

Active information seeking using the Approximate Number System

Jinjing (Jenny) Wang (jinjing.jenny.wang@gmail.com)

Elizabeth Bonawitz (lbaraff@gmail.com)

Department of Psychology, Rutgers University, Newark, NJ 07102

Abstract

Human adults share the ability to approximate large quantities without counting with newborn infants and non-human species. This ability is supported by the Approximate Number System (ANS) - a primitive and domain-specific cognitive system that supports noisy numerical decisions. How does the ANS support active exploratory decisions? Using a numerical comparison task, we found that the amount of active information seeking does not simply increase as the decision becomes more difficult. Instead, there seems to be an inverted U-shaped relationship between trial difficulty and how much one chooses to seek information. Additionally, this effect is not modulated by participants' performance, suggesting that participants' exploratory decisions based on ANS representations are driven by the utility of information seeking actions.

Keywords: Information Seeking; Active Learning; Approximate Number System; Decision Making

Introduction

How the mind processes sensory data and interprets the physical world is the hallmark question in cognitive science. However, the data we receive from the physical world is not readily interpretable. Rather than passively absorb all the information that is available, humans and animals actively explore and selectively attend to aspects of the world (Gottlieb, Oudeyer, Lopes, & Baranes, 2013). This kind of active exploration and selective attention is essential to effective learning and proper cognitive functioning. What determines when we want to explore and to what we choose to attend?

It has been widely demonstrated that observers, humans and animals alike, are drawn to novel and surprising events, which is often explained by a motivation to decrease errors in prediction (Loewenstein, 1994; Schultz & Dickinson, 2000). According to Loewenstein, an observer's desire to learn about a specific topic is driven by a discrepancy between the observer's existing knowledge and what they would like to know. Consistent with this account, infants as young as 10 months old can form expectations about object behavior, and explore more when these expectations are violated (Stahl & Feigenson, 2015). Relatedly, school-age and preschool children prefer to play with toys whose functionality are ambiguous or unexpected (Bonawitz, van Schijndel, Friel, & Schulz, 2012; Schulz & Bonawitz, 2007). This kind of prediction errors cannot only be

mathematically defined, but has also been decoded from neuronal activities (Bromberg-Martin & Hikosaka, 2011).

In addition to novelty and surprise, humans and animals are also drawn to more complex stimuli or more difficult situations (Berlyne, 1966). For example, when confined in a minimally-stimulated space, adults prefer to produce light patterns that are the most diverse and unpredictable (Jones, Wilkinson, Braden, 1961). In another experiment, when probed about their curiosity about facts related to different animal species, adult participants were more curious about facts that they knew less about (Berlyne, 1954). These phenomena, that exploration is driven by novelty, surprise, and complexity, are consistent with the information processing account that defines information gain by uncertainty (Berlyne, 1960).

However, this tendency to be drawn to situations with maximum uncertainty (and to reduce it through learning actions) seems counterproductive in many cases. In particular, when the gap between one's current epistemic state and the information provided by the environment is too big, actions of learning and exploration can yield little benefit. For example, no matter how much effort a reader puts into staring at some foreign words without knowing the language or having access to a dictionary, the reader would still have no clue what the words mean.

Instead of linearly increasing exploratory actions as uncertainty increases, numerous studies have demonstrated a trade-off between the cost and benefit of information seeking actions (Coenen, Nelson, & Gureckis, 2018). When reading and rating contentful questions, such as "what instrument was invented to sound like human singing," adult participants' rated level of curiosity was the highest for questions that they had intermediate levels of confidence, and their level of curiosity was the lowest for questions in which they either had extremely low confidence or extremely high confidence (Kang et al., 2009). In a different exploratory situation, where each task option was initially hidden from participants, participants' exploratory decisions also followed a similar U-shaped pattern - they explored the most when the task was moderately difficulty, and explored less when the task was either too easy or too hard (Baranes, Oudeyer, & Gottlieb, 2014).

Consistent with these results, the field of developmental robotics suggests that exploration is based on dynamic changes in the rate of learning (Gottlieb, Oudeyer, Lopes, & Baranes, 2013; Oudeyer, Kaplan, & Hafner, 2007). Robots with this rate-based learning system can efficiently learn

skills in high dimensions without being distracted by activities that are either well learnt or unlearnable (Baranes & Oudeyer, 2013; Pape et al., 2012). Exploration increases as the rate of information increases. In cases when there is very low certainty (or high uncertainty), any particular action may produce new information, but if the problem space is complex enough, then additional information may not produce significant shifts in belief weights -- thus highly complex environments may not produce information that supports learning rates. Instead, learning rate may be highest in the Goldilock's spot (Kidd, Piantadosi, & Aslin, 2012), in which any particular action produces information to support a steeper learning rate. This predicts that, rather than a direct linear relationship, exploration should be lowest at both extremely low and extremely high levels of uncertainty, and exploration should be the highest at an intermediate levels, where information has the highest rate of return.

Results from infants' preference for object complexity are consistent with this account. Seven- and 8-month-old infants' probability of looking at an event was the lowest when looking at either highly predictable or highly surprising content (Kidd, Piantadosi, & Aslin, 2012; 2014; Piantadosi, Kidd, & Aslin, 2014; see also Pelz & Kidd, *in prep*). These results suggest that infants are able to direct their attention to maintain an intermediate rate of information absorption. It is possible that this kind of attentional mechanism is in place to prevent infants, who arguably have the most to learn and the least resources, from wasting cognitive resources on either overly predictable or overly unpredictable information.

One open question is whether adults reveal such trade-offs in active exploratory decision making situations. It is possible that this kind of balance between cognitive resource and exploration is unique to childhood. Another open question is whether such trade-offs are unique to novel learning environments or tasks that require higher-level conceptual reasoning, such as deciding what questions to ask or which route to take in a novel environment. When performing familiar activities using acquired skills, one may not need to adjust exploration based on uncertainty. Alternatively, the expected information gain from exploratory actions may explain information seeking behavior beyond these contexts.

To address these questions, the current study uses adults' exploratory decisions using a primitive cognitive system as a case study to test the relationship between problem difficulty and adults' exploratory decisions. Upon seeing 20 dots and 10 dots, without counting, we can immediately tell which array has more dots. This ability to automatically and effortlessly discriminate large numerosities is supported by the Approximate Number System (ANS; Dehaene, 1997), which produces noisy and ratio-dependent representations in

human adults (Halberda, Ly, Wilmer, Naiman, & Germine, 2012), newborn infants (Izard, Sann, Spelke, & Streri, 2009), as well as non-human species (Cantlon, Platt, & Brannon, 2009; Dehaene, Dehaene-Lambertz, & Cohen, 1998). With ANS representations, discriminating 20 dots from 10 is just as easy as discriminating 40 from 20 (a ratio of 2), but both are easier than discriminating 15 from 10 (a ratio of 1.5). The discriminability of numerosities is determined by the numerical ratio, instead of set size, non-numerical dimensions (such as size of individual dots). In other words, the Approximate Number System strictly obeys Weber's Law (Dehaene, 2003). This well-established law allows us clean control over the difficulty and uncertainty of the trials - the less discriminable the trials (the closer the ratio), the more uncertainty. Additionally, infants and adults are able to maintain multiple numerical representations at once (Feigenson, 2008; Zosh, Halberda, & Feigenson, 2011).

This intuitive and automatic cognitive system provides a case study for testing the scope of the expected information gain account - whether adults' decision making using the intuitive and automatic numerical representations also demonstrate a cost and benefit trade-off of information seeking actions. It has been recently suggested that adults and children are sensitive to their internal confidence in numerical decisions (Baer, Gill, & Odic, 2018; Halberda & Odic, 2015), and numerical precision can be influenced by the order of trial difficulty (Odic, Hock, & Halberda, 2014; Wang, Libertus, & Feigenson, 2018; Wang, Odic, Halberda, & Feigenson, 2016). It is possible that this sensitivity to internal confidence or uncertainty drives adults' exploratory decisions in a way that balances the cost and benefit of information seeking actions. On the other hand, neuroimaging studies revealed that the encoding of ANS signals are extremely rapid - as fast as 180ms in the bilateral occipital-parietal sites (Hyde & Spelke, 2009; Park, DeWind, Wordoff, & Brannon, 2015). It is possible that this automatic encoding of numerical information leaves little room for improvement from exploratory actions, and hence adults may show no cost-benefit tradeoff in their exploratory decisions in a numerical task.

To test this, we designed a nonverbal numerical comparison task with four alternative forced choices. This design ensures that the numerical representations can be maintained in adults' working memory (i.e., about four items; Epelboim & Suppes, 2001; Luck & Vogel, 1997). On the other hand, a four-alternative-forced-choice paradigm lowers the chance level to 25%, which increases the performance gap between random guessing and effortful performance by 25% compared to two-alternative-forced choice tasks (which has a 50% chance level), potentially providing more utility for information seeking actions.




Figure 1: Schematic of the experimental procedure.

To further reduce random guessing, we also offered participants a small reward bonus depending on their performance. Participants can choose to see any one of four large arrays of dots, each for only 200ms which is too brief a window to count the dots. The key difference from traditional numerical comparison tasks is that participants are given the option to re-explore each array as many times as they would like before deciding the largest array. Participants then decide when they are ready to choose the array with the largest numerosity. Numerical comparison tasks allow us to systematically quantify and vary uncertainty and the difficulty of the task by changing the ratio between the numerosities.

If exploration is driven by the utility of information seeking actions, we would expect to see an inverted U-shaped relation between trial difficulty and the amount of exploratory actions. Alternatively, if exploration is driven by performance or error rate, then we should expect participants to explore most in trials in which the difficulty of the trials is the highest. Such an account would reveal information seeking to be linearly related to trial difficulty. Finally, if adults are not sensitive to the uncertainty of the trials, then exploration should not vary with the complexity of the trials.

Method

Participants Forty-two adults were recruited online through Amazon Mechanical Turk.

Stimuli Stimuli consisted of series of arrays containing collections of blue, red, yellow, and cyan dots on a grey background. During all the trials, three of the four arrays always contained the same number of dots (in different layout and configuration), and the fourth array differed from the remaining three with variable ratios. Difficulty was manipulated by changing the ratio between the largest number and the remaining number. Ratios varied between 1 (i.e., all four arrays were the same; the correct answer was pre-determined and randomly generated) and 2 (i.e., the larger number was twofold the smaller number), with at least 6 trials of each ratio. There were a total of 128 trials a participant could possibly complete.

Procedure After reading the instructions, each participant received an untimed practice session with eight practice trials. Participants then completed a timed test session where they had five minutes to complete as many trials correctly as possible. Participants were compensated based on how many trials they answered correctly during the test trials. Each trial started with four empty boxes outlined with distinct colors and paired with a reminding message about which key to press to “flip” the box and reveal the dots.




Figure 2: Average accuracy as a function of ratio (larger/ smaller). Error bars represent confidence interval of the mean.

After each keypress, dots appear for 200ms in the chosen box. For example, pressing the “R” key showed blue dots in the blue box, and pressing “U” showed red dots in the red box (Figure 1). During the flipping phase, the participant could press the spacebar to indicate that they were ready to move onto the decision phase at any point. Once the participant had moved to the decision phase, they were prompted to press a key to indicate which box contained the most dots. Feedback was provided after each trial. Participants on average completed 37.67 test trials ($SD = 21.80$).

Results

We first examined participants’ accuracy in the decision phase. On average, participants performed correctly 62% of the time, well above chance (25%; binomial exact test $p < .001$).

We then averaged each participants’ performance for each ratio to analyze the effect of ratio on accuracy. If participants used the ANS to solve the task, their performance should show the ratio-dependent signature of

the ANS. Alternatively, it is possible that participants were able to count or maintain more precise representation of the numerical arrays after seeing them multiple times. As shown in Figure 2, participants’ accuracy increases significantly as the ratio becomes easier. A log-linear regression model predicting accuracy using ratio explains over 72% of the variance, $\beta = .86$, $t = 6.54$, $p < .001$, suggesting that participants primarily relied on ANS representations in the current numerical comparison task, even when they could receive additional information about the numerical stimuli. Consistent with previous research on adults’ ANS precision, participants’ accuracy on the task plateaued at about 1.5 ratio (Halberda & Feigenson, 2008).

The central question of the current study is how the difficulty of numerical decisions impacts people’s information seeking. To test this, we examined the relationship between ratio and the number of boxes participants flipped before the decision phase.




Figure 3: Average number of boxes flipped before decision as a function of ratio (larger/ smaller). Error bars represent confidence interval of the mean.




Figure 4: Average number of boxes flipped before decision as a function of accuracy. Error bars represent confidence interval of the mean.

As shown in Figure 3, overall, participants sought more additional information when the trials were more difficult. However, instead of a simple linear increase in number of flips as the ratio decreases, there is an inverted U-shaped relationship between ratio and flips when the ratio was between 1 and 1.25. Indeed, it is precisely in this range that participants steeply shift from near chance performance to near ceiling performance. This supports

the claim that exploration is driven by expected information gain.

To test for a quadratic trend of ratio on number of flips, or an inverted U-shaped relationship between exploration and numerical ratio, we ran a model using both linear and quadratic ratio terms. This revealed a significant effect for both ratio ($\beta = -1.08, t = -3.49, p < .001$) and ratio-squared ($\beta = .91, t = 2.91, p = .004$). However, we found no relationship between average accuracy and the number of boxes flipped before decision (Figure 4). This

suggests that, rather than perceived performance or general motivation, participants' expected information gain drives their information seeking.

Conclusions

The current study investigated the relationship between the difficulty of numerical decisions and exploratory decisions. We found that adults' active search for additional information was the highest for trials with intermediate difficulty, and the lowest when the trials were either too easy or impossibly hard. Moreover, exploration has no clear relationship with numerical performance. These results suggest that numerical difficulty drives adults' exploratory decisions, showing a trade-off between cost of exploratory action and expected benefit from exploration.

Previous research has shown that infants seem to prefer an intermediate flow of information when exploring the environment (Kidd et al., 2012), and adults in novel or complex exploratory tasks explore the most when the task is at intermediate level of uncertainty (Kang et al., 2009; Baranes et al., 2014). These results have been taken to suggest that in learning and exploration, the observer have a tendency to optimize the cost of action and the gained information (Coenen et al., 2018). The current results extends this literature by suggesting that adults remain motivated to show such trade-off in their exploratory decisions even when using the primitive Approximate Number representations that have been active since infancy.

These results are consistent with both the idea that adults can balance the cost and benefit of exploratory actions, and that the rate of information gain drives exploratory behavior adults' numerical decision making. One possibility is that adults were making immediate decisions about whether to explore more solely based on the difficulty of each trial. Alternatively, it is possible that adults were dynamically adapting their exploratory decisions based on observed performance change, or their observed rate of learning, from previous explorations. Future research exploring the benefit of the exploratory actions, such as performance change with and without exploration, will help clarify the mechanisms by which adults make their exploratory decisions.

Where does this ability to dynamically adapt exploration to our own uncertainty come from? The similar U-shaped pattern in infants' attention suggests that infants are able to respond to probabilistic uncertainty in the environment (e.g. Kidd et al., 2012). However, it is possible that the ability to monitor the uncertainty in one's cognitive representations, such as numerical precision, may require more advanced metacognitive skills. Alternatively, infants may already come equipped with implicit representations of their uncertainty in numerical decisions. Recent work suggests that infants as young as 6 months old perform differently in a numerical change detection task as as the order of trial difficulty

changes (Wang, Libertus, & Feigenson, 2018). It remains to be tested whether infants can adapt their exploratory behavior when using Approximate Number representations, and whether their exploration has the same kind of relationship with trial difficulty.

Another important question raised by the current study is whether active information seeking boosts numerical precision. In general, we found no relationship between overall accuracy and information seeking. It is possible that seeing the dot arrays more does not actually significantly impact people's accuracy at making numerical decisions. On the other hand, it remains possible that more complex interactions exist between information seeking and numerical precision. Future work examining the difference between people's numerical accuracy with and without information seeking will help test these possibilities.

References

- Baer, C., Gill, I. K., & Odic, D. (2018). A domain-general sense of confidence in children. *Open Mind*, 2(2), 86-96.
- Baranes, A. F., Oudeyer, P. Y., & Gottlieb, J. (2014). The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Frontiers in neuroscience*, 8, 317.
- Berlyne, D. E. (1954). A theory of human curiosity. *British Journal of Psychology. General Section*, 45(3), 180-191.
- Berlyne, D. E. (1960). Conflict, arousal, and curiosity.
- Berlyne, D. E. (1966). Curiosity and exploration. *Science*, 153(3731), 25-33.
- Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64(4), 215-234. <https://doi.org/10.1016/j.cogpsych.2011.12.002>
- Bromberg-Martin, E. S., & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nature Neuroscience*, 14(9), 1209-1216. <https://doi.org/10.1038/nn.2902>
- Cantlon, J. F., Platt, M. L., & Brannon, E. M. (2009). Beyond the number domain. *Trends in cognitive sciences*, 13(2), 83-91.
- Coenen, A., Nelson, J. D., and Gureckis, T. M. (2018). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic bulletin & review*, pages 1-41.
- Dehaene, S. (1997). *The number sense: How the mind creates mathematics*. Oxford University Press, USA.
- Dehaene, S. (2003). The neural basis of the Weber-Fechner law: a logarithmic mental number line. *Trends in cognitive sciences*, 7(4), 145-147.
- Epelboim, J., & Suppes, P. (2001). A model of eye movements and visual working memory during problem solving in geometry. *Vision Research*, 41(12), 1561-1574. [https://doi.org/10.1016/S0042-6989\(00\)00256-X](https://doi.org/10.1016/S0042-6989(00)00256-X)

- Feigenson, L. (2008). Parallel non-verbal enumeration is constrained by a set-based limit. *Cognition*, *107*(1), 1–18. <https://doi.org/10.1016/j.cognition.2007.07.006>
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, *17*(11), 585–593. <https://doi.org/10.1016/j.tics.2013.09.001>
- Halberda, J., Ly, R., Wilmer, J. B., Naiman, D. Q., & Germine, L. (2012). Number sense across the lifespan as revealed by a massive Internet-based sample. *Proceedings of the National Academy of Sciences*, *109*(28), 11116–11120.
- Hyde, D. C., & Spelke, E. S. (2009). All numbers are not equal: an electrophysiological investigation of small and large number representations. *Journal of cognitive neuroscience*, *21*(6), 1039–1053.
- Izard, V., Sann, C., Spelke, E. S., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences*, *106*(25), 10382–10385.
- Jones, A., Wilkinson, H. J., & Braden, I. (1961). Information deprivation as a motivational variable. *Journal of Experimental Psychology*, *62*(2), 126.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T. Y., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, *20*(8), 963–973.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks Effect: Human Infants Allocate Attention to Visual Sequences That Are Neither Too Simple Nor Too Complex. *PLOS ONE*, *7*(5), e36399. <https://doi.org/10.1371/journal.pone.0036399>
- Kidd, C., Piantadosi, S.T., & Aslin, R.N. (2014.) The Goldilocks effect in infant auditory cognition. *Child Development*, *85*(5):1795-804.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, *116*(1), 75–98. <https://doi.org/10.1037/0033-2909.116.1.75>
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. <https://doi.org/10.1038/36846>
- Odic, D., Hock, H., & Halberda, J. (2014). Hysteresis affects approximate number discrimination in young children. *Journal of Experimental Psychology: General*, *143*(1), 255.
- Oudeyer, P., Kaplan, F., & Hafner, V. V. (2007). Intrinsic Motivation Systems for Autonomous Mental Development. *IEEE Transactions on Evolutionary Computation*, *11*(2), 265–286. <https://doi.org/10.1109/TEVC.2006.890271>
- Park, J., DeWind, N. K., Woldorff, M. G., & Brannon, E. M. (2015). Rapid and direct encoding of numerosity in the visual stream. *Cerebral cortex*, *26*(2), 748–763.
- Pelz, M. & Kidd, C. (In prep.) The dynamics of attentional switching in a complex environment.
- Piantadosi, S.T., Kidd, C., & Aslin, R.N. (2014) Rich Analysis and Rational Models: Inferring individual behavior from infant looking data. *Developmental Science*, *17* (3): 321–337.
- Schultz, W., & Dickinson, A. (2000). Neuronal Coding of Prediction Errors. *Annual Review of Neuroscience*, *23*(1), 473–500. <https://doi.org/10.1146/annurev.neuro.23.1.473>
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045–1050. <https://doi.org/10.1037/0012-1649.43.4.1045>
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants’ learning and exploration. *Science*, *348*(6230), 91–94. <https://doi.org/10.1126/science.aaa3799>
- Wang, J., Libertus, M. E., & Feigenson, L. (2018). Hysteresis-induced changes in preverbal infants’ approximate number precision. *Cognitive Development*, *47*, 107–116. <https://doi.org/10.1016/j.cogdev.2018.05.002>
- Wang, J. J., Odic, D., Halberda, J., & Feigenson, L. (2016). Changing the precision of preschoolers’ approximate number system representations changes their symbolic math performance. *Journal of Experimental Child Psychology*, *147*, 82–99.
- Zosh, J. M., Halberda, J., & Feigenson, L. (2011). Memory for multiple visual ensembles in infancy. *Journal of Experimental Psychology-General*, *140*(2), 141.