Attention shifts to more complex structures with experience

Tess Allegra Forest¹, Noam Siegelman^{2*}, and Amy S. Finn^{1*}

1. Department of Psychology, University of Toronto
100 St. George Street, 4th Floor
Sidney Smith Hall
Toronto, ON M5S 3G3 Canada
416-978-5201

2. Haskins Laboratories
300 George Street, Suite #900
New Haven, CT, 06511, USA
203-865-6163

Corresponding Author:

Tess Allegra Forest
100 St. George Street. 4th floor
Sidney Smith Hall
Toronto, ON M5S 3G3
Email: tess.forest@mail.utoronto.ca

Introduction and Discussion: 1992 words

Author contributions statement: All authors conceptualized the research idea. TAF created experimental stimuli and set-up, and completed data collection. TAF and NS completed statistical analyses. TAF drafted manuscript, and all authors worked together to edit manuscript. Asterisk (*) denotes NS and ASF contributed equally to the work.

Abstract

Our environments are saturated with learnable information. What determines which of this information gets prioritized for limited attentional resources? While previous studies suggest that learners prefer medium complexity information, here we argue that what counts as medium should change as someone learns an input's structure. Specifically, we examined the hypothesis that attention is directed towards more complicated structures as learners gain more experience with the environment. Participants watched four simultaneous streams of information which varied in complexity. Reaction times to intermittent search trials (Ex. 1, N=75) and eye-tracking (Ex. 2, N=45) indexed where participants attended over the experiment. Using a subject- and trial-specific measure of complexity, we demonstrated that participants attended to increasingly complex streams over time. Individual differences in structure learning also modulated attention allocation, with better learners attending to complex structures from earlier in learning, suggesting the ability to prioritize different information over time gates structure learning success.

Keywords: Attention allocation; statistical learning; entropy; individual differences.

Our environments are saturated with information which varies in complexity. Some of this information is likely familiar, and some novel. Some information is easily learnable, while other information is more opaque or even absent of structure altogether. But what determines which information is prioritized for processing with our limited attentional resources, in what order, and when? Although infants (Kidd et al., 2012) and children (Cubit et al., 2021) attend to information that is neither overly simplistic nor overly complex, past studies have not fully considered the dynamic nature of the environment and the learning process. That is, not only does our environment change over time, but we also gradually learn its structure, suggesting that what is prioritized for learning might also shift over time. One intriguing possibility is that a learner could tackle the less complex aspects of their environment first, later progressing to more complicated structures. Here, we ask whether learners flexibly shift which structures they attend to based on their experience in a complex environment with multiple sources of structured information.

While we do not know how different sources of information are prioritized for limited attentional resources over time, there is a special relationship between attention and learning, such that we attend to learnable sources of information. For example, visual attention is biased towards stimuli that occur in a more frequently experienced color over those that occur in a random color (Alamia & Zénon, 2016), and attention is preferentially directed towards structured (and therefore learnable) aspects of an environment over random ones (Zhao et al., 2013). When presented with four streams of images simultaneously, one of which was structured and three of which were random, participants responded faster to visual search trials when the target was embedded in the structured stream as compared to the random streams, indicating that attention was directed to structure (Zhao et al., 2013). This relationship between attention and learning is

bidirectional: in visual search tasks, learners shift their attention towards locations that have a high probability of containing a target based on previously-experienced structure (Li & Theeuwes, 2020; Wang et al., 2019), and in reinforcement learning paradigms learners shift their attention based on the task dimension that is currently most rewarding (Leong et al., 2017). The spatial scale of the regularities someone experiences also determines how they attend to visual scenes—participants exposed to visual displays containing local shape structures made judgements about individual shapes more quickly than participants exposed to visual displays containing global structure (Zhao & Luo, 2017). Thus, not only do regularities in our environment capture attention, but learning these regularities shapes future attention allocation towards those same structures.

That said, it is currently unknown how this interplay between attention and learning unfolds over time when multiple structured sources of information are present. One exciting possibility is that learners shift their attention to different sources of information based on *how much* structure there is to learn from a given source. One standard way to measure how much learnable structure is present is by calculating the entropy (or the degree of uncertainty; less deterministic patterns have a higher entropy) in a source of information (Shannon, 1948). Since the 1950s, there has been great interest in understanding how human processing is impacted by an input's entropy (Miller, 1956; Shannon, 1951). While explaining human behaviour as a function of input complexity has had mixed success (Luce, 2003), it is now clear that humans are sensitive to the entropy of their environments (Hasson, 2017; Itti & Baldi, 2009; Nastase et al., 2014; Strange et al., 2005). These diverse works suggest learners are equipped to monitor how much structure there is to learn from a particular source—potentially a crucial step in attending to the most learnable information first.

Still, it is not entirely clear how sensitivity to structure, and greater attention to structured over random information, translates to navigating simultaneous sources of structured information—both in regard to what is attended first, and how attention shifts with experience. For example, which structures are prioritized for initial learning? One possibility is that learners prefer information sources which provide slightly more information than they have already learned. This idea has been applied to understanding how infants allocate their attention: infants attend longer to visual or auditory displays which are moderately different from what they habituated to (Kinney & Kagan, 1976), and both infants (Kidd et al., 2012) and children (Cubit et al., 2021) are least likely to disengage from medium complexity information. This privileged status of medium complexity information is observed across domains: adults' aesthetic preferences favor medium complexity images to simple or complex ones (Berlyne et al., 1971), and curiosity research has shown people prefer information that is useful for their learning (Gottlieb et al., 2013; Loewenstein, 1994; Wade & Kidd, 2019). One theoretical explanation for this phenomenon is that people attend to medium complexity sources because they are neither too simple nor too complex, and therefore most learnable or useful for that person.

Clearly, learners prefer structured information, and are sensitive to entropy variations to attend to mid-level complexity. Still, the *information that is mediumly complex depends on a learner's own experience*, and should therefore shift as learners gain more experience in their environment. This is because humans rapidly learn the statistical structure of their environments (Armstrong et al., 2017; Batterink, 2017; Forest et al., 2021; Saffran et al., 1996). Regardless of the structure prioritized initially, continued experience with this structure renders it redundant as learning occurs.

We therefore asked how learners' attention shifts to different levels of complexity, and, importantly, how this changes as a function of experience. We used a modified version of a statistical learning paradigm introduced by Zhao et. al. (2013). Our task presented learners with multiple sources of structured information simultaneously, which varied in complexity. Then, we measured what learners attended to over the experiment using reaction time (Ex. 1 and 2) and eye-tracking (Ex. 2). In order to understand how a learner's unique experience modulated their attention to different levels of complexity over time, we calculated a trial- and subject-specific measure of entropy based on the images each subject had seen thus far ("real-time entropy"). We then used real-time entropy to model how learners shift their attention as a function of their experience. Finally, because individuals vary in their structure learning ability (Misyak & Christiansen, 2012; Siegelman et al., 2017), we explored the potential consequences of this attention allocation for successful learning in complex environments. We predicted that (1) learners would gradually shift their attention towards more complex aspects of their environment with experience, and (2) preferentially attending to different complexity sources would relate to individual differences in statistical learning.

Experiment 1

Methods

Participants

75 young adults (mean age = 19.71 years, SD = 2.98 years) participated in exchange for credit in an introductory psychology class at the University of Toronto. All participants had normal or corrected to normal vision and reported no history of neurological impairments. Because no prior research has directly addressed the question at hand, we did not have an a priori estimate of the

effect size of the crucial time by complexity interaction for a formal power analysis. Still, we decided upon a sample size of 75 prior to data collection as this is enough to detect moderate effect sizes of Cohen's $d = 0.4$ or more in 80% power (Batterink et al., 2015; Forest et al., 2021), which is conservatively smaller than typical effects observed in various statistical learning paradigms, and in line with past research including studies designed to answer questions about individual differences (e.g., Siegelman et al., 2017). Sample size and analyses were preregistered on the Open Science Framework (OSF) website (<https://osf.io/nwhk4>). All data collection was approved by the ethics board at the University of Toronto.

Procedure

Exposure & Visual Search. Participants first completed an exposure and visual search phase, in which they watched a series of images on two types of trials—learning trials and intermittent visual search trials (described below). All phases of the experiment were completed on Apple desktop computer using PsychoPy (Peirce, 2008). We adapted a previous study in which learners watched one structured and three random sources of information (Zhao et al., 2013), to instead include multiple sources of information with varying levels of complexity (described below). Other changes were made so that the task will work with children in future studies (using a cover story that includes looking for ‘lucky’ four-leaf clovers).

Learning Trials. The most frequent (480/576) trials were ‘learning trials’. These stimuli consisted of four streams, each comprised of eight unique black and white shapes (Fiser & Aslin, 2002; Figure 1). In each stream, the 8 shapes appeared in a sequence which varied in its complexity (described below; *Stream structure*). Each stream appeared in its own, unique

location (top, bottom, left, or right, randomized by participant with an equal likelihood of any stream appearing in any location for each person) throughout the study. On each trial, participants saw one shape from each stream. Thus, on learning trials, participants saw one black and white shape in each of four fixed locations on the screen (top, bottom, right, and left, *Figure 1a*). On each learning trial, the shapes appeared simultaneously and remained on screen for 800 ms (200 ms ISI). There were 480 learning trials over the course of the entire study. Participants were instructed simply to watch the shapes on these trials.

Stream structure. The shapes used on learning trials consisted of 32 shapes that were divided into four streams of eight shapes each. Within each stream, the eight shapes were grouped into four pairs of two shapes each (AB, CD, EF, GH). In each pair, one shape was assigned to be the first item in the pair (Shape A), and the other to be the second (Shape B). During exposure, shapes were presented one at a time. The groups of eight shapes and the pairs of shapes were consistent for all participants.

The four streams varied in how reliably the two items in each pair appeared together, or more formally, in the transitional probabilities (TP) between shapes within pairs. In one stream, the two shapes in each pair always followed one another, such that Shape A preceded Shape B 100% of the time (just as C preceded D, E-F, and G-H 100% of the time; TP between shapes within a pair = 1.0, *Figure 1b*). In another stream, paired shapes occurred in direct succession 80% of the time (e.g., Shape A preceded Shape B 80% of the time, likewise for all other pairs; TP within all pairs = 0.8). In a third stream, paired shapes occurred in direct succession 40% of the time (e.g., Shape A preceded Shape B 40% of the time, likewise for all other pairs; TP within pair = 0.4). The final stream was random, such that each shape could follow any other (with a

sole constraint on repetitions of the same shape), and thus the transitional probability between any two shapes was 0.14, or 1/7 (because there were eight shapes total). In the two structured streams with a $TP < 1$, when Shape A did not precede its paired Shape B (20% or 60% of the time for streams with $TP=0.8$ and $TP=0.4$, respectively), Shape B was replaced by the second item from another pair (e.g. Shape D, likewise for all other pairs; see also Bogaerts, Siegelman, & Frost, 2016; Siegelman, Bogaerts, & Frost, 2019). This shape was from each of the other three pairs an even number of times. In the random stream the same shape could not repeat twice in a row. Overall, the four streams varied in their degree of complexity or unpredictability: From the simplest (i.e., most predictable, $TP = 1$), to simple ($TP = 0.8$), to complex ($TP = 0.4$) to the most complex (random) stream (see *Figure 1a*).

Visual Search Trials. Intermittently, participants were presented with a visual search trial instead of a learning trial. On these trials, participants saw a colored clover in each of the same four locations in which shapes appeared on learning trials (*Figure 1a*). Each color (pink, blue, or yellow) was used an equal number of times throughout the study, and on any given search trial all flowers were the same color. These clovers had either a circular or a flat stem base. On each search trial, three of the four clovers had three leaves and one had four leaves. Participants' job was to locate the four-leaf clover (target) from amongst the three-leaf clovers (distractors) and indicate whether the stem of that clover had a circular or a flat base, as quickly and accurately as possible. Clovers remained on screen until participants responded. The reaction time data collected on these search trials allowed us to have an index of where participants were attending over time: if participants are attending to a particular location, they should be faster to respond

on search trials which appear in that spot (see also Zhao et al., 2013). After participants responded, the experiment progressed to the next learning trial.

Visual search trials were interleaved directly between learning trials. The frequency of visual search trials was distributed such that participants experienced a visual search trial on average every five learning trials (range between three and eight learning trials), to encourage participants to engage with the study for the entire course of learning without interrupting learning trials too frequently. There were 96 visual search trials across the entire course of study. The order of visual search trials relative to learning trials was unique for each participant.

The target of the visual search trial appeared equally frequently in each location on screen. We also ensured an identical number of search trials of each color, and in each location in each half of the learning phase (i.e. in learning trials 1-240 and 241-480). Moreover, the stem base of the target was flat and round an equal number of times. The distractors on each visual search trial consisted of flat and round stem bases, such that on any given visual search trial there were two flat and two round stem-based clovers. Visual search trials appeared equally frequently after A and B shapes, so they could not be used to segment the structured streams into pairs.

Testing knowledge of each stream offline. After the exposure phase, participants were tested on their knowledge of the patterns in each of the three structured streams using a two-alternative forced choice recognition test. Test questions asked participants to indicate which of two pairs of shapes seemed more like what they would have seen during the exposure phase. They then chose between two items that had formed a pair during exposure, or a first item shape followed by a second item shape from a different pair (in the same stream). Each pair of shapes from exposure was tested against four incorrect pairings of shapes, for a total of 16 test items per stream and 48

test items total. The 48 test items were presented in a random order that was uniquely determined for each participant.

Figure 1. a) Stimulus presentation during the learning phase. Either four black and white shapes (learning trials) or four colored clovers (search trials) appeared on each trial. In each location, black and white shapes appeared in a patterned order, ranging from the simplest pattern (darkest border; borders were not shown to participants and are used for visualisation purposes only) to the most complex (lightest border). b) Top panel visualises the transitional probabilities between shapes with a pair in each stream (related to complexity). Lower panel graph demonstrates how the real-time entropy value in each location changes over time, and uniquely for each subject. Two representative subjects' real-time entropy values are highlighted in pink and the rest are colored to match the border hues in the other components of the figure.

Analysis Approach

We were primarily interested in understanding how attention to different levels of predictability shifted as participants gained experience. Thus, we modeled how attention (as indexed by faster reaction time on search trials), shifts as a function of time and stream complexity. To do this, we ran two mixed effects models (detailed below) which predicted changes in reaction time as a function of time and complexity. Mixed-effect models included the maximal random effect structure that converged, with gradual removal of random effects when necessary based on those that were associated with the least variance (Barr et al., 2013). Models were fitted using the *lme4* package in R (Bates et al., 2015), and significance of fixed-effects was estimated using *lmerTest* (Kuznetsova et al., 2017).

Prior to running either model, we filtered reaction time data to not include incorrect trials (4.9% of trials), as well as trials with RTs shorter than 300 milliseconds, longer than 5 seconds, or more than ± 2 SDs away from a participant's mean (4.4% of correct trials).

Static Complexity. The first model that we ran included fixed effects for stream complexity (an effect-coded variable with four levels as visualized in Figure 1), trial number (centered and scaled), and their interaction. We also included a fixed effects term for target location to control for potential biases in RTs to targets in some locations; as expected, participants were significantly faster to respond to targets which appeared in the top location, see below. This pre-registered model also included random by-participant intercepts. Because some of our control variables contain multiple levels, we will report model comparison results (using the *anova()* function in R) which reflect the overall effect of the fixed effect across these levels. We also report the Beta values and standard error terms from the raw model outputs as a measure of

effect size (note that because 'stream' was effect-coded, beta values reflect the estimated difference between each level and the grand mean).

Real-time Entropy. While the model above consider differences in the overall complexity between the four streams, we were also interested in modelling how participants' own experience modulated what they were attending to at a particular point in time. For any given participant, complexity varies as a function of time. In fact, complexity is higher earlier in exposure, when participants have limited experience with each of the streams and there is still much noise in the information. Conversely, complexity is lower later in exposure, once participants have experienced enough repetitions of the pairs of a particular stream for the structure to have stabilized.

Because of this, we also calculated a measure of complexity that reflects what participants have experienced so far at any given moment, which we refer to as *real-time entropy*. Real-time entropy was calculated for each person, for each location, on each trial, by calculating the entropy in that location, given the information that had been presented up until that point. Operationally, this was done by creating a matrix of all possible transitions between shapes in each stream (size: 8x8). Then, for each subject, this matrix was updated to reflect the number of times a transition between any two shapes had occurred. This count matrix was then transformed into a TP matrix ($p(i|j)$ for each possible transition) and fed into a standard formula for entropy calculation (see below). At trial 1, all elements in the matrix were pre-defined as 1, reflecting equi-probable prior probabilities of observing all transitions. Then, at trial 2, the transition between the first and second shape was counted and entropy was re-calculated. To compute entropy from TP matrices, we used the Markov Entropy formula, which estimates the

degree of uncertainty (entropy) in a stream given an item's probability and between-item TPs, and is defined mathematically as: $-\sum_{i=1}^n p(i) \sum_{j=1}^n p(j|i) * \log p(j|i)$ (Nastase et al., 2014; Siegelman et al., 2019)). To reiterate, this calculation was repeated after each trial across the learning phase (ignoring search trials) for each location and for each subject—thus, because each of our participants saw each stream of shapes in unique order, this measure was subject, trial, and stream specific.

Using real-time entropy, we then ran the same mixed effects model described above (*Static Complexity*) using real-time entropy (centered and scaled) instead of stream as our independent measure of interest. All results will be reported in the same way as described for the *Static Complexity* model. Full model specifications are available on OSF (<https://osf.io/2h6m5/>).

Measuring learning offline. Because we were interested in understanding how attention to each stream might relate to eventual knowledge of the patterns in each stream, we also compared performance on the alternative forced choice test to chance (50%) using one sample t-tests, prior to examining how individual attention allocation to a particular stream related to later alternative forced choice performance. However, group level performance was not above chance on any of the streams and thus this second analysis was not carried out.

Results

Attention shifts to more complex locations over time. Our central question was whether attention would shift towards more complex structures as participants gained experience in a complex structured environment. Our *Static Complexity* model indicated that both stream and trial number independently predicted reaction time on visual search trials (stream: $F_{(3, 6434)}=3.32$,

$p = 0.02$, coefficients (relative to the simplest (TP = 1) stream), $\beta_{TP=0.8} = 0.03$, $SE = 0.01$, $\beta_{TP=0.4} = 0.01$, $SE = 0.01$, $\beta_{rand} = 0.02$, $SE = 0.01$; trial number: $F_{(1, 6435.3)} = 229.26$, $p = 0.02$, $\beta = -0.05$, $SE = 0.01$), such that participants were faster towards the end of the study. As expected (Previc, 1990), target location also independently predicted reaction times ($F_{(3, 6434)} = 70.98$, $p < .001$, coefficients (relative to left side location: $\beta_{top} = -0.03$, $SE = 0.01$, $\beta_{right} = 0.04$, $SE = 0.01$, $\beta_{bottom} = 0.09$, $SE = 0.01$), such that participants were fastest to respond on visual search trials where the target appeared on the top of the screen. Importantly, even when controlling for target location, and in line with our prediction that participants should attend to increasingly complex locations over time, there was also a significant interaction between stream complexity and trial number ($F_{(3, 6463.9)} = 3.37$, $p = 0.02$, coefficients (relative to simplest (TP = 1) stream), $\beta_{TP=0.8 \times Trial} = 0.01$, $SE = 0.01$, $\beta_{TP=0.4 \times Trial} = -0.01$, $SE = 0.01$, $\beta_{Rand \times Trial} = -0.02$, $SE = 0.01$) such that reaction times to more complex streams reduced more than for less complex streams as the experiment went on (as trial number increased). This suggests participants shifted their attention towards more complex locations as they gained experience.

Real-time entropy also predicts attention. Second, we were interested in understanding whether participants' individual experience shaped what they attended to. As with static complexity, real-time entropy predicted faster response times and trial number independently predicted slower reaction time on visual search trials (real-time entropy: $F_{(1, 6438.5)} = 19.23$, $p < .001$, $\beta = 0.02$, $SE = .004$; trial number: $F_{(1, 6438.7)} = 106.92$, $p < .001$, $\beta = -0.04$, $SE = .004$). And, as expected, target location again also predicted reaction times in this model ($F_{(3, 6438)} = 702.39$, $p < .001$, coefficients (relative to left side location: $\beta_{top} = -0.04$, $SE = 0.01$, $\beta_{right} = 0.04$, $SE = 0.01$, $\beta_{bottom} = 0.09$, $SE = 0.01$). Most central to our hypothesis, and while controlling for target

location, there was a significant interaction between real-time entropy and trial number ($F_{(1, 6439.1)}=31.90, p < .001, \beta = -0.02, SE = .004$) such that complex locations experienced a bigger drop in reaction times over the course of the study than simple locations did (Figure 2). This suggests that when taking into account the actual stimuli every participant was exposed to up until each point in time, attention shifted towards more complex aspects of the environment as they gained experience.

Figure 2. Model predictions for the interaction between trial number (X-axis) and real-time entropy (line color and dash thickness) on logged reaction time values (Y-axis) over the course of Experiment 1 (left graph), and Experiment 2 (right graph). Visualization depicts model predictions for real-time entropy values 1 standard deviation (SD) below the mean (darkest, thinnest dashed line), at the mean (middle dashed line), and 1(SD) above the mean (lightest, solid line) of real-time entropy. The interaction between trial number and real-time entropy was significant at $p < 0.001$ in both experiments.

Offline learning not above chance. We were interested in understanding how attention allocation related to eventual knowledge of each structure, so we compared participants' performance to chance (50%) at the group level, prior to examining how individual attention

allocation to a particular stream related to later alternative forced choice performance. However, group level performance was not above chance on any of the streams (simplest (TP=1) stream: 46.5% correct, $t(74) = -2.19$, $p = 0.03$; simple (TP=0.8) stream: 47.5% correct, $t(74) = -1.82$, $p = 0.07$; complex (TP = 0.4) stream: 50.6% correct, $t(74) = 0.41$, $p = 0.69$; note that the one significant test here reflects *below* chance performance levels, which are most likely spurious), and thus analyses relating learning to shifts in attention were not carried out. However, we were still interested in understanding how statistical learning ability might relate to attention in complex environments, and thus we ran a second experiment in which we gathered a second, separate measure of statistical learning ability for each participant.

Experiment Two

Methods

Experiment 1 suggested that learners shift their attention to more complex aspects of their environment with experience. However, the measure of attention used in Experiment 1 was collected on intermittent visual search trials rather than continuously, and was not linked to statistical learning ability as measured by an offline test. We therefore ran an additional experiment including eye tracking to replicate our findings using a more direct measure of attention. Moreover, we included an additional statistical learning task to better understand individual differences in the relationship between attention allocation and learning ability.

Participants

45 young adults (mean age = 19.08 years, SD = 0.84 years) participated in exchange for credit in an introductory psychology class at the University of Toronto. All participants had normal or

corrected to normal vision and reported no history of neurological impairments. Sample size was determined prior to data collection based on results of a power analysis using the data from Ex. 1, and preregistered on OSF (<https://osf.io/anxvd>). To calculate the sample size needed here, we ran a sensitivity analysis of the interaction between entropy and trial from Ex. 1 in the real-time entropy model across a range of sample sizes ($N=10-75$, randomly sampled subset of Ex. 1 participants, a total of 1,000 bootstrapped iterations at each N). This provided us with an expected power value at each sample size from 10-75, which allowed us to determine that after ~35 participants, we continuously achieved 90% power to detect the estimated interaction (real-time entropy * trial) from Ex. 1. Our design for Ex. 2 involved creating 24 counter-balanced conditions for the on-screen location of each stream (very simple, simple, complex, very complex), and thus, we preregistered a sample size of 48 (two subjects in each counter-balanced condition). However, we were forced to halt data collection due to COVID-19 restrictions after data had been collected from 45 subjects. Because our original power analysis suggested that even with 45 subjects we were well powered to detect our effect of interest (namely the real-time entropy * trial interaction), we chose to analyse the data at that point rather than collect the last three subjects' data.

Procedure

All materials and procedures were identical to those described for Ex. 1 except that eye tracking data was also collected while participants viewed the stimuli and completed visual search trials. The experiment was therefore presented using the Experiment Builder software (S.R. Research, 2011), while participants' eye-movements were recorded using an EyeLink 1000+ eye tracker. Following the offline measure of learning described in Ex. 1, participants

additionally completed a separate measure of statistical learning ability which was completed without an eye tracker using the same setup described for Experiment 1.

Statistical learning ability. Because we were interested in understanding how attention allocation relates to statistical learning ability, participants additionally completed a second, simpler, visual statistical learning task after completing the exposure and test phases described in Experiment 1 (albeit with eye-tracking, here after “main task”), with a short break in between. This secondary task was also presented on an Apple desktop computer using PsychoPy (Peirce, 2008). Aside from the use of different stimuli, this task was based on a published study (Siegelman et al., 2017) which was designed specifically for the measurement of individual differences. Thus, it includes an exposure phase including 8 triplets (from two levels of complexity, $TP=1$ and $TP=0.33$) followed by a test phase with 42 questions that vary in their task demands (pattern completion vs. recognition), number of response options (2-, 3- and 4-alternative-forced-choice), and the complexity of both the correct responses ($TP=0.33$ vs. $TP=1$) and the foils (ranging from $TP=0$ to $TP=0.5$). These experimental parameters result in a measure that is more reliable compared to other tasks in the field (Siegelman et al., 2017). Importantly, the stimuli used for this task were novel, colorful objects totally distinct from those used in the first part of the study. These data were analysed to provide a separate metric of individual statistical learning ability (i.e., number of correct responses from the 42 test questions) to include as a covariate in our models.

Analysis Approach

Our analysis approach largely mirrored Ex 1. As real-time entropy was a better predictor of attention allocation over time in Ex. 1 (and we felt it was a better construct of complexity in the first place), our preregistered analysis plan (<https://osf.io/anxvd>) and sample size calculations for Ex. 2 were based on that metric of complexity. First, to replicate Ex 1., we ran the same Real-time entropy model described above (*Ex. 1, Real-time Entropy*) using the reaction time data from search trials in Ex. 2. As in Ex. 1, we filtered reaction time data to not include incorrect trials (0.96% of trials), as well as trials with RTs shorter than 300 milliseconds, longer than 5 seconds, or more than ± 2 SDs away from a participant's mean (4.2% of correct trials).

Next, we ran a second model which used the trial-wise fixation time for each location (i.e. sum of fixation durations to each location on each trial) as a trial-by-trial metric of attention to each location. We included trial-wise fixation time as the dependent variable (instead of reaction time, henceforth “Fixation Time”) in the model to test whether attention varied over the course of the experiment in relation to complexity. For this model, fixations shorter than 80ms, or longer than 1 second were deemed implausible (i.e. likely the result of a tracker error because trial length was 1 second) and were therefore removed from the eye-movement analyses (3.7% of data points). We also used our separate measure of statistical learning ability to examine whether general attention allocation is related to performance on a separate statistical learning task. This statistical learning ability score was included in this fixation time model, as a continuous predictor (scaled and centered). The two- and three-way interactions between statistical learning ability, real-time entropy, and trial number were also included.

Results

As in Ex. 1, the offline alternative forced choice performance from the main task showed no evidence of learning at the group level (Simplest (TP=1) stream: 51.8% correct, $t(44) = 0.91$, $p = 0.37$; simple (TP=0.8) stream: 51.5% correct, $t(44) = 0.77$, $p = 0.44$; complex (TP=0.4) stream: 49.9% correct, $t(44) = -0.08$, $p = 0.94$). Because this was predicted based off of Ex. 1, and we included a separate statistical learning task to account for no group level learning of the original structures, it was not examined further.

Like in Experiment 1, Real-time entropy predicts attention.

Replicating the real-time entropy effects observed in Ex. 1, we found that real-time entropy, trial number, and target location all independently predict reaction times (Real-time entropy model; real-time entropy: $F_{(1, 56.9)} = 12.37$, $p = .001$, $\beta = 0.03$, $SE = 0.008$; trial number: $F_{(1, 3868.1)} = 117.28$, $p < .001$, $\beta = -0.05$, $SE = .005$; target location: $F_{(3, 2400.2)} = 33.69$, $p < .001$, coefficient (relative to left side location): $\beta_{\text{top}} = -0.03$, $SE = 0.01$, $\beta_{\text{right}} = 0.08$, $SE = 0.01$, $\beta_{\text{bottom}} = 0.003$, $SE = 0.01$), such that participants were faster to respond to higher complexity locations, later trials, and the top location on the screen. Importantly, there was also an interaction between trial number and real-time entropy ($F_{(1, 3880.5)} = 34.31$, $p < .001$, $\beta = -0.03$, $SE = .004$, *Figure 2*) mirroring that of Ex. 1, such that higher entropy locations were increasingly attended to a greater degree than lower entropy locations later on in the study.

Trial-wise fixation time also shifts with experience

After having replicated the main finding from Ex. 1, we next asked whether an additional index of attention—trial-wise fixation time—would also be shaped by experience in a complex

environment. The results of this fixation time model indicated that trial number and location on screen both independently predicted trial-wise fixation time (trial number: $F_{(1, 22, 911.1)}=62.40, p < .001, \beta = 0.04, SE = 0.005$; location: $F_{(3, 22, 913.6)}=2215.40, p < .001$, coefficients (relative to left side location): $\beta_{\text{top}} = 0.62, SE = 0.01, \beta_{\text{right}} = -0.03, SE = 0.01, \beta_{\text{bottom}} = 0.06, SE = 0.01$), such that participants fixated for longer towards the end of study and looked most at the top location.

Unlike the analyses using reaction time on search trials as the dependent variable, there was no main effect of real-time entropy on fixation time ($F_{(1, 22, 897.2)}=1.64, p = 0.20, \beta = -0.006, SE = 0.005$). Crucially however, there was a significant interaction between real-time entropy and trial number ($F_{(1, 22, 891.8)}=11.50, p = .001, \beta = 0.01, SE = .004$), conceptually replicating Ex. 1 (see *Figure 3a*): the more complex a location was, the bigger the shift in attention to that location over the course of the study, indicating that participants attended to more complex parts of their environment as they gained experience in that environment.

Figure 3. Model predictions for the interaction between trial number (X-axis) and real-time entropy (line color and dash thickness) on logged trial-wise fixation time values (Y-axis) over the course of the experiment. Visualization depicts model predictions for real-time entropy values 1 standard deviation (SD) below the mean (darkest, thinnest dashed line), at the mean (middle dashed line), and 1(SD) above the mean (lightest, solid line) of real-time entropy. The interaction between trial number and real-time entropy was significant at $p < 0.001$. b) Model predictions for the three-way interaction between trial number (X-axis), real-time entropy (line color and dash thickness), and statistical learning ability on logged trial-wise fixation time values (Y-axis), visualized for learners 1 standard deviation (SD) below the mean (“Bad Learners”) and 1 SD above the mean (“Good Learners”) of statistical learning ability. Note that statistical learning ability was a continuous variable—plotting separately above and below the mean is for visualization purposes only.

Statistical learning ability modulates the relation between complexity and attention allocation

In addition to the effects described above, our fixation-time model also included information about participants' visual statistical learning ability on our second task. There was no main effect of statistical learning ability ($F_{(1, 43.9)}=0.00$, $p = 0.995$, $\beta = .0002$, $SE = 0.03$), indicating that people's statistical learning abilities were unrelated to their fixation times as observed in the main task. There was, however, a significant two-way interaction with statistical learning ability and real-time entropy ($F_{(1, 22,886)}=61.96$, $p < .001$, $\beta = 0.04$ $SE = 0.005$), such that better learners looked at more complicated aspects of their environment. There was also an interaction between trial number and statistical learning ability ($F_{(1, 22,917.1)}=24.14$, $p < .001$, $\beta = 0.02$, $SE = 0.005$), such that fixation time was higher at the end of study for good learners than bad (although all learners looked longer at the end of the study, see *Trial-wise fixation time also shifts with experience*). Most interestingly, there was also a significant three-way interaction between statistical learning ability, real-time entropy, and trial number ($F_{(1, 22,893.3)}=4.99$, $p = 0.03$, $\beta = -0.009$, $SE = 0.004$), such that participants' statistical learning ability changed the relationship between real-time entropy and trial number. Specifically, participants with higher statistical learning scores showed less of a shift towards more complex locations over time, relative to participants with lower statistical learning scores. Instead, better learners preferred more complex structures from earlier on in the learning process and maintained that preference throughout learning (*Figure 3b*; note that our statistical learning score is continuous; the split into good and bad learners [± 1 SD from the model prediction mean] is for visualization purposes only).

Discussion

We showed that as learners gain experience in an environment, they attend to different parts of that environment. In Experiment 1, participants' reaction times to search trials in locations with more complex structures decreased more than those in less complex locations. In Experiment 2, we replicated this finding, and further demonstrated with fixation time that participants increased their attention, over the experiment, to locations that were more complex. These results provide strong support for the hypothesis that as learners gain experience with environmental regularities, they attend to successively more complex parts of their input. Moreover, the manner in which our participants shifted their attention varied by statistical learning ability: better learners attended to more complex structures throughout learning, while worse learners initially explored low-complexity structures for more time before shifting to higher complexity locations. Thus, not only do learners reallocate their attention with experience, but the method in which they do so directly relates to how successfully they learn from structured environments.

The finding that participants reallocate their attention as a function of experience complements work showing that learners prioritize stimuli of medium complexity (Cubit et al., 2021; Kidd et al., 2012). As suggested earlier, as a learner gains experience with particular structures, the parts of their environment that were originally of a medium complexity should become increasingly simple. Thus, with experience, more complex structures should shift to being perceived by the learner as less complex. Our results show exactly this using a measure of complexity that is *unique to each subject's own experience*, which prevents having to assume that all participants' experience is equivalent, or that each source of information is equally complex over the course of a learning experience.

Our findings also complement work examining how long-term expertise in a particular domain changes people's preferences. For example, children explore their world to reduce uncertainty about their environments, like by playing with objects that violated their expectations (Stahl & Feigenson, 2015) or learning word-meanings that reduce referential ambiguity (Zettersten & Saffran, 2020). At longer timescales, professional artists prefer more complex artwork than amateur artists (Silvia & Barona, 2009), and an analysis of Beethoven's lifetime work illustrated that the entropy (complexity) of his sonatas increased over time (Daikoku, 2019). This suggests that our finding that attention to complexity changes with short-term experience matches long-term, experience-dependent changes in complexity preferences.

Moreover, an expanding body of literature suggests there are multiple ways to measure statistical learning (Batterink et al., 2015), which may map onto differently accessible kinds of knowledge about the statistical structure. Importantly, these different measures do not often correlate, nor do they report similar amounts of learning. Similar to this, our finding that offline measures of knowledge—which did not differ from chance in this complex learning paradigm—did not relate to our primary dependent measures of online reaction time or eye-fixation, even though trial-wise fixation time changed as a function of experience. This implies that online eye-tracking methodologies may be useful for understanding subtle experience-driven changes to statistical learning, and could be added to a growing battery of indirect measurements which can gauge learning rate (Batterink, 2017; Choi et al., 2020).

In support of this, fixation time did relate to a much simpler, independent measure of statistical learning such that better learners attended to more complex structures throughout learning whereas worse learners initially explored low-complexity structures before shifting to more complex locations. We suggest two possible explanations for this pattern of results. First,

better learners may have learned from simple locations faster than worse learners, thus enabling them to attend to more complex structures earlier. In possible conjunction with this, worse learners may monitor their environment in serial, moving from one structure to another in succession, while better learners may be able to attend to more complex locations consistently by monitoring multiple information sources in parallel; that is, by sampling all of the locations at more regular intervals. Preliminary evidence for this is the smaller magnitude of difference in better learners' looking times to different complexity locations as compared with worse learners' (as visualized in *Fig. 3b* by the larger gap between lines for bad learners than good), but future designs need to be optimized for measuring possible differences in sampling behaviour and linking these directly to learning. Regardless, the counterintuitive idea that better learners may sample more broadly (and therefore be less focused) has implications for development: given children's less focused attention as compared with adults (Plebanek & Sloutsky, 2017; Plude et al., 1994), they may also be capable of monitoring multiple sources simultaneously and therefore attend to more complex locations consistently, like the better learners here. This could be one key to understanding more about children's unique learning advantages (Gualtieri & Finn, 2021).

While it has long been clear that humans are capable of learning the structures present around them, we have demonstrated here that this learning ability interacts with sophisticated attentional selection mechanisms that are sensitive to what a learner has already experienced. In showing that learners shift their attention towards successively more complex aspects of their environment as a function of their own, unique experience, and that this relates to structure learning success, we have prompted a better understanding of a fundamental learning system humans leverage to sort through the immense complexity that surrounds them.

References

- Alamia, A., & Zénon, A. (2016). Statistical regularities attract attention when task-relevant. *Frontiers in Human Neuroscience*, 10(FEB2016), 1–10. <https://doi.org/10.3389/fnhum.2016.00042>
- Armstrong, B. C., Frost, R., & Christiansen, M. H. (2017). The long road of statistical learning research: past, present and future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160047. <https://doi.org/10.1098/rstb.2016.0047>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 1–43. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Machler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Batterink, L. J. (2017). Rapid Statistical Learning Supporting Word Extraction From Continuous Speech. *Psychological Science*, 28(7), 921–928. <https://doi.org/10.1177/0956797617698226>
- Batterink, L. J., Reber, P. J., Neville, H. J., & Paller, K. A. (2015). Implicit and explicit contributions to statistical learning. *Journal of Memory and Language*, 83, 62–78. <https://doi.org/10.1016/j.jml.2015.04.004>
- Berlyne, D. E., Ogilvie, J. C., & Parham, L. C. C. (1971). The dimensionality of visual complexity, interestingness, and pleasingness. *Canadian Journal of Psychology*, 25(3), 195–206. <https://doi.org/10.1037/h0082381>
- Bogaerts, L., Siegelman, N., & Frost, R. (2016). Splitting the variance of statistical learning performance: A parametric investigation of exposure duration and transitional probabilities. *Psychonomic Bulletin and Review*, 23(4), 1250–1256. <https://doi.org/10.3758/s13423-015-0996-z>
- Choi, D., Batterink, L. J., Black, A. K., Paller, K. A., & Werker, J. F. (2020). Preverbal Infants Discover Statistical Word Patterns at Similar Rates as Adults: Evidence From Neural Entrainment. *Psychological Science*, 31(9), 1161–1173. <https://doi.org/10.1177/0956797620933237>
- Cubit, L. S., Canale, R., Handsman, R., Kidd, C., & Bennetto, L. (2021). Visual Attention Preference for Intermediate Predictability in Young Children. *Child Development*, 00(0), 1–13. <https://doi.org/10.1111/cdev.13536>
- Daikoku, T. (2019). Depth and the uncertainty of statistical knowledge on musical creativity fluctuate over a composer's lifetime. *Frontiers in Computational Neuroscience*, 13(April). <https://doi.org/10.3389/fncom.2019.00027>
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, 99(24), 15822–15826. <https://doi.org/10.4324/9780203770351>
- Forest, T. A., Finn, A. S., & Schlichting, M. L. (2021). General precedes specific in memory for structured events. *Journal of Experimental Psychology: General*.
- 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>
- Gualtieri, S., & Finn, A. S. (2021). (Accepted) The sweet spot: When children's developing abilities, brains, and knowledge make them better learners Samantha. *Perspectives on Psychological Science*.

- Hasson, U. (2017). The neurobiology of uncertainty: Implications for statistical learning. In *Philosophical Transactions of the Royal Society B: Biological Sciences* (Vol. 372, Issue 1711). Royal Society of London. <https://doi.org/10.1098/rstb.2016.0048>
- Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49(10), 1295–1306. <https://doi.org/10.1016/J.VISRES.2008.09.007>
- 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), 1–8. <https://doi.org/10.1371/journal.pone.0036399>
- Kinney, D. K., & Kagan, J. (1976). Infant Attention to Auditory Discrepancy. *Child Development*, 47(1), 155–164.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package : Tests in Linear Mixed Effects. *Journal of Statistical Software*. <https://doi.org/10.18637/jss.v082.i13>
- Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments. *Neuron*, 93(2), 451–463. <https://doi.org/10.1016/j.neuron.2016.12.040>
- Li, A. S., & Theeuwes, J. (2020). Statistical regularities across trials bias attentional selection. *Journal of Experimental Psychology: Human Perception and Performance*, 46(8), 860–870. <https://doi.org/10.1037/xhp0000753>
- 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>
- Luce, R. D. (2003). Whatever Happened to Information Theory in Psychology? *Review of General Psychology*, 7(2), 183–188. <https://doi.org/10.1037/1089-2680.7.2.183>
- Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97. <https://doi.org/10.1037/h0043158>
- Misyak, J. B., & Christiansen, M. H. (2012). Statistical Learning and Language: An Individual Differences Study. *Language Learning*, 62(1), 302–331. <https://doi.org/10.1111/j.1467-9922.2010.00626.x>
- Nastase, S., Iacovella, V., & Hasson, U. (2014). Uncertainty in visual and auditory series is coded by modality-general and modality-specific neural systems. *Human Brain Mapping*, 35(4), 1111–1128. <https://doi.org/10.1002/hbm.22238>
- Peirce, J. W. (2008). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2(January), 1–8. <https://doi.org/10.3389/neuro.11.010.2008>
- Plebanek, D. J., & Sloutsky, V. M. (2017). Costs of Selective Attention: When Children Notice What Adults Miss. *Psychological Science*, 28(6), 723–732. <https://doi.org/10.1177/0956797617693005>
- Plude, D. J., Enns, J. T., & Brodeur, D. (1994). The development of selective attention: A life-span overview. *Acta Psychologica*, 86, 227–272.
- Previc, F. H. (1990). Functional specialization in the lower and upper visual fields in humans: Its ecological origins and neurophysiological implications. *Behavioral and Brain Sciences*, 13(3), 519–542. <https://doi.org/10.1017/S0140525X00080018>
- S.R. Research. (2011). *Experiment Builder 1.10.165*.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928. <https://doi.org/10.1126/science.274.5294.1926>
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(4), 623–656. <https://doi.org/10.1002/j.1538-7305.1948.tb00917.x>

- Shannon, C. E. (1951). Prediction and Entropy of Printed English. *Bell System Technical Journal*, 30(1), 50–64. <https://doi.org/10.1109/9780470544242.ch12>
- Siegelman, N., Bogaerts, L., & Frost, R. (2017). Measuring individual differences in statistical learning: Current pitfalls and possible solutions. *Behavior Research Methods*, 49(2), 418–432. <https://doi.org/10.3758/s13428-016-0719-z>
- Siegelman, N., Bogaerts, L., & Frost, R. (2019). What Determines Visual Statistical Learning Performance? Insights From Information Theory. *Cognitive Science*, 43(12). <https://doi.org/10.1111/cogs.12803>
- Silvia, P. J., & Barona, C. M. (2009). Do People Prefer Curved Objects? Angularity, Expertise, and Aesthetic Preference. *Empirical Studies of the Arts*, 27(1), 25–42. <https://doi.org/10.2190/em.27.1.b>
- 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>
- Strange, B. A., Duggins, A., Penny, W., Dolan, R. J., & Friston, K. J. (2005). Information theory, novelty and hippocampal responses: Unpredicted or unpredictable? *Neural Networks*, 18(3), 225–230. <https://doi.org/10.1016/j.neunet.2004.12.004>
- Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin and Review*, 26(4), 1377–1387. <https://doi.org/10.3758/s13423-019-01598-6>
- Wang, B., Samara, I., & Theeuwes, J. (2019). Statistical regularities bias overt attention. *Attention, Perception, and Psychophysics*, 81(6), 1813–1821. <https://doi.org/10.3758/s13414-019-01708-5>
- Zettersten, M., & Saffran, J. R. (2020). Sampling to learn words: Adults and children sample words that reduce referential ambiguity. *Developmental Science*. <https://doi.org/10.1111/desc.13064>
- Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention Is Spontaneously Biased Toward Regularities. *Psychological Science*, 24(5), 667–677. <https://doi.org/10.1177/0956797612460407>
- Zhao, J., & Luo, Y. (2017). Statistical regularities guide the spatial scale of attention. *Attention, Perception, and Psychophysics*, 79(1), 24–30. <https://doi.org/10.3758/s13414-016-1233-1>