However, the ANOVA on precision found the model with the greatest evidence as the one which included context, task version and the interaction between the two (winning model $BF_{Pre} = 2.47 \times 10^{20}$).

Variable	Fixed CP	Fixed OB	Random CP	Random OB
$x_t^{(1)} - \mu_t^{(1)}$ Absolute discrepancy	16.978	13.244	16.744	13.339
$1/\sigma_t^{(1)}$ Precision	0.063	0.065	0.055	0.061
π Relative discrepancy	341.295	271.006	327.603	281.689

Table 1: Means of the individual, trial-by-trial, computed variables for the fixed (column 2 & 3) and random (column 4 & 5) task version, for both contexts (CP; change-point, and OB; oddball). $x_t^{(1)} - \mu_t^{(1)}$ reflects participants' accuracy in terms of absolute discrepancy through the difference in their estimated mean and the true helicopter mean, $1/\sigma_t^{(1)}$ reflects participants' certainty in terms of precision through the variability of their estimate of the mean; π reflects their performance in terms of relative discrepancy, incorporating both the absolute discrepancy and precision. NOTE: *lower* values of absolute and relative discrepancy indicate *better* performance as they represent a *smaller* difference between the estimated mean and the actual mean.

Variable	Bayes Factor CP v OB	Bayes Factor Fixed v Random
$x_t^{(1)} - \mu_t^{(1)}$ Absolute discrepancy	6.58×10^{121}	0.0170
$1/\sigma_t^{(1)}$ Precision	2.55×10^{8}	1.169×10^{35}
π Relative discrepancy	5.84×10^{59}	0.0165

Table 2: BFs from a t-test of differences in the means between the two conditions are provided in column 2 - these means are for all trials in the change-point condition compared to all trials in the oddball condition, regardless of task version. BFs from a t-test of differences in the means between the two task versions are provided in column 3 - these means are for all trials in the fixed condition compared to all trials in the random condition, regardless of context.

Discussion

Summary

We found that participants could perform well in the task, with generally small differences between their estimate of the mean and the actual mean, regardless of which context they were in or which task structure they were assigned to. Interestingly, the oddball context resulted in better performance (smaller relative discrepancy) than the change-point context, both in terms of absolute discrepancy and precision. When considering the experimental uncertainty present in the two contexts, one might expect the change-point condition to be easier, as the mean location remains stable and consistent across multiple trials before changing, whilst the mean location changes randomly on each trial for the oddball condition. A mean that has longer periods of stability should be easier to estimate, particularly towards the end of that period. However, this is under the assumption that the agent and process of estimating the mean are aware of this stability and will leverage it accordingly. If the model assumes that the mean changes and thus must be updated each trial, then data which are generated by such a changing mean are now more likely to result in accurate and precise estimates, compared to data generated by a different process. As seen in the example participant's data in Figure 3, the model generated μ follows a

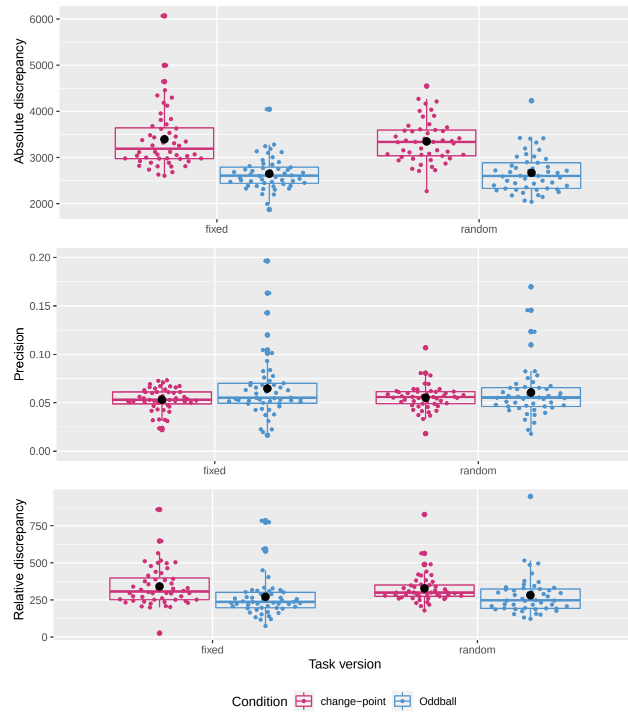

Figure 4: The trial-by-trial variables noted in Table 1 summarised at the participant level, across contexts (Condition; change-point = pink and left, oddball = blue and right) and task versions (fixed = left column, random = right column). Absolute discrepancy ($x_t^{(1)} - \mu_t^{(1)}$; top row) is summed across trials, whilst Precision ($1/\sigma_t^{(1)}$; middle row) and Relative discrepancy (π; bottom row) are the means across trials. The dots are quasi-random samples from the distribution of data; black dots represent the mean, whilst the box plots represent the median and inter-quartile range.

similar pattern over trials across the two contexts, despite the experimentally generated mean having two distinctly different patterns (one with steps between stable periods, one in accordance with a random walk).

It is important to remember that whilst participants knew there were periods where the helicopter remained stable for a while, they did not know when those periods were. It is possible, that they may have opted for one strategy for estimating the helicopter across all contexts, with that strategy more aligned with the oddball context. In our previous work using this paradigm, we showed that participants were sensitive to the context changes such that they were capable of differentially responding to large errors based on the context (Wall et al., 2023). However, participants could have utilised one strategy for estimating the mean regardless of context, and then used another strategy for determining what to do with large errors. The inclusion of a mixture process in the response portion of the model could help to investigate this possibility in the future.

The main aim of this study was not to investigate differences in performance across contexts, but rather differences in performance across task structures, to examine the possible influence of longer-term context learning. Our computed measure of performance (relative discrepancy, π), did not show any group differences, nor did the difference between the estimated and generative means. These results are in line with our previous modelling (Wall et al., 2023), where we did not find any group differences in the parameters of our modified version of the normative model developed by Nassar et al. (2010). The measure of precision, computed from the HGF parameter σ, however, did find group differences, in the direction we hypothesised, with greater precision for participants given the consecutive fixed task structure, than the randomised structure. It is worth highlighting, that although the Bayes Factors for the ANOVAs and t-tests were decisive (and were supported by frequentist analysis), the effect size of the group difference was very small, and difficult to observe visually. It is questionable then, how useful such a difference might be when extrapolating the findings to more real-world contexts.

Limitations and future directions

Whilst the HGF is a natural fit for the question posed in this paper about higher-order learning, this analysis did not fully utilise the hierarchical structure of the task environment by only having two levels and by only examining the outputs from the first level. In this study, we have established that the participant's estimate of their mean (μ_1), and the variability (σ_1) and precision ($1/\sigma_1$) of that estimate are useful indicators of differences in performance between different experimental conditions and contexts. Whilst it was beyond the scope of this paper, investigation of differences in participants' estimates of the volatility of the helicopter for each trial, could potentially be more sensitive to differences between the two task versions. Further, the model could be extended to include a third level, which may help to estimate the longer term volatility of the context. Future research could compare the fit and outputs of a 2 vs 3 level HGF.

Whilst these are interesting avenues for future research, this application of a simple two-level HGF to this data is still a worthwhile contribution. Most applications of the HGF involve much simpler data, often with binary stimuli and responses, and with more explicit conditions, which minimises the uncertainty present for participants and produces cleaner data. This study provided an important test that the HGF could be applied to this task where participants were not explicitly told what context they were in, and therefore, had a wide scope of possible strategies they could have employed to complete the task under this additional uncertainty, resulting in richer, but more difficult to model, data.

Also, although the HGF is a promising model for these types of tasks, there are other models we could have chosen to look at learning over time. We could have used a particle filter model, as used in Brown and Steyvers (2009), where the number of particles in the filter represents the fidelity of one's hypothesis space about parameters of the task environment. We could similarly use a sampling-based approach to determine the influence of samples from past experiences, such as used in Bornstein, Khaw, Shohamy, and Daw (2017). An interesting avenue for future research would be to compare these different models on different variants of tasks such as this, to determine if different models better fit different aspects of learning and prediction under uncertainty.

Conclusions

In conclusion, we showed that the HGF is a suitable model for examining learning across changing contexts and that whilst consistency in context does not influence the accuracy of performance, it can produce more precise estimates within that context.

Acknowledgements

This research was partially supported by the National Health and Medical Research Council (NHMRC) Ideas grants scheme.

References

Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, *8*(1), 15958.

Brown, S., & Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, *58*, 49–67.

Bruckner, R., Heekeren, H. R., & Nassar, M. R. (2022). Understanding learning through uncertainty and bias. *PsyArXiv*.

d'Acremont, M., & Bossaerts, P. (2016). Neural mechanisms behind identification of leptokurtic noise and adaptive behavioral response. *Cerebral Cortex*, *26*(4), 1818–1830.

Marković, D., & Kiebel, S. J. (2016). Comparative analysis of behavioral models for adaptive learning in changing environments. *Frontiers in Computational Neuroscience*, *10*, 33.

Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the hierarchical gaussian filter. *Frontiers in human neuroscience*, *8*, 825.

Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related eeg signals and learning. *Elife*, *8*, e46975.

Nassar, M. R., Waltz, J. A., Albrecht, M. A., Gold, J. M., & Frank, M. J. (2021). All or nothing belief updating in patients with schizophrenia reduces precision and flexibility of beliefs. *Brain*, *144*(3), 1013–1029.

Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, *30*(37), 12366–12378.

Wall, L., Cooper, G., Hawkins, G., Brown, S., & Todd, J. (2023, Jan). *Consistency is the key! learning to adapt in a multi-context predictive inference task.* PsyArXiv. Retrieved from `psyarxiv.com/pu7dm`