Age-related Differences in Learning to Selectively Attend

Angela Radulescu
Princeton Neuroscience Institute
Princeton University
Princeton, NJ 08540
angelar@princeton.edu

Reka Daniel
Princeton Neuroscience Institute
Princeton University
Princeton, NJ 08540
rekad@princeton.edu

Yael Niv
Department of Psychology
Princeton Neuroscience Institute
Princeton University
Princeton, NJ 08540
yael@princeton.edu

Abstract

When confronted with many stimuli in a complex world, how does the aging brain learn where to focus attention? Previous work shows that older adults have more difficulty switching between different task-relevant dimensions. It remains unclear, however, whether and how the cognitive strategies they use differ from those employed by younger adults. Here we focus on age-related differences in the dynamics of representation learning, where participants learn which stimulus features are relevant to each task through trial and error. We compare the behavior of older and younger adults in a multidimensional reinforcement learning task designed to study how subjects update their representations on-line, and propose a series of models that implement various forms of selective attention. Model-based analysis of choice patterns shows that both younger and older adults employ attention during learning. However, older adults seem to maintain a narrower attentional filter, a cognitive strategy that might reflect an adaptation to changes in the interaction between the dopaminergic system and the prefrontal cortex.

Keywords: aging, representation learning, selective attention

Acknowledgements

This research was supported by a Ellison Medical Foundation New Scholar in Aging award to YN. We thank Yuan Chang Leong and Andra Geana for invaluable discussions.

1 Introduction

Older adults typically encounter difficulties with a variety of tasks broadly classified as indicators of cognitive flexibility [1]. One prevailing explanation for these multi-domain cognitive deficits is that with age, it becomes more difficult to represent, update and maintain task-relevant information [2]. In line with work that charts the decline in the dopaminergic (DA) system [3] and the deterioration of frontal lobes during healthy aging [4], it has been suggested that changes in the interaction between DA and the prefrontal cortex (PFC) can account for observed differences in attentional modulation [5] and inhibition of irrelevant stimuli [6]. Recent fMRI studies have also begun harnessing the explanatory power of reinforcement learning (RL) theory [7] to examine age differences in learning scenarios thought to be DA dependent. For example, work by Chowdhury and colleagues demonstrated that trial-by-trial BOLD dynamics in a probabilistic learning task are sensitive to pharmacological manipulations of DA levels [8]. Yet so far, surprisingly few RL models have directly explored the hypothesized functional link between attention and learning in the context of aging. This is especially limiting because dissociating neural substrates of prediction errors from those of attentional shifts is not always straightforward [9].

We previously proposed that in order to reduce computational costs, the brain might restrict learning to only a few relevant aspects of a stimulus [10], and suggested that attention processes shape a useful (if not statistically optimal) filtered representation of the world. Notably, this attentional filter is itself dynamically modulated by feedback, a process termed *representation learning* [11]. Here we develop a class of simple RL algorithms designed to ask whether aging brings about a shift in representation learning strategies. We show that in older adults, models that favor more focused representations give a better account of behavioral choice data. Older adults seem to rely on a narrower attentional filter that is less sensitive to new feedback, perhaps to compensate for reduced dopaminergic efficiency.

2 Participants and methods

To study age-related changes in strategies of on-line attention allocation, we tested 25 young (mean age = 23.2, range = 18 - 35) and 25 older adults (mean age = 69.3, range = 65 - 80) on a previously developed ”Dimensions Task” [10]. On each trial of the Dimensions Task (Fig. 1A), participants selected one of three visual stimuli differing along three dimensions (color, shape and texture). At any time point, one ”target” feature within one of the three dimensions was more rewarding than the others. Participants were made aware of this in the instructions. If the participant chose the stimulus containing the target feature, he or she had a 75% chance of receiving 1 point (or 0 points otherwise). If the chosen stimulus did not contain the target feature, the probability of obtaining 1 point was 25%. In order to maximize the number of points earned, subjects had to learn the identity of the target feature and use it to select the correct stimulus on each trial. To acquire repeated measurements of learning within each subject, we divided the task into 10 games (Fig. 1B), each consisting of a minimum of 8 and a maximum of 25 trials. We defined a correct trial as one in which the subject chose the stimulus containing the target feature. Once a minimum of 8 trials and a criterion of 6 consecutive correct trials was reached, the game had a 50% chance of ending on any subsequent trial (games ended after 25 trials even if criterion was not reached). At the end of a game, the participants were given a self-paced break and were notified that the target feature has now changed.

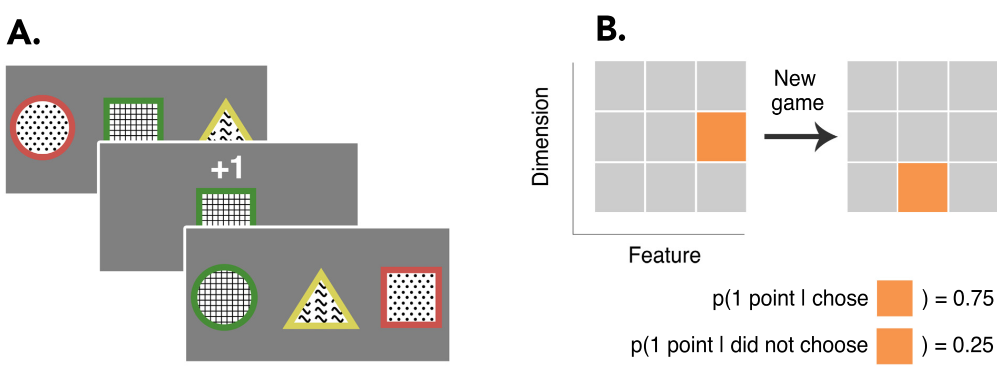

Figure 1: *The Dimensions Task*. A. Subjects chose between stimuli that differed in color, shape and texture. B. In each game, they had to learn by trial and error which of the 9 feature-dimension pairs is the highly rewarding one.

3 The family of models considered

Of particular interest to the comparison between younger and older adults are reinforcement learning models of the Dimensions Task that simplify the state space using *function approximation* (FA) [10, 12]. These models share a common architecture in that they calculate the value of each stimulus as a linear combination its feature values. For our purposes, we can incorporate different assumptions about selective attention into the function approximator and use trial-by-trial model fitting [13] to precisely test whether varying the width of the attentional filter is more likely to improve fits in older adults.

3.1 Reinforcement learning with uniform attention

The simplest variant that we considered assumes constant, uniformly distributed attention. We define Φ as a vector of dimension weights

$$\Phi = \begin{array}{cl} \begin{bmatrix} 0.33 \\ 0.33 \\ 0.33 \end{bmatrix} & \begin{array}{l} color \\ shape \\ texture \end{array} \end{array} \tag{1}$$

that determines both attention for choice, that is, how feature weights (w) combine to determine $V(S)$, the value of a stimulus S,

$$V_t(S) = \sum_{d=1}^{3} w_{t,d}\Phi_d \tag{2}$$

and attention for learning, that is, how feature weights of the chosen stimulus are updated:

$$w_{t+1,d}(\text{chosen}) = w_{t,d}(\text{chosen}) + (\frac{1}{2t} + \eta)(r_t - V_t(\text{chosen}))\Phi_d \tag{3}$$

where $\delta_t = r_t - V_t(\text{chosen})$ is the reward prediction error, η is fit and the effective step size $(\frac{1}{2t} + \eta)$ decreases over the course of a game. The idea here is simple: learn about chosen features, equally distributing attention over dimensions.

3.2 Reinforcement learning with value-driven attention

This model is identical to the above model, except in how the dimension weights are computed. Instead of assuming constant uniform weighting, Φ_d is updated on every trial proportional to how much the feature in that dimension contributed to the value of the previously chosen stimulus:

$$\Phi_{d,t+1} = \frac{e^{\theta w_{t,d}(chosen)}}{\sum_{i=1}^{3} e^{\theta w_{t,i}(chosen)}} \tag{4}$$

θ is a free parameter that determines what proportion of the attention is dedicated to the previously highest-valued feature, that is, how narrow the attentional filter should be. A large θ is equivalent to a narrow filter that heavily biases both choice and learning in the direction of dimensions containing features with high values.

3.3 Reinforcement learning with value-driven attention and decay

We add one additional mechanism for implementing selective attention by introducing the decay of features that are not consistently chosen. The idea here is to learn about chosen features, and uniformly decay the rest towards 0 at a fit rate of η_k. Thus in addition to Eq. 3 we have:

$$w_{t+1,d}(\neg chosen) = (1 - \eta_k)w_{t,d}(\neg chosen) + \eta_k \times 0 \tag{5}$$

As specified in Eq. 2 above, attention for choice still favors those dimensions that contain features with higher values. This is, in some practical senses, similar to the *serial hypothesis testing* strategy described in [11] in that it favors relying on consistently chosen features. However, it also affords "incidental learning" (i.e., if choosing a red square with polka dots based on red, the model also updates the values of polka dots and square).

3.4 Model fitting and comparison

For all models, a "softmax" policy with a fit inverse temperature β was used to compute the probability of making a particular choice on each trial:

$$\pi_t(chosen) = \frac{e^{\beta V_t(chosen)}}{\sum_{i=1}^{3} e^{\beta V_t(i)}} \tag{6}$$

Model parameters were optimized by maximizing the log likelihood of the participant's data given the model. These parameters were then used to compute the Bayesian Information Criterion (BIC) approximation of model evidence [14], E_M:

$$E_M \approx log(p(D|M, \hat{\theta}_M) - \frac{||\hat{\theta}||}{2} log\ N \tag{7}$$

where $p(D|M, \hat{\theta}_M)$ is the likelihood of data D given model M and parameters $\hat{\theta}_M$, $||\hat{\theta}||$ is the number of free parameters in the model and N is the number of data points (trials). To provide a more intuitive measure of model evidence, we divided the total score for each model by the number of trials for which a participant provided a response, and exponentiated it, to yield a complexity-corrected average likelihood per trial that varies between 0 and 1.

4 Results

The two groups were evenly matched on overall performance in the Dimensions Task: younger adults learned 48% of games on average, while older adults learned 42% of games (paired-sample t-test, t(48) = 1.27, p = .21). Mean reaction times were significantly higher in older than in younger adults ($.89s \pm .006$, $.65s \pm .008$, paired-sample t-test, t(48) = 24.3, p<.001). To investigate whether older adults learn using more focused representations, we fit computational models to choice data. Collapsing over the two groups, we find that assuming a narrower attentional filter improves fits in general (Fig. 2A): with a mean BIC-corrected per-trial likelihood of 0.55 (SEM = .008), the "attention & decay" model's predictions were significantly above chance (t(98) = 30.06, p $<$ .001), as well as above those of the "uniform attention" model (t(98) = 8.84, p $<$ .001). However, assuming more focused attention improved fits significantly more for older than for younger adults (Fig. 2C). Additionally, pairwise comparisons between the inferred best fit parameters of the winning model revealed a significant difference in the θ parameter (t(48) = 2.70, p $<$ .01) (Fig. 2B).

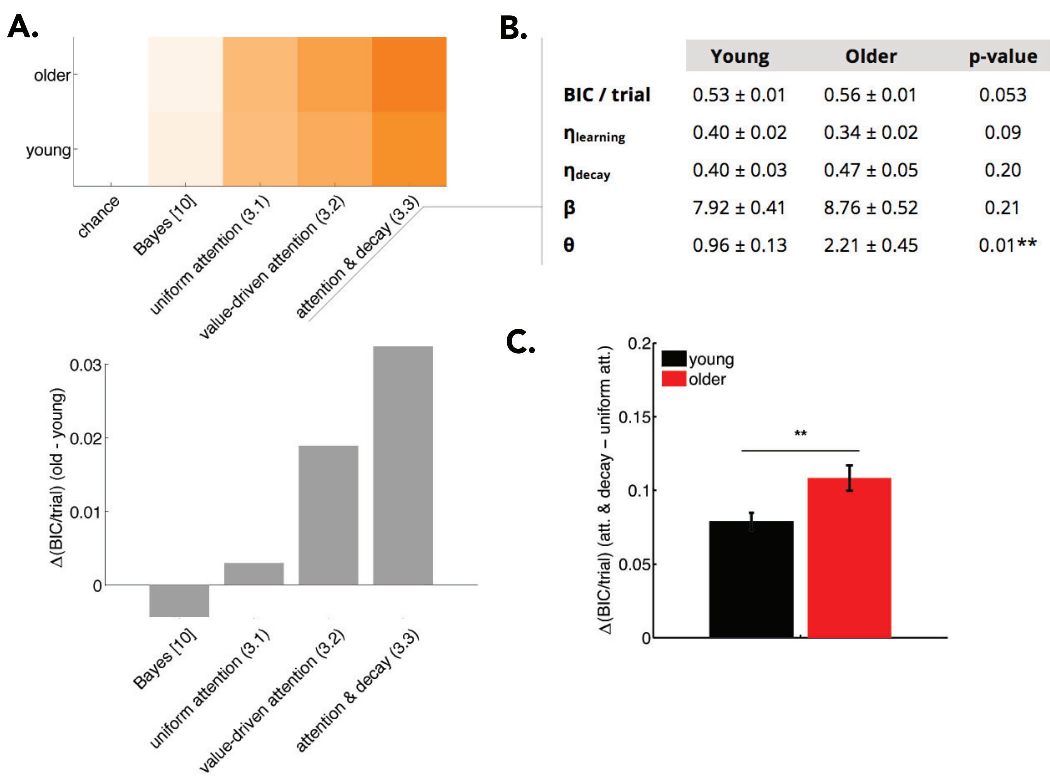

	Young	Older	p-value
BIC / trial	0.53 ± 0.01	0.56 ± 0.01	0.053
$\eta_{learning}$	0.40 ± 0.02	0.34 ± 0.02	0.09
η_{decay}	0.40 ± 0.03	0.47 ± 0.05	0.20
β	7.92 ± 0.41	8.76 ± 0.52	0.21
θ	0.96 ± 0.13	2.21 ± 0.45	0.01**

Figure 2: *Group comparisons of models assuming different attentional filters*. A. Top: mean group BIC-corrected likelihood per trial, models ordered from an "ideal observer" learner [10] to a progressively narrower focus of attention (darker shades correspond to higher likelihoods). Bottom: group difference in average BIC-corrected likelihood per trial. B. Average BIC-corrected likelihoods per trial (chance is 33%), and fit parameter estimates (mean±SEM across subjects) for the "attention & decay" model. Paired-sample t-tests revealed a significant difference in the θ parameter, suggesting the older adults' reliance on a learning strategy in which attention is more biased by high-valued features. C. The difference in average BIC-corrected likelihood between the "attention & decay" model and the "uniform attention" model was significantly higher in older adults (paired-sample t-test, t(48) = 2.84, p $<$ .01).

5 Summary and conclusions

We model choice data in a multidimensional decision making task using RL with function approximation and operationalize selective attention during learning as a mechanism which biases the learner towards sparser representations. Models that assume a narrower attentional filter are significantly better at predicting the trial-by-trial choices of older adults than those of young adults. This suggests that when faced with the problem of reducing computational costs, older adults are closer to a sub-optimal strategy akin to *serial hypothesis testing* that enables less incidental learning. Moreover, the significant difference in the θ parameter suggests a tendency for older adults to dynamically allocate attention to features with high learned values. While less efficient, this could reflect an adaptation to a "noisier" striatal DA learning system and a reliance on past representations of a stimulus when selecting between options. Interestingly, this finding seems to agree with a large body of behavioral research describing an age-related bias towards positive valence during information processing [15]. Perhaps rather than reflecting an asymmetry in learning from gains versus losses [16], this "positivity bias" could result from a tendency to maintain attentional focus on those representations that have previously predicted reward.

Building on these findings, we lay the groundwork for a number of interesting questions: does more focused attention in older adults also mean less frequent switching? And if so, is there an age difference in how much feedback is needed to redirect attention, as suggested by studies that report a tendency to perseverate in older adults [17]? While the RL models presented here do not explicitly track participants' focus of attention, we can overcome this limitation by comparing older and younger adults' choice patterns using the model developed by Wilson and Niv [11], which uses the reward history to determine when to switch away from the currently attended feature. Finally, our framework provides a promising avenue for precisely characterizing underlying neural changes in the functional interaction between the PFC and the DA system and testing whether the age-related shift in strategy evidenced here is indeed related to weaker interactions between striatal prediction error activity and attentional modulation in the PFC.

References

[1] S. Li, U. Lindenberger, B. Hommel, G. Aschersleben, W. Prinz, and P. B. Baltes, "Transformations in the couplings among intellectual abilities and constituent cognitive processes across the life span," *Psychological Science*, vol. 15, no. 3, pp. 155–163, 2004.

[2] T. S. Braver and D. M. Barch, "A theory of cognitive control, aging cognition, and neuromodulation," *Neuroscience and biobehavioral reviews*, vol. 26, no. 7, pp. 809–817, 2002.

[3] S. Li, U. Lindenberger, and L. Bäckman, "Dopaminergic modulation of cognition across the life span.," *Neuroscience and biobehavioral reviews*, vol. 34, no. 5, p. 625, 2010.

[4] R. Cabeza and N. A. Dennis, "Frontal lobes and aging: deterioration and compensation," in *Principles of frontal lobe function* (D. Stuss and R. T. Knight, eds.), Oxford University Press USA, 2013.

[5] A. Gazzaley, J. W. Cooney, J. Rissman, and M. D'Esposito, "Top-down suppression deficit underlies working memory impairment in normal aging," *Nature neuroscience*, vol. 8, no. 10, pp. 1298–1300, 2005.

[6] K. L. Campbell, C. L. Grady, C. Ng, and L. Hasher, "Age differences in the frontoparietal cognitive control network: Implications for distractibility," *Neuropsychologia*, 2012.

[7] Y. Niv, "Reinforcement learning in the brain," *Journal of Mathematical Psychology*, vol. 53, no. 3, pp. 139–154, 2009.

[8] R. Chowdhury, M. Guitart-Masip, C. Lambert, P. Dayan, Q. Huys, E. Düzel, and R. J. Dolan, "Dopamine restores reward prediction errors in old age," *Nature neuroscience*, 2013.

[9] M. R. Roesch, D. J. Calu, G. R. Esber, and G. Schoenbaum, "All that glitters: dissociating attention and outcome expectancy from prediction errors signals," *Journal of neurophysiology*, vol. 104, no. 2, pp. 587–595, 2010.

[10] S. J. Gershman, J. D. Cohen, and Y. Niv, "Learning to selectively attend," in *32nd Annual Conference of the Cognitive Science Society*, 2010.

[11] R. C. Wilson and Y. Niv, "Inferring relevance in a changing world," *Frontiers in human neuroscience*, vol. 5, 2011.

[12] R. Sutton and A. Barto, *Reinforcement learning: An introduction*, vol. 116. Cambridge University Press, 1998.

[13] N. D. Daw, "Trial-by-trial data analysis using computational models," *Decision making, affect, and learning: Attention and performance XXIII*, vol. 23, p. 1, 2011.

[14] G. Schwarz, "Estimating the dimension of a model," *The annals of statistics*, vol. 6, no. 2, pp. 461–464, 1978.

[15] M. Mather, L. L. Carstensen, *et al.*, "Aging and motivated cognition: The positivity effect in attention and memory," *Trends in cognitive sciences*, vol. 9, no. 10, pp. 496–502, 2005.

[16] M. J. Frank, L. Kong, *et al.*, "Learning to avoid in older age," *Psychology and aging*, vol. 23, no. 2, p. 392, 2008.

[17] M. G. Rhodes, "Age-related differences in performance on the wisconsin card sorting test: a meta-analytic review.," *Psychology and Aging*, vol. 19, no. 3, p. 482, 2004.