1. Supplemental Data

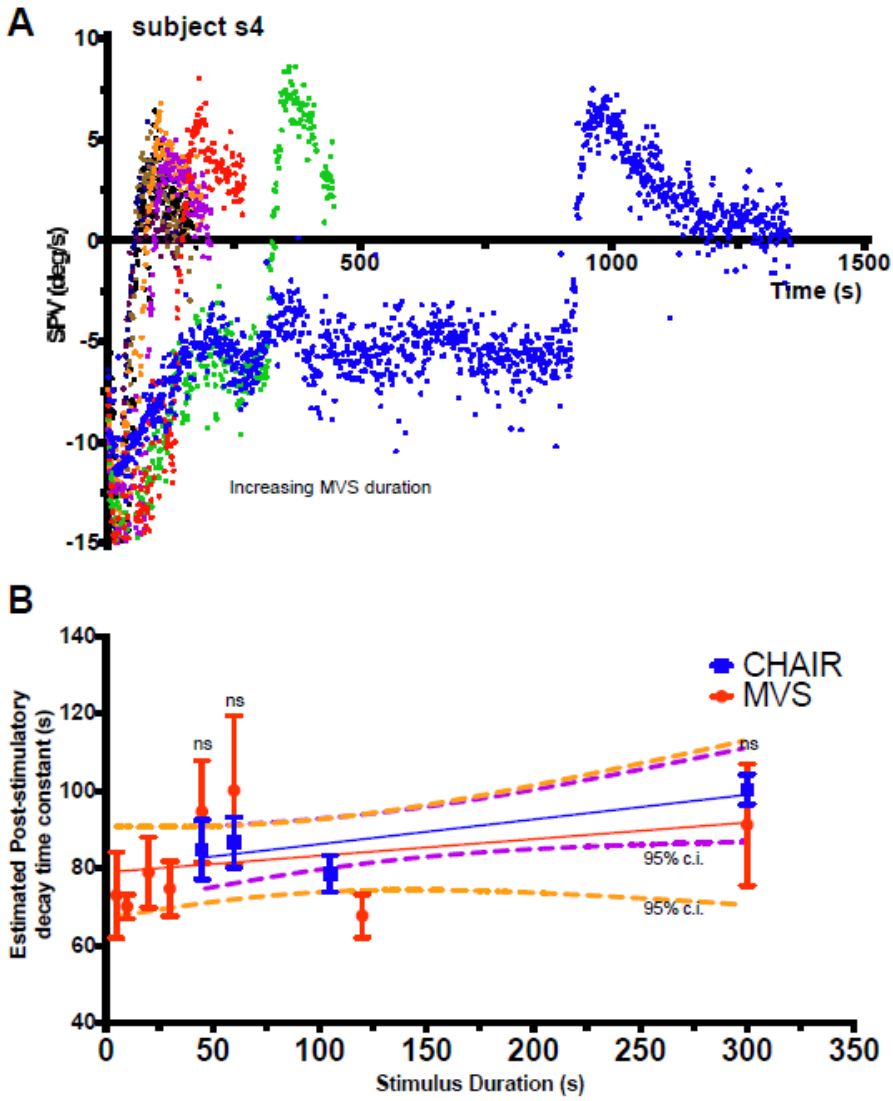

Figure S1 – related to Figure 1
A: Responses of subject s4 to increasing MVS durations of 5s, 10s, 20s, 30s, 45s, 60s, 2mins, 5mins and 15mins. B: Comparison of the estimated post-stimulatory decay time constants, as estimated from the slope of the ln(|SPV|)-time graph from the first 200s of each phase between chair acceleration and MVS at equivalent stimulation durations of 45s, 60s and 300s using a student t-test, n=5 subjects. No significant differences were seen (p=0.31,0.06,0.58). Error bars are standard errors from the mean and dashed lines are 95% confidence intervals (purple for chair, orange for MVS). *Abbreviations*: SPV = slow-phase velocity; CHAIR = Rotatory stimulation; MVS = magnetic vestibular stimulation.

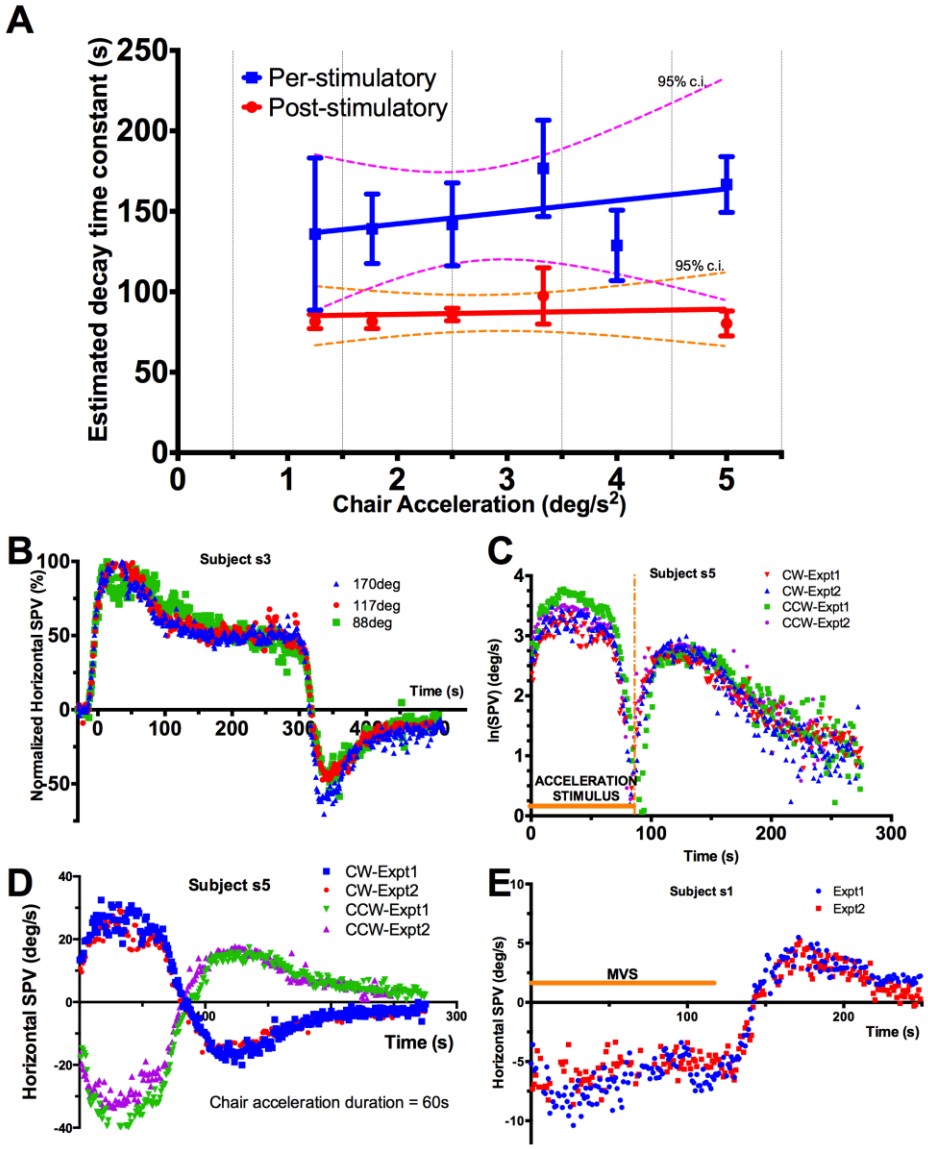

Figure S2 – related to Figure 1

A: Chair acceleration amplitude does not affect estimated decay time constants of per- and post-stimulatory responses, as estimated from the slope of the ln(|SPV|)-time graph during the first 200s of each phase Error bars are standard errors from the mean (five subjects) and dashed lines are 95% confidence intervals (purple for chair, orange for MVS) B: The normalized and absolute value responses of a single subject (s3) at three different head pitches in the MRI bore – two head pitches were above the subject's null point and one was below. Head pitch does not affect the adaptation response. C-D: Response to clockwise and counterclockwise rotatory chair accelerations of fixed amplitude and duration. Experiments in each direction were repeated twice (Expt1 and Expt2): the comparison is represented in (C) as a ln(|SPV|)-time graph to show overlap, and the raw data plotted in (D). Stimulus direction does not affect adaptation response. E: To control for possible habituation across trials, responses of s7 to 120s MVS in two separate trials are shown: experiment 1 (red) at the beginning of the experimental period and experiment 2 (blue) at the end. *Abbreviations*: SPV = slow-phase velocity; CHAIR = Rotatory stimulation.

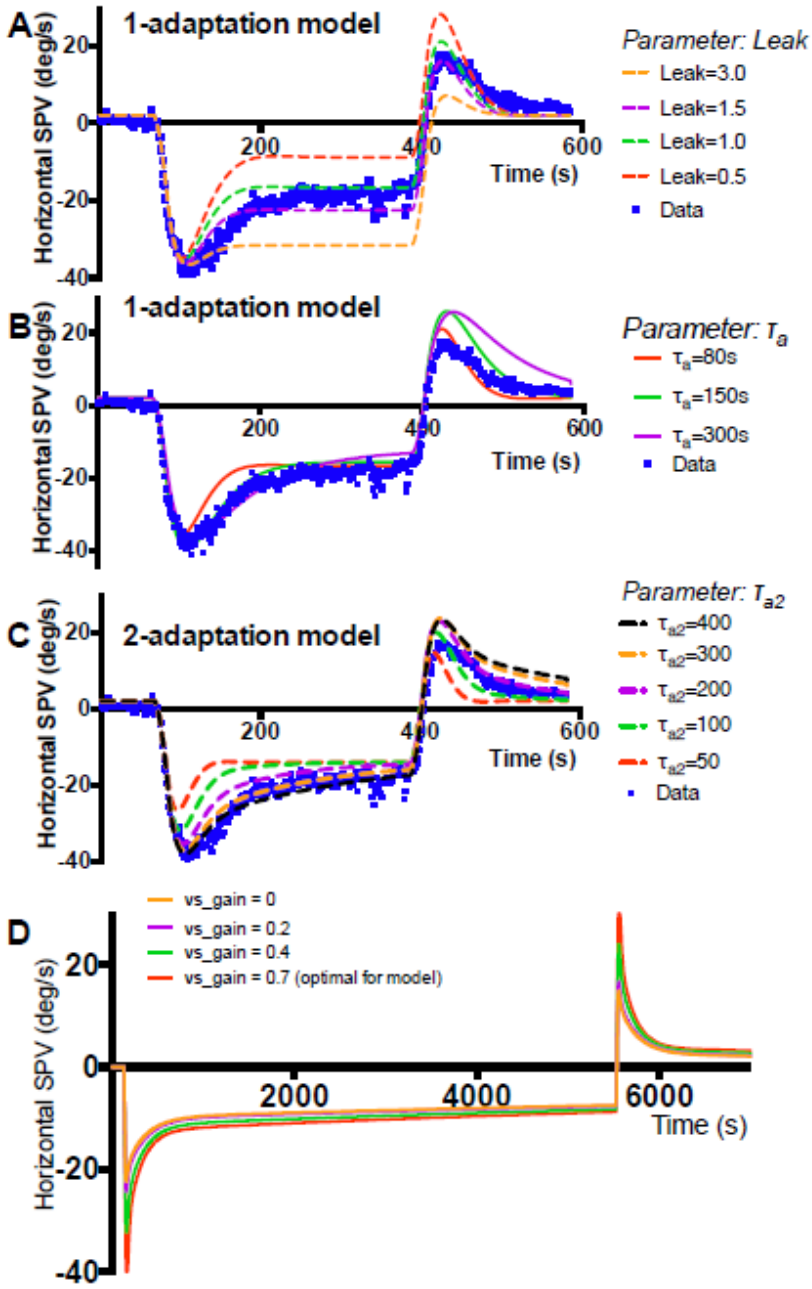

Figure S3 – related to Figure 3 and Figure 4

Changing parameters of a VOR model with a single adaptation operator (1-adaptation model) for a 300s MVS exposure of subject s3. Each parameter is changed in turn while all other parameters are kept constant.
A: Changing the *leak* of an adaptation operator changes the degree of incompleteness of adaptation. The lower the *leak*, the more complete the adaptation. B: Changing the adaptation time constant (T_a) does not simulate the response profile. C. Changing the parameters of a VOR model with two adaptation operators for the same data used in A and B. *C*: Changing the second adaptation time constant (T_{a2}) reveals that T_{a2} around 300s best approximates the data. *D:* Changing the gain of velocity storage (vs_gain) only affects the dynamics of the rise of the response, as it should, but not the adaptive response itself.

(A) Set-point adaptation:
Engineering Control Systems Perspective

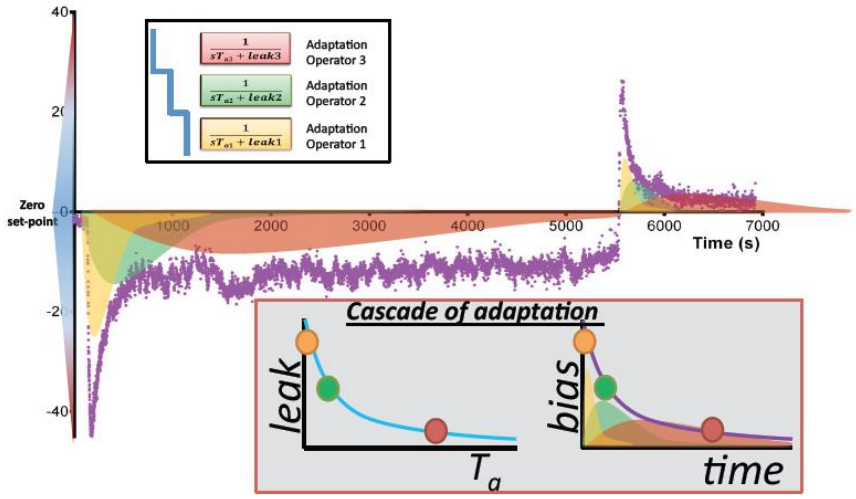

(B) Set-point adaptation:
Bayesian Perspective

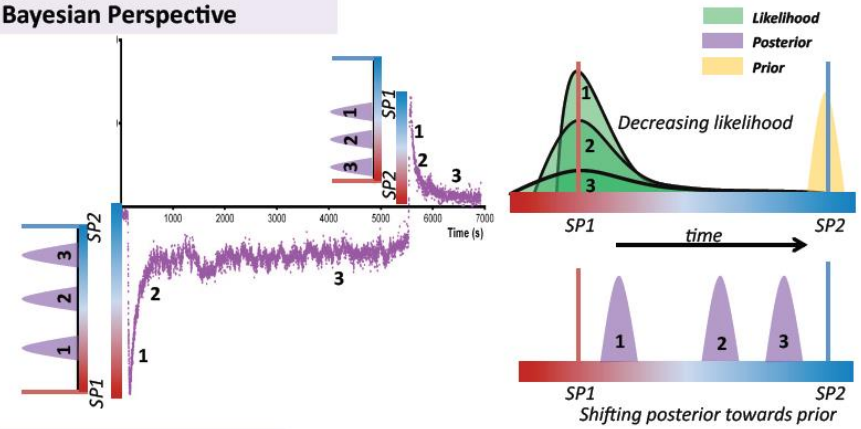

(C) Set-point adaptation:
Skinnerian Perspective

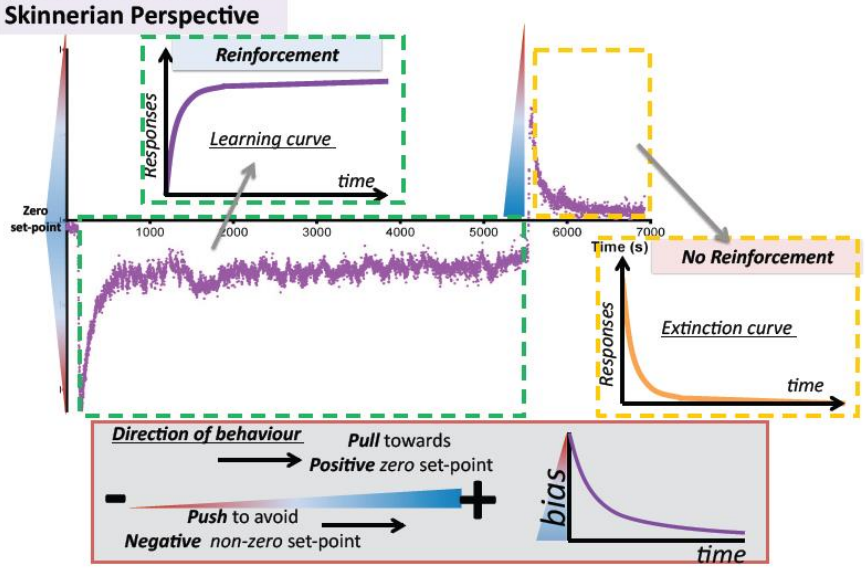

Figure S4 – Set-point adaptation: Single phenomenon, multiple perspectives, related to Figure 3

A: *Engineering control systems perspective:* A cascade of adaptation operators (integrators) with progressively slower dynamic properties shifts the system toward a new zero set-point by eliminating bias. Slowly-changing, lower-frequency stimuli beget slower but more enduring adaptation as reflected in the persistence of the after-effect. This is represented as adaptation operators with slower dynamic properties that are less leaky and eliminate bias over a longer period of time. Shaded areas reflect the relative effects of the different adaptation operators. Purple dots are the slow-phase velocity nystagmus response.
B: *Bayesian Perspective:* During stimulation, the system shifts from one set point (SP1) to another (SP2), and the brain may infer the current set-point using Bayesian inference - the posterior is proportional to the prior times the likelihood. Here we speculate how the concept of multiple time courses of learning with different dynamics can be integrated with Bayesian approach.The yellow curve represents the prior (i.e., SP2). The green curves represent the likelihood as estimated by the adaptation operators. The AUC (area under the cuve) of the likelihood corresponds with certainty of the estimate of the set-point (i.e., SP2). Over time, the likelihood of SP1 becomes less certain and the posterior, represented by the purple curves, shift closer to the prior (SP2) when more information about the environment is collected over longer time periods.
C: *Skinnerian perspective:* A general observation from behavioural reinforcement experiments is that the longer the history of reinforcement (learning curve), the more protracted the extinction curve. What is new here is the integration of concept of multiple time courses of learning with different dynamics. Set-point adaptation can be formulated, in behavioralism terms, as the reinforcement of behaviour that moves a system from an *aversive* non-zero set-point ('bias') to a rewarding zero set-point (with no bias). Therefore, any bias-eliminating behaviour is reinforced. As Skinnerian learning and extinction curves have been demonstrated in many behavioural paradigms in many animal species, we speculate that our finding of multiple time courses of learning can be extended to this approach.

Table S1 – Optimised parameter values, related to Figure 4

MVS 30 min trials

	Ta1	CI-L	CI-U	**La1**	CI-L	CI-U	**Ta2**	CI-L	CI-U	**La2**	CI-L	CI-U	**Ta3**	CI-L	CI-U	**La3**	CI-L	CI-U	**rmse**
subject s1																			
opt1	149.8	141.8	157.7	0.28	0.26	0.30													0.6811
opt2	75.2	67.5	82.9	1.31	1.19	1.43	437.1	499.2	474.0	0.16	0.14	0.19							0.4238
opt3	82	82	82	3.47	3.47	3.47	198.9	198.9	198.9	0.89	0.89	0.89	770.3	770.2	770.3	0.1	0.1	0.1	0.3779
subject s3																			
opt1	225.4	220.6	230.1	0.64	0.66	0.68													2.2911
opt2	81.3	75.8	86.8	3.36	3.24	3.47	445.5	434.9	457.0	0.73	0.72	0.74							1.2726
opt3	46.2	39.2	53.1	5.1	4.9	5.3	306.8	297.9	315.7	1	0.9	1.1	4261	2610.2	5824.2	0.1	0	0.2	1.0495
subject s4																			
opt1	102.6	97.52	107.7	0.7	0.69	0.72													0.784
opt2	80	78	82	3.11	3.01	3.21	303.1	297.5	308.1	1	0.9	1.1							0.673
opt3	55	53	57	4	4	4	183	160	210	0.1	0.1	0.1	2769.2	2769.2	2769.2	0.1	0.1	0.1	0.647
subject s8																			
opt1	215.2	207.6	222.9	0.62	0.6	0.64													1.486
opt2	50	43.1	56.9	2.5	2.33	2.66	351.8	333.5	370.1	0.56	0.54	0.57							1.162
opt3	55.2	55	55	2.75	2.75	2.75	532.6	532.7	532.7	0.91	0.91	0.91	800.1	800	800	1.39	1.39	1.39	1.178

MVS 90 min trials

	Ta1	CI-L	CI-U	**La1**	CI-L	CI-U	**Ta2**	CI-L	CI-U	**La2**	CI-L	CI-U	**Ta3**	CI-L	CI-U	**La3**	CI-L	CI-U	**rmse**
subject s3																			
opt1	270.8	266.2	275.4	0.45	0.44	0.45													2.127
opt2	142.4	133.8	151.0	2.99	2.87	3.11	500	487.1	512.9	0.48	0.48	0.49							1.550
opt3	63	54	72	5.7	5.5	6	308.5	302	315	0.7	0.7	0.7	5449	4798	6181	0.1	0.1	0.1	1.094
subject s8																			
opt1	245.4	238.7	253.0	0.49	0.48	0.5													1.335
opt2	70	60.31	79.7	2.92	2.75	3.08	506.2	485.8	526.6	0.54	0.53	0.55							1.120

Legend: Parameters optimised using least means square curve fitting method (*lsqcurvefit* matlab function, along with *nlparci* matlab function to estimate the confidence intervals of each parameters for the model). Optimised models: *opt1* = 1-adaptation model, *opt2* = 2-adaptation model, *opt3* = 3-adaptation model. All parameters were held constant except parameters that were optimised: *opt1* optimises Ta1 and La1, *opt2* optimises Ta1, La1, Ta2, La2, *opt 3* optimises Ta1, La1, Ta2, La2, Ta3, La3.
rmse = root-mean-square-error. CI-L = Confidence Interval (lower), CI-U = confidence interval (upper).

2. Supplemental Experimental Procedures

MRI procedures.

All experiments were performed in accordance with an approved protocol by The Johns Hopkins University Institutional Research Board. In the MRI subjects lay supine in a Philips Achieva 7T MRI magnetic field (Philips research, Hamburg, Germany) for trials. Eye movements were recorded in darkness using MRI-compatible goggles fitted with infrared video-oculography (VOG, horizontal and vertical) captured at 30 frames/s with 640 x 480 resolution (Resonance Technology, Inc., Los Angeles, CA) while the subjects remained still. No MRI images were taken during the study. Next to the subject's head, a gauss meter (AlphaLab/Trifield GM2, range up to 3T) measured absolute field strength, and a custom-built search coil (75 turns of AWG36 magnet wire on 12mm circular frame) measured change in magnetic field over time (dB/dt) as the table moved into the bore. The magnetic field vector B was directed from the head of the subject toward the feet when entering the magnet supine and head-first. A string potentiometer sensor cable (UniMeasure, VP series) was attached to the scanner table to monitor its position and velocity as it moved into and out of the bore. Data were collected using a single custom-written program (using Microsoft Visual C++) which synchronised the eye movements and analog sensor data for later analysis using Matlab (MathWorks, Natick, MA).
The subject was first placed supine on the MRI table with their head near the bore in a neutral neck position as if for a routine MRI head scan. The pitch angle was measured using a non-metallic protractor with a bubble level. The external reference was taken as the line from the lateral canthus of the eye to the tragus, approximating Reid's horizontal plane such that the lateral semicircular canals are angled about 20 degrees upward from this line. Before each subsequent entry into the magnetic field bore, the angle of the head pitch of the subject was measured and in some cases altered to change the orientation of the labyrinth with respect to the magnetic field and so influence the amplitude of induced nystagmus. Eye movements were recorded from the right eye of the subject using infrared VOG, which remained fixed on the head throughout collection of all data files. The eye movements of the subjects were calibrated outside the magnetic field while supine with the head neutral on the table and looking directly at a target screen above. The VOG goggles calibration was repeated whenever the goggles were removed or repositioned. After calibration, the room lights were turned off and vision was prevented by covering the head of the subject with a double layer of black felt cloth. The baseline eye movement recordings were then taken outside the bore. The field strength outside the bore near the ears of the subject was approximately 0.7 tesla.

MRI experimental paradigms

In each trial, after the baseline data was recorded, the subject was moved into the MRI bore, head first, using the fixed-speed table motor drive (10.8 cm/s over 2m travel). Once at the centre of the bore, they remained there for fixed durations between 5s and 90mins. After the fixed duration inside the bore, the subject was moved out of the bore to the original starting position and remained still for another 300s or longer while the reversal nystagmus was recorded or until the eye movements appeared similar to baseline. Before each trial began, just outside the bore, we made sure that any residual nystagmus had returned to the baseline level. All subjects undertook a 300s duration trial inside the bore for standardized comparison with rotatory chair data. The head of the subject was positioned at the same pitch angle in all trials unless indicated otherwise. In trials in which the head pitch was varied, the head was positioned at a different pitch angle by placing pads under the neck and shoulders of the subject or at the back of the head. Subjects were instructed to stare in a straight ahead direction of gaze throughout the trial and this was monitored online by the experimenter.

Chair setup protocol

Subjects sat upright on a motorized rotatory chair within a cubic coil frame (1.02m on a side) that generated three orthogonal magnetic fields of varying frequencies. A custom made chin and forehead stabilizer kept the head of the subject fixed to the chair to reduce any relative movement between the head and the chair. Horizontal movements of one eye were recorded using the magnetic-field search-coil technique. Eye drops were provided, as needed, for comfort. A second coil was taped to the subject's forehead to record head movements. A third coil was fixed to the chair to record its rotation. Raw coil signals were filtered in hardware (90Hz low-pass bandwidth, digitized (1kHz) and saved on computer for later analysis. We performed calibration at the beginning of each session. The subjects sat facing a row of nine red LEDs located 124cm away in an otherwise darkened room. The subjects were required to make 13 sequences of nine fixations (each lasting 1s). During each visual target sequence, the chair remained

fixed. Between two target sequences the chair moved to a new position according to its own sequence (chair position 0,4,8,12,8,4,0,-5,-8,-12,-8,-4 and 0 degrees).

Chair acceleration step paradigm

The eye movements of the subject were recorded for 30s looking for any baseline nystagmus. In order to mimic the 15s acceleration ramp as the subject entered the MRI bore (from 0.7T to 7T magnetic field strength), we applied a 15s jerk (derivative of acceleration) at the onset during which the acceleration magnitude was increased from 0 to its final value. The chair was rotated at constant acceleration for a fixed duration between 45s and 300s. Then a second 15s jerk was applied as the acceleration reduced to 0 to mimic exiting the MRI bore. When the chair stopped accelerating, the subject was maintained at final maximum chair velocity for another 240-300s until the nystagmus returned to baseline. The subject was then rotated in the opposite direction at a constant acceleration of a different fixed duration, to bring the subject to the final minimum chair velocity. Again, when the chair stopped accelerating, the subject was maintained at the final minimum chair velocity until the nystagmus returned to baseline. Finally, the subject was rotated in the opposite direction at a constant acceleration until the final velocity returned to zero. We recorded the nystagmus until it reached baseline. In summary, each chair-acceleration trial is equivalent to three consecutive MRI trials of different durations. The acceleration duration combinations we used were 45s-105s-45, 45-385s-45s, and 60s-300s-60s. Some subjects undertook the trials in both directions. All subjects undertook the 60s-300s-60s paradigm to allow for comparison with MRI trials.

Data analysis

For recording eye movements in the magnetic field a custom Matlab pupil tracker program selects pixels in the saved eye camera video frames based on a brightness threshold, then returns the centroid of the resulting black pupil area. A behavioral calibration while looking at known targets is used to compute eye angles in degrees. Custom-written MatLab programs analysed the horizontal and vertical components of eye movement and dB/dt data. Nystagmus was manually marked near the beginning and end of each slow-phase, and a least-squares line was automatically fit to the data between the marked points. The slope of each slow-phase line became a single slow-phase velocity data point and outlier data points were removed.The torsional and vertical components of nystagmus were usually quite small and were not analyzed. Saccades, quick phases and blinks were ignored.

Simulations and model fits

The VOR model was created using the Simulink module of MatLabA custom-written Matlab program allowed visual comparison of the model output and experimental data plot on the same graph, and permitted parameter optimization using Levenberg-Marquardt optimization technique (Matlab function Isqnonlin). We optimize parameters quantitatively using a least means square curve fitting method (lsqcurvefit matlab function, along with nlparci function to estimate the confidence intervals of each parameters for models with 1, 2 or 3 different time courses). We then calculate the goodness-of-fit between the optimised models and the data, i.e. the root mean square error (RMSE). The model and experimental results from Matlab were transferred to GraphPad Prism for statistical analysis and visual display.