

Generalisation of prior information for rapid Bayesian time estimation

Neil W. Roach^{1*}, Paul V. McGraw¹, David Whitaker^{2,3} & James Heron²

1. Visual Neuroscience Group, School of Psychology, The University of Nottingham, United Kingdom, NG7 2RD. 2. Bradford School of Optometry and Visual Science, University of Bradford, United Kingdom, BD7 1DP. 3. School of Optometry and Vision Sciences, Cardiff University, United Kingdom, CF24 4HQ.

Submitted to Proceedings of the National Academy of Sciences of the United States of America

To enable effective interaction with the environment, the brain combines noisy sensory information with expectations based on prior experience. There is ample evidence showing that humans can learn statistical regularities in sensory input and exploit this knowledge to improve perceptual decisions and actions. However, fundamental questions remain regarding how priors are learned and how they generalise to different sensory and behavioural contexts. In principle, maintaining a large set of highly specific priors may be inefficient and restrict the speed at which expectations can be formed and updated in response to changes in the environment. On the other hand, priors formed by generalising across varying contexts may not be accurate. Here we exploit rapidly induced contextual biases in duration reproduction to reveal how these competing demands are resolved during the early stages of prior acquisition. We show that observers initially form a single prior by generalising across duration distributions coupled with distinct sensory signals. In contrast, they form multiple priors if distributions are coupled with distinct motor outputs. Together, our findings suggest that rapid prior acquisition is facilitated by generalisation across experiences of different sensory inputs, but organised according to how that sensory information is acted upon.

Bayesian inference | Time perception | Sensorimotor learning

Like all complex animals, humans rely on their senses to extract information about the environment and guide decision making and behaviour. Often however, sensory information is ambiguous. Signals transmitted to the senses can be weak or degraded, such as patterns of reflected light under low illumination or speech sounds in noisy environments. Moreover, sensory representations of even the most high-fidelity signals tend to be variable (1) and are insufficient to completely disambiguate different distal causes (2). Mounting empirical evidence indicates that when forming decisions and planning actions, the brain combines uncertain sensory information with expectations based on prior knowledge (3-6). For example, a variety of biases in visual perception have been shown to be consistent with reliance on prior knowledge regarding statistical regularities in the environment, such as the distribution of local orientations (7) and speeds (8) in natural scenes and the positioning of light sources (9). In many studies, perception and behaviour have been shown to be well described by near-optimal integration of sensory evidence and prior knowledge according to the principles of statistical decision theory.

Prior knowledge can be acquired over a range of different timescales. Priors specifying stable statistical characteristics of the environment are typically thought to be either innate or the consequence of life-long implicit learning (5,7). However, context-specific priors can also be formed based on recent experiences. Studies using simple sensorimotor tasks suggest that human participants are adept at learning the distribution of sets of stimuli and integrating this acquired knowledge with sensory evidence when making decisions and planning actions (e.g. 4, 10-13). For example, a classic finding in temporal reproduction experiments is that judgments converge towards a central value - participants consistently overestimate shorter durations and underestimate

longer ones (14-15). The centre of convergence at which there is no constant reproduction error (termed the *'indifference point'*) is not fixed, but rather shifts depending on the set of durations presented (16-17). Recent studies have demonstrated that these central tendency biases can be parsimoniously explained by the integration of noisy sensory estimates with acquired knowledge of the stimulus distribution in a manner that maximises reproduction precision (17-19).

Considerable advances have been made toward understanding how priors are learned and updated over time (20-23) and the degree of complexity in stimulus distributions that can be represented (19, 23-25). However, our understanding of how prior knowledge is organised by the brain remains limited. Recent research indicates that participants provided with extended training are able to learn multiple priors for stimuli presented in different contexts. For example, Kerrigan and Adams (26) demonstrated that distinct light position priors can be learned for different coloured illumination in a few hours. Similarly, Gekas and colleagues (24) showed that distinct priors for motion direction can be learned for sets of dot stimuli of different colours. Multiple prior learning has also been demonstrated using stimulus sets demarcated by way of spatial location (27) and symbolic visual cues (28), suggesting that sensory specificity may be a key organising principle for representations of prior knowledge.

In principle, maintaining high levels of prior specificity should ensure that expectations about different objects and events in the external environment remain accurate, even when they are characterised by distinct statistical properties. However, application of this strategy in complex real-world environments poses fundamental challenges. In a typical sensorimotor experiment, subjects are presented with stimuli that vary along a small number of dimensions and asked to perform a single task. In contrast,

Significance

When sensory information is uncertain, humans rely on prior knowledge of the environment when making perceptual judgments and planning actions. Humans can rapidly learn new priors, but the rules governing how recent experiences are grouped together are unknown. Here we show that observers initially form single priors by generalising across distinct sensory signals, but form multiple priors when they are coupled with different motor outputs. Thus, our results reveal that internal models of temporal statistics are initially organised not according to the nature of the sensory input, but the way it is acted upon.

Reserved for Publication Footnotes

137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204

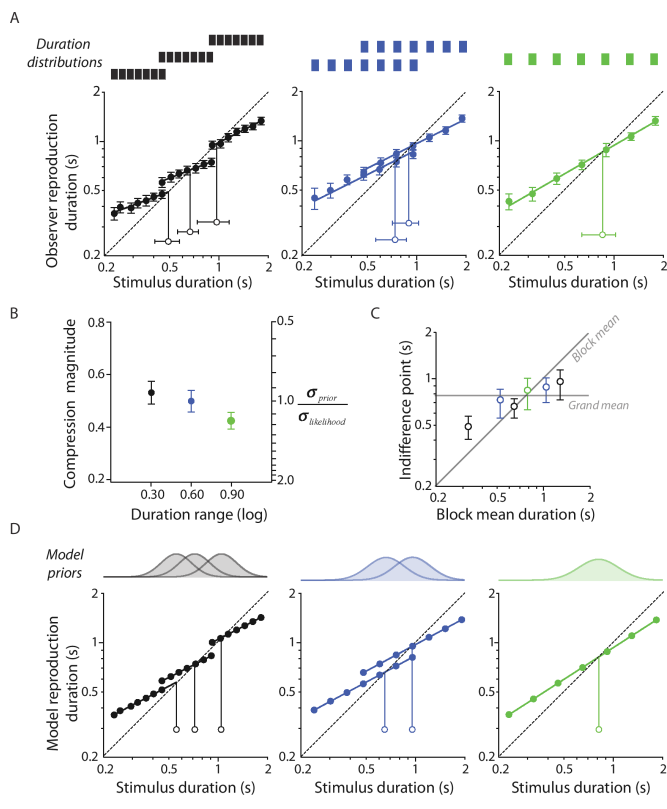


Fig. 1. Rapidly induced central tendency biases in human and Bayesian model observers. **(A)** Mean reproduction durations as a function of stimulus duration (filled symbols) for six stimulus distributions presented in separate testing blocks. Error bars show ± 1 SEM calculated across observers. Solid lines show best fitting power functions for each dataset, whereas the dotted diagonal line denotes veridical (unbiased) performance. Unfilled symbols in this and subsequent figures represent the estimated indifference points (centre of compressive bias) along with bootstrapped 95% confidence intervals. **(B)** Mean compression magnitude ($1 - \text{power exponent}$), plotted as a function of duration range (error bars show 95% CIs). The color of data points indicates associated conditions in Figure 1A and 1B. The right hand ordinate scale shows the ratio of prior standard deviation to likelihood standard deviation (both log-scaled) required to produce equivalent levels of bias in the Bayesian model. **(C)** Estimated indifference points as a function of the mean duration for a given stimulus distribution (error bars show 95% CIs). The solid horizontal line indicates the global mean calculated across all conditions. **(D)** Simulated data for a Bayesian observer-actor that combines sensory evidence with acquired knowledge of the stimulus distribution (for details see Supporting Information). Inset distributions represent inferred approximations of the prior distribution for each condition (see Supporting Information for more details).

people encounter a wide range of sensory inputs outside the laboratory and use this information to interact with the environment in a variety of ways. Trying to learn a specific prior for every combination of stimulus and behavioural context would be unfeasible, since no two situations are ever identical. Moreover, increased specificity in prior representations has the negative side-effect of limiting the rate at which priors can be learned and updated in response to changes in the environment.

Here we exploit central tendency biases in duration reproduction to investigate how these competing demands are resolved during the early stages of prior acquisition. We first demonstrate that individuals can rapidly form priors that approximate the distribution of stimulus durations presented within a single testing session. To reveal the rules governing the initial structuring of prior knowledge, we then interleave stimuli from duration distributions that result in distinct central-tendency biases when presented in separate sessions. Across a series of experiments, we

manipulate the sensory characteristics of each set of stimuli as well as the nature of the motor response required.

Results

Rapidly induced central tendency biases in duration reproduction.

Observers were asked to reproduce the duration of a simple visual stimulus (isotropic Gaussian patch) by pressing and holding a key for an equivalent interval of time. Within each 70-trial testing block, stimuli were presented for one of 7 different durations, sampled from a log-uniform distribution. Filled symbols in Figure 1A show mean reproduction durations for six different distributions of varying range and central tendency. Compared to the dashed diagonal lines denoting veridical performance, clear evidence of a compressive bias can be seen in each condition. To summarise the pattern of biases observed, we fitted each dataset with a power function and derived two statistics: the magnitude of the compressive bias ($1 - \text{exponent}$ of the power function) and the indifference point (the duration at which the fitted power function intersects the diagonal, see Supporting Information for details). Both metrics were systematically affected by short-term manipulation of the stimulus distribution, albeit in different ways. The magnitude of compression was dependent on the variability of durations presented within a given testing block, with greater mean compression observed for smaller duration ranges than larger ones (Figure 1B). Indifference points were pulled towards the mean duration encountered in the associated testing block (Figure 1C).

In keeping with recent suggestions (18-19), we found that these biases can be well described by a Bayesian ideal observer model in which acquired knowledge of the stimulus distribution is used to support inferences made about stimulus duration when available sensory measurements are imprecise. Rather than learning the discrete physical stimulus distribution, we assumed that observers acquire a smoothed approximation of the distribution (see Supporting Information). Figure 1D shows simulated datasets for each duration distribution condition, along with the best-fitting priors. Differences in the magnitude of compression between conditions are captured by a change in the width of the prior, consistent with previous evidence showing that observers can implicitly learn the variance of a stimulus distribution (22-23, 29). Changes in the indifference point are consistent with a lateral shift in the prior distribution towards recently presented stimulus values. While extensive training can result in priors that accurately approximate the stimulus distribution (17), the limited opportunity afforded by each of our brief testing blocks resulted in partial prior recalibration - note how indifference points fall close to the mid-point between the mean of the relevant duration distribution and that of the entire stimulus set. Critically however, the fact that quantifiable changes in the pattern of bias can be induced over this short timescale provides us with a means of probing the characteristics of the underlying prior-updating mechanisms.

Generalisation across distributions paired with distinct sensory inputs.

To investigate generalisation during rapid prior formation, we took the approach of interleaving stimuli drawn from duration distributions that produce distinct patterns of reproduction bias when presented in separate blocks. The filled black symbols in Figure 2A show mean duration reproductions for two such stimulus sets, which have been replotted from the left panel of Figure 1A. Indifference points for these short (S) and long (L) duration distributions are significantly different (open black symbols), confirming that prior expectations differed between testing blocks. A contrasting pattern of results was obtained when the two distributions were randomly intermixed within a testing block (red symbols). Interleaving increased the overestimation

205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272

273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340

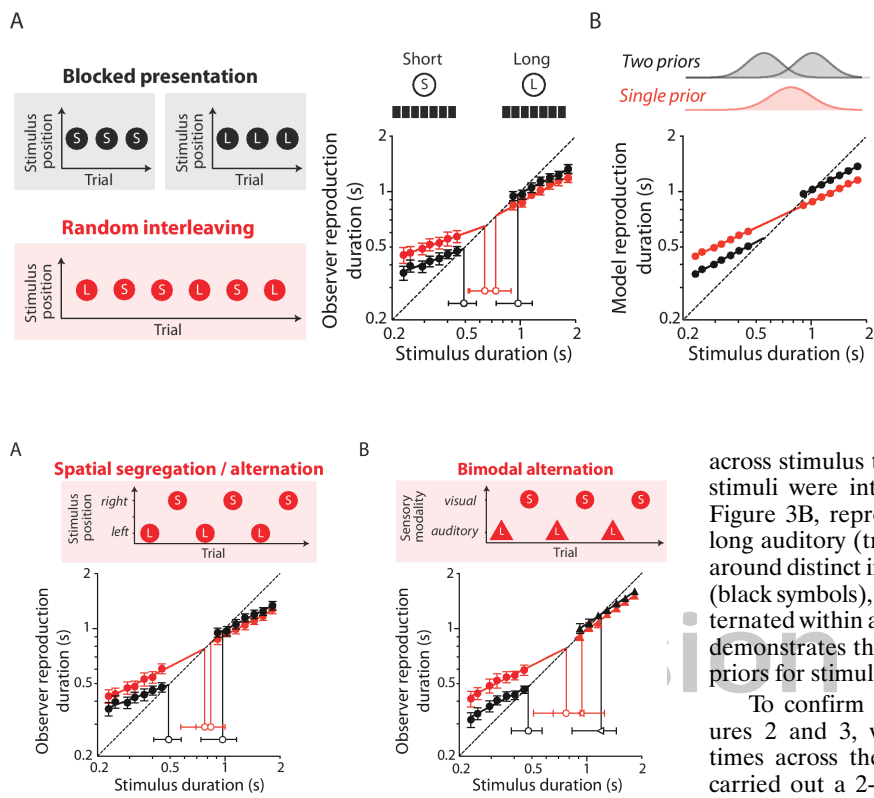


Fig. 2. Generalisation across interleaved duration distributions. **(A)** Stimulus durations sampled from short (S) and long (L) distributions were either presented in blocks (black) or randomly interleaved (red). Interleaving altered the pattern of biases, producing a convergence of indifference points (unfilled symbols) towards a common duration. **(B)** Simulated performance of the Bayesian observer-actor model for blocked (black) and interleaved (red) conditions. In the interleaved condition, a singular prior was implemented by generalising across both stimulus distributions.

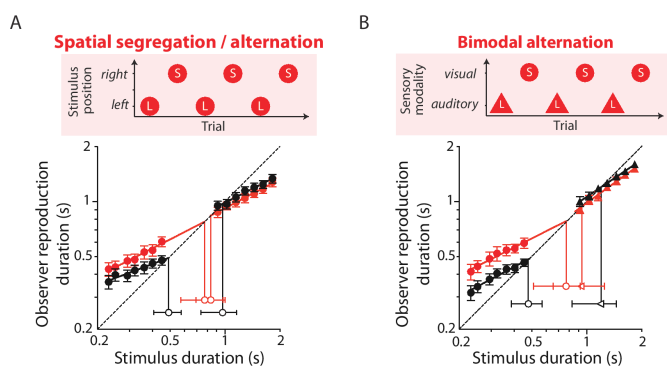


Fig. 3. Rapidly learned priors generalise across distributions paired with distinct sensory inputs. **(A)** Mean reproduced durations for interleaved presentation in which short (S) and long (L) distributions were presented at different locations and alternated from trial-to-trial to remove spatial and temporal uncertainty (red symbols). For comparison black symbols show data obtained when stimulus distributions were presented in separate blocks (replotted from Figure 1A). **(B)** Mean reproduced durations of visual (circular symbols) and auditory (triangular symbols) under blocked (black) and alternated (red) presentation conditions. Error bars for mean reproductions show $\pm 1\text{SEM}$, whereas error bars for indifference points show bootstrapped 95% CIs.

bias in reproduction of short durations and underestimation bias in long durations, resulting in indifference points which converge towards a point of central tendency for the combined stimulus range. These results are consistent with the formation of a single composite prior distribution, and can be parsimoniously modeled using the same prior as previously inferred from results with a wide continuous duration range (Figure 2B). This convergence of indifference points provides a useful marker of when prior formation generalises across the two interleaved stimulus sets; if observers formed separate priors for each stimulus set, no difference in the pattern of biases would be expected between blocked and interleaved conditions.

Since observers did not have a means of distinguishing the different stimulus sets in the random interleaving condition, it perhaps unsurprising that acquired prior expectations were formed by generalising across them. Interestingly however, we found that the same pattern of results holds even if the stimuli used for the two interleaved duration distributions are clearly discriminable. Figure 3A shows results obtained when short and long visual stimuli where presented at different spatial locations (left and right of fixation) and alternated from trial to trial to remove any spatial and temporal uncertainty regarding set membership. Despite these changes to the stimulus sequence, indifference points in the interleaved condition again converged towards a common central duration, suggesting that observers' prior expectations did not distinguish between the two clearly demarcated stimulus sets. Even more striking evidence for generalisation

across stimulus type was found when sets of visual and auditory stimuli were interleaved within a single session. As shown in Figure 3B, reproduction of short visual (circular symbols) and long auditory (triangular symbols) produced compressive biases around distinct indifference points when tested in separate blocks (black symbols), but converged towards a common point when alternated within a single testing session (red symbols). This finding demonstrates that observers do not automatically form separate priors for stimuli presented to different sensory modalities.

To confirm the statistical robustness of the results in Figures 2 and 3, we computed participants' mean reproduction times across the seven durations in each duration range and carried out a 2-way repeated measures ANOVA. In each case we found a significant interaction between presentation condition (blocked or interleaved) and duration range (short or long) (Random interleaving: $F(1,8)=17.75, p<0.05$; Spatial segregation/alternation: $F(1,8)=13.84, p<0.05$; Bimodal alternation: $F(1,8)=9.47, p<0.05$). Decomposition of these significant interactions into simple effects revealed that interleaving significantly increased reproduction times for the short range condition in each experiment (see Table S1). Interleaving had a more modest, contracting effect on reproduction times in the long range, which reached statistical significance in one of the three conditions. Together, these results demonstrate a clear tendency for rapidly acquired duration priors to generalise across different stimulus types, even when they are clearly discriminable and are characterised by distinct temporal statistics.

Our results differ from previous studies, where participants have been shown to acquire stimulus-specific priors when provided with more extensive training (24, 26-27). To check that this also occurs in our experimental paradigm, we conducted an additional experiment in which observers completed repeated testing sessions with alternative presentation of short and long stimulus sets at different spatial locations (Figure 4A). In early sessions (red symbols), biases converged towards a common central duration, replicating the reliance on a single, generalised prior in Figure 3A. However, indifference points systematically diverged with further training (blue symbols), consistent with the gradual employment of stimulus-specific priors over time. This transition is clearly illustrated in Figure 4B, which plots changes in the ratio of indifference points for short and long stimulus sets over the course of testing.

No generalisation across distributions paired with different motor responses

In striking contrast to the results obtained when interleaving different stimulus types, we found no evidence of prior generalisation when stimulus sets were paired with different types of motor response. Figure 5A shows results obtained when observers were only required to make a duration reproduction for one of the two interleaved stimulus sets. Short and long visual stimuli were alternatively presented to the left and right of fixation and observers

341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408

409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476

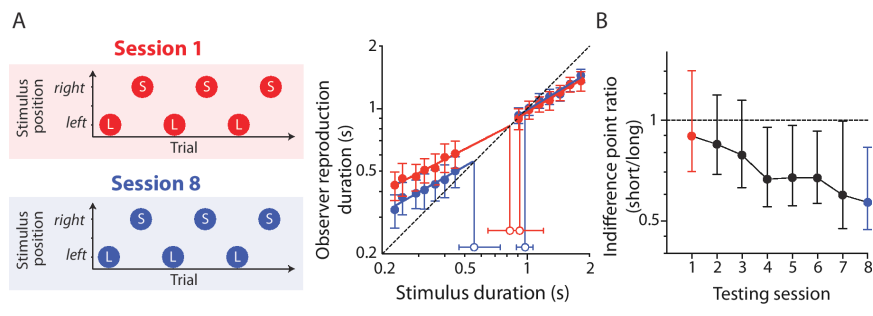


Fig. 4. Emergence of stimulus-specificity with extended training. **(A)** Observers were required to manually reproduce the duration of short (S) and long (L) visual stimuli presented alternatively to the right and left of fixation. Each observer completed 8 sessions, each comprising 140 trials. Mean reproduced durations in the first session (red filled symbols) replicate the pattern of compressive bias shown in Figure 3A. Indifference points (unfilled red symbols) for the two duration distributions converge towards a common duration, consistent with reliance on a single generalised prior. However by the final session, mean reproduced durations for short and long stimuli (blue filled symbols) were compressed around distinct indifference points (blue unfilled symbols). Error bars for mean reproductions show $\pm 1\text{SEM}$, whereas error bars for indifference points show bootstrapped 95% confidence intervals. **(B)** Ratio of indifference points for short and long duration distributions, plotted as a function of testing session. The systematic divergence of indifference points suggests that observers' prior expectations become increasingly location-specific over time. Error bars indicate bootstrapped 95% confidence intervals.

477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544

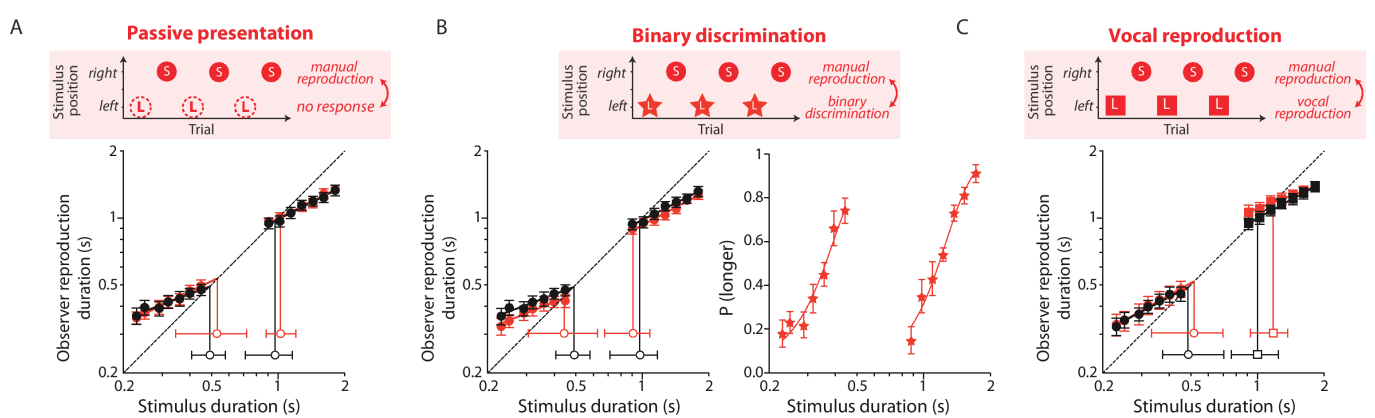


Fig. 5. Rapidly learned priors distinguish between distributions paired with different motor responses. **(A)** Mean reproduction durations for short (S) and long (L) distributions when presented in separate blocks (black symbols) or when reproduction trials for one stimulus set were alternated with passive presentation of the other stimulus set (red symbols). In contrast to the data shown in Figure 2, there is no systematic convergence of indifference points (open symbols) in the interleaved condition. **(B)** The left panel shows mean reproduction durations for a control experiment in which reproduction trials for one stimulus set were alternated with trials requiring binary duration discrimination judgments on the other stimulus set (red symbols). For comparison black symbols show reproduction data obtained when short and long duration distributions were presented in separate blocks. The right panel shows the group-mean psychometric functions for the binary judgments trials, along with the best-fitting logistic functions. The proportion of trials on which stimuli were judged to be longer than the mean duration of the stimulus set is plotted as a function of stimulus duration. **(C)** Mean reproduced durations using button press (circular symbols) and vocalisation (square symbols) methods, measured in separate blocks (black symbols) or interleaved presentation (red symbols). Error bars for mean reproductions and binary choice probability show $\pm 1\text{SEM}$, whereas error bars for indifference points show bootstrapped 95% CIs.

were instructed to reproduce the duration of stimuli presented on one side, whilst passively observing stimuli presented at the other spatial location. Reproduction data for both short and long duration distributions was obtained by reversing the mapping between spatial location and response in different sessions. Despite the stimulus sequences being identical to that described previously (e.g. Figure 3A), biases obtained in interleaved conditions were indistinguishable from those obtained with blocked presentation (condition \times duration range interaction: $F(1,8)=0.60$, $p>0.05$, see Table S2). This result demonstrates that prior formation is not a simple function of sensory history. Rather, it is only the set of previously reproduced stimuli that determines observers' prior expectations on subsequent reproduction trials; passive trials do not contribute.

A potential limitation of asking observers to passively view a set of stimuli is that it is difficult to know with any certainty how well these stimuli were attended. To address this concern, we conducted a control experiment in which observers made binary judgments about the duration of stimuli presented to one side of fixation ('is this stimulus shorter or longer than the average duration seen at that location?'), alternated with reproductions of

the duration of stimuli presented to the other. As shown in Figure 5B, responses on binary judgment trials were sufficient to allow construction of systematic psychometric functions for duration discrimination, confirming that observers attending to the non-reproduced stimuli, forming estimates of stimulus duration and maintaining an accurate internal representation of the mean. As was the case for passive presentation however, the duration of these stimuli did not systematically impact upon the pattern of biases in interleaved reproduction trials (condition \times duration range interaction: $F(1,8)=0.99$, $p>0.05$). Thus, we can be confident that the lack of generalisation across stimulus sets reflects the fact that they were paired with different types of response.

To further test the limits of response-specificity during rapid prior formation, we also investigated the effects of interleaving duration reproduction trials requiring different motor responses. Observers were asked to reproduce the duration of some stimuli by making a vocal utterance ('ba') into a microphone for a corresponding length of time. The duration of speech segments was extracted from audio recordings of each trial using an automated speech detection algorithm (see Experimental Procedures). As shown in Figure 5C, the duration of observers' vocal reproduc-

tions (square symbols) showed compressive biases in blocked conditions comparable to those observed with button press responses (circular symbols). In interleaved conditions (red symbols), observers reproduced the stimulus duration on every trial, but alternated between vocal and keypress methods (paired with stimuli presented to the left and right of fixation respectively). Again, we found no systematic difference in the pattern of biases observed with blocked and interleaved presentation (condition \times duration range interaction: $F(1,8)=0.26, p>0.05$), indicating an absence of generalisation across experiences of stimuli paired with different types of reproduction response. This experimental result is important for two reasons. First, it indicates that it is not simply the types of judgement made by observers that is critical for driving response-specificity in prior expectations, but also the specific motor actions involved in executing those judgements. It also provides a strong demonstration that observers concurrently maintain multiple motor-specific priors in interleaved conditions because, unlike passive and binary discrimination experiments, reproduction data for both short and long duration distributions was collected within the same session.

Discussion

Generalisation is an essential ingredient in the formation of prior expectations within complex environments. In everyday life we encounter a diverse range of sensory stimuli, making it impossible to learn statistical regularities without grouping experiences together in some meaningful way. Our results indicate that recent experiences of stimulus duration are initially routed into separate prior representations not according to the type of sensory input, but according to the way in which observers act upon this input. Observers consistently formed single priors by generalising across stimulus distributions with distinct statistics, even when they were clearly distinguished by way of spatial location, temporal sequence or sensory modality. We propose that this broad generalisation across sensory inputs is a default mode which acts to widen the data acquisition 'net' for initial prior acquisition, allowing approximations of stimulus distributions to be rapidly learned and modified. This strategy is not fixed however - in line with previous findings (24, 26-27), we find that observers are able to learn stimulus-specific priors with extended training. Together, this work indicates that the structuring of prior knowledge is dynamic and that emphasis shifts from flexibility to specificity as learning progresses. Characterising the dynamics of this transition provides a challenge for future behavioural studies, as well as computational models describing how priors are learned and updated over time (20-23). An appealing idea is that the brain performs a type of cluster analysis, partitioning sensory inputs into groups with different statistical properties. The emergence of prior specificity may reflect the accumulation of sufficient evidence to justify clustering.

In contrast to the generalisation observed across different sensory inputs, interval duration priors appear to be specific to motor response from the outset of learning. We found no evidence of generalisation across passive and active trials, trials requiring different forms of duration judgement, or trials requiring reproduction judgements involving different motor systems. In these experiments, we were careful to use identical stimulus sequences to those that produced generalisation when paired with a consistent manual reproduction response. Accordingly, we can be confident that the dissociation between stimulus and response generalisation is not an artefact of changes in the experimental design, such as the complexity of the duration distributions (19; 23-25) or discriminability of the stimulus sets. Whether this strategy provides any functional advantage is not yet clear. One possibility is that the multi-dimensional nature of sensory input heightens the chance of forming redundant prior representations based on uninformative distinctions between stimulus character-

istics. Initially grouping prior experiences according to a smaller set of goal-directed actions may be less problematic in this regard.

At present, our understanding of the neural substrates of prior knowledge is limited. Priors specifying long-term structural regularities in the environment have been suggested to be implicitly encoded in sensory representations, via inhomogeneities in the number, distribution and/or bandwidth of tuning curves (7, 30) or patterns of spontaneous firing activity (31-32). In contrast, recent functional neuroimaging results suggest that recently acquired priors are represented in distinct brain regions to sensory likelihoods (33). But why should actions be used to partition prior knowledge of interval duration?

There is growing evidence indicating that the encoding of temporal information relies upon brain areas that are involved in motor control, most notably the basal ganglia, cerebellum and supplementary motor area (SMA) (34-35). In many instances there appears to be a tight coupling between sensory processing and motor planning. For example, it has been proposed that the SMA can be subdivided into two structurally and functionally distinct areas: a rostrally located pre-SMA involved in encoding of temporal structure and a more caudally positioned SMA-proper, which uses this information to implement motor actions (36). Interestingly, groups of neurons in SMA have been shown to exhibit interval tuning that is invariant to sensory modality (37) and activity that is action selective (38). Recruitment of neuronal populations with similar properties during early prior acquisition could potentially underlie the pattern of specificity and generalisation found in the present study. To test this hypothesis, future studies could take a similar approach to characterise the selectivity of rapidly acquired priors in non-temporal tasks. Central-tendency biases provide a particularly useful testbed for examining these issues; robust effects have been reported across a wide range of magnitude estimation tasks including length, distance and loudness (16, 39).

If motor-specificity proves to be a general property of prior acquisition in temporal and non-temporal domains, it would imply broader integration of sensory analysis and motor control in the brain. In Bayesian models of magnitude estimation, perceptual decision making and action planning are typically implemented in serial stages: likelihoods and priors are first combined to obtain a posterior probability distribution over the possible states of the world; an estimate is then formed taking into account the relative costs and benefits associated with alternative choices; and finally an appropriate motor action is planned and executed (17, 18, 20). However, an alternative suggestion is that the brain continuously processes sensory information to specify several potential actions in parallel (40-42). According to this 'intentional' framework of information processing, perceptual decision making and motor planning are intrinsically linked. In support of this view, activity consistent with accumulation of evidence for a decision has been identified in a variety of neural populations tuned for motor output parameters (42-43). Moreover, decision-related activity has been shown to occur in parallel in areas associated with different motor responses (44-46). Motor specificity could arise naturally if signals used to form and update prior representations are distributed in a similar manner.

Materials and Methods

Observers.

A total of 27 observers participated in the study, comprising one of the authors (NWR) and 26 participants who were naive to the specific purpose of the experiments (age range 19-36). All had normal, or corrected visual acuity and no history of hearing difficulties. 6-9 observers completed each of the sub-experiments (see Table S4 for complete break down), collecting a minimum of 140 trials per combination of duration distribution and condition (total trials >30,000). All experiments were conducted with the written consent of each observer and in accordance with the School of Psychology Ethics Committee at the University of Nottingham.

Stimuli.

Visual stimuli were isotropic Gaussian patches ($\sigma = 1$ deg.) generated in Matlab using PsychoPhysics Toolbox extensions (47-48) and presented on a linearised 22 inch NEC MultiSync FP1370 monitor (100Hz framerate, 47 cd/m² background luminance). Viewing distance was fixed via headrest at 103cm at which 1 pixel subtended 1 arcmin of visual angle. In different experiments visual stimuli were centered either at fixation or 3.33 deg. to the left or right of fixation. Auditory stimuli were 500Hz tones with 5ms cosine on and off ramps, generated at a sampling rate of 44.1KHz and presented diotically via Sennheiser HD-265 headphones at approximately 60 dB SPL.

Procedure.

Stimulus durations were sampled from discrete log-uniform distributions, each comprising 7 levels. For the three distributions shown in Figure 1A, the median durations (i.e. 4th levels) were 320ms, 640ms and 1280ms respectively and duration levels were each separated by 0.05 log units. For Figure 1B the distributions were centered on 480ms and 960ms with a step size of 0.1 log unit, and Figure 1C depicts a distribution centered on 640ms with a 0.15 log unit step size. For blocked conditions, each of the 7 stimulus durations were presented 10 times in a pseudo-random order.

For the experiments depicted in Figures 2-5, only the small step size distributions (0.05 log units) centered on 320ms (short) and 1280ms (long) were used. Interleaved sessions comprised 140 trials, 70 for each duration range. The order of presentation was either pseudo-randomised (Figure 2A) or else alternated from trial to trial (all other experiments). For all conditions, the delay between a response and the presentation of the next stimulus was

jittered between 500ms and 1200ms. Observers were not informed as to the nature of the duration distributions and no feedback was provided.

On manual reproduction trials, observers were instructed to press and hold a button on the keyboard with their index finger to indicate the perceived duration of the stimulus. Vocal reproduction was achieved by asking observers to make vocalisations of the speech sound /ba/ which were recorded using a Coles 4104 Ribbon Microphone. The duration of vocal reproductions was extracted using an automated voice activity detection algorithm from the VOICEBOX speech processing toolbox (49). On binary discrimination trials, observers judged whether a stimulus was longer or shorter than the mean duration seen at that location by pressing one of two keyboard buttons (method of single stimuli, 50). To provide observers the opportunity to build up a stable internal representation of the mean, the first seven trials consisted of one presentation of each of the duration level in the set and responses from these trials were discarded.

Bayesian model.

To simulate patterns of reproduction bias, we assumed that observers combine sensory estimates of stimulus duration with a learned approximation of the stimulus distribution. A full description of the model and procedures for parameter setting can be found in Supporting Information.

Acknowledgements

This work was supported by a Wellcome Trust Research Fellowship to NWR (WT097387). We thank Thomas Veale for assisting with data collection.

- Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nat Rev Neurosci*, 9(4), 292–303.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference. *Ann Rev Psychol*, 55(1), 271–304.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci*, 27(12), 712–719.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247.
- Seriès, P., & Seitz, A. R. (2013). Learning what to expect (in visual perception). *Front Hum Neurosci*, 42, 2757–2772.
- Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn Sci*, 10(7), 319–326.
- Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat Neurosci*, 14(7), 926–932.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nat Neurosci*, 5(6), 598–604.
- Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. *Nat Neurosci*, 7(10), 1057–1058.
- Miyazaki, M., Yamamoto, S., Uchida, S., & Kitazawa, S. (2006). Bayesian calibration of simultaneity in tactile temporal order judgment. *Nat Neurosci*, 9(7), 875–877.
- Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *J Neurosci*, 26(40), 10154–10163.
- Chalk, M., Seitz, A. R., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *J Vis*, 10(8), 2:1–18.
- Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the origins of suboptimality in human probabilistic inference. *PLOS Comp Biol*, 10(6), e1003661.
- Lejeune, H., & Wearden, J. H. (2009). Vierordt's The Experimental Study of the Time Sense (1868) and its legacy. *Eur J Cogn Psychol*, 21(6), 941–960.
- Vierordt, K. (1868). *Der Zeitsinn nach Versuchen*. Tübingen: Laupp.
- Hollingworth, H. L. (1910). The central tendency of judgment. *J Phil, Psychol Sci Meth*, 7(17), 461–469.
- Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nat Neurosci*, 13(8), 1020–1026.
- Cicchini, G. M., Arrighi, R., Cecchetti, L., Giusti, M., & Burr, D. C. (2012). Optimal encoding of interval timing in expert percussionists. *J Neurosci*, 32(3), 1056–1060.
- Acerbi, L., Wolpert, D. M., & Vijayakumar, S. (2012). Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLOS Comp Biol*, 8(11), e1002771.
- Petzschner, F. H., & Glasauer, S. (2011). Iterative Bayesian estimation as an explanation for range and regression effects: a study on human path integration. *J Neurosci*, 31(47), 17220–17229.
- Kwon, O.-S., & Knill, D. C. (2013). The brain uses adaptive internal models of scene statistics for sensorimotor estimation and planning. *Proc Natl Acad Sci USA*, 110(11), E1064–73.
- Verstynen, T., & Sabes, P. N. (2011). How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J Neurosci*, 31(27), 10050–10059.
- Berniker, M., Voss, M., & Körding, K. P. (2010). Learning priors for Bayesian computations in the nervous system. *PLoS ONE*, 5(9): e12686.
- Gekas, N., Chalk, M., Seitz, A. R., & Seriès, P. (2013). Complexity and specificity of experimentally-induced expectations in motion perception. *J Vision*, 13(4), 8.
- Sanborn, A. N., & Beierholm, U. R. (2016). Fast and Accurate Learning When Making Discrete Numerical Estimates. *PLOS Comp Biol*, 12(4), e1004859.
- Kerrigan, I. S., & Adams, W. J. (2013). Learning different light prior distributions for different contexts. *Cognition*, 127(1), 99–104.
- Nagai, Y., Suzuki, M., Miyazaki, M., & Kitazawa, S. (2012). Acquisition of multiple prior distributions in tactile temporal order judgment. *Front Psychol*, 3, 276.
- Petzschner, F. H., Maier, P., & Glasauer, S. (2012). Combining symbolic cues with sensory input and prior experience in an iterative Bayesian framework. *Front Int Neurosci*, 6, 58.
- Sohn, H., & Lee, S. H. (2013). Dichotomy in perceptual learning of interval timing: calibration of mean accuracy and precision differ in specificity and time course. *J Neurophys*, 109(2), 344–362.
- Wei, X.-X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain "anti-Bayesian" percepts. *Nat Neurosci*, 18(10), 1509–1517.
- Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013), 83–87.
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn Sci*, 14(3), 119–130.
- Vilares, I., Howard, J. D., Fernandes, H. L., Gottfried, J. A., & Körding, K. P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. *Curr Biol*, 22(18), 1641–1648.
- Coull, J. T., Cheng, R.-K., & Meck, W. H. (2011). Neuroanatomical and neurochemical substrates of timing. *Neuropsychopharmacology*, 36(1), 3–25.
- Merchant, H., Harrington, D. L., & Meck, W. H. (2013). Neural basis of the perception and estimation of time. *Ann Rev Neurosci*, 36(1), 313–336.
- Kotz, S. A. E., & Schwartz, M. (2011). Differential input of the supplementary motor area to a dedicated temporal processing network: functional and clinical implications. *Front Int Neurosci*, 5, 86.
- Merchant, H., Pérez, O., Zarco, W., & Gámez, J. (2013). Interval tuning in the primate medial premotor cortex as a general timing mechanism. *J Neurosci*, 33(21), 9082–9096.
- Fujii, N., Mushiaki, H., & Tanji, J. (2002). Distribution of eye- and arm-movement-related neuronal activity in the SEF and in the SMA and Pre-SMA of monkeys. *J Neurophysiol*, 87(4), 2158–2166.
- Petzschner, F. H., Glasauer, S., & Stephan, K. E. (2015). A Bayesian perspective on magnitude estimation. *Trends Cogn Sci*, 19(5), 285–293.
- Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos Trans R Soc of London B, Biol Sci*, 362(1485), 1585–1599.
- Shadlen, M. N., Kiani, R., Hanks, T. D., & Churchland, A. K. (2008). Neurobiology of decision making: An intentional framework. In C. Engel & W. Singer (Eds.), *Better than Conscious Decision Making, the Human Mind, and Implications for Institutions* (pp. 71–101). Cambridge, MA.
- Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Ann Rev Neurosci*, 33, 269–298.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Ann Rev Neurosci*, 30(1), 535–574.
- Calton, J. L., Dickinson, A. R., & Snyder, L. H. (2002). Non-spatial, motor-specific activation in posterior parietal cortex. *Nat Neurosci*, 5(6), 580–588.
- Cui, H., & Andersen, R. A. (2007). Posterior parietal cortex encodes autonomously selected motor plans. *Neuron*, 56(3), 552–559.
- de Lafuente, V., Jazayeri, M., & Shadlen, M. N. (2015). Representation of accumulating evidence for a decision in two parietal areas. *J Neurosci*, 35(10), 4306–4318.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spat Vis*, 10(4), 433–436.
- Kleiner, M., Brainard, D., & Pelli, D. (2007) What's new in Psychtoolbox-3? *Perception*, 36 Supplement, 14.
- Brookes, M. ed. VOICEBOX: Speech Processing Toolbox for MATLAB Available at: <http://www.ee.ic.ac.uk/hp/staff/dmb/voicebox/voicebox.html>.
- Morgan, M. J., Watamaniuk, S. N., & McKee, S. P. (2000). The use of an implicit standard for measuring discrimination thresholds. *Vision Res*, 40(17), 2341–2349.
- Gibbon, J. (1992). Ubiquity of scalar timing with a poisson clock. *J Math Psychol*, 35, 283–293.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. *Ann N Y Acad Sci*, 423, 52–77.