



Published in final edited form as:

Cerebellum. 2021 February ; 20(1): 62–73. doi:10.1007/s12311-020-01183-x.

Reinforcement signaling can be used to reduce elements of cerebellar reaching ataxia.

Amanda S. Therrien¹, Matthew A. Statton², Amy J. Bastian^{2,3}

¹Moss Rehabilitation Research Institute, Elkins Park, PA

²Kennedy Krieger Institute, Baltimore, MD

³Department of Neuroscience, Johns Hopkins School of Medicine, Baltimore, MD.

Abstract

Damage to the cerebellum causes a disabling movement disorder called ataxia, which is characterized by poorly coordinated movement. Arm ataxia causes dysmetria (over- or undershooting of targets) with many corrective movements. As a result, people with cerebellar damage exhibit reaching movements with highly irregular and prolonged movement paths. Cerebellar patients are also impaired in error-based motor learning, which may impede rehabilitation interventions. However, we have recently shown that cerebellar patients can learn a simple reaching task using a binary reinforcement paradigm in which feedback is based on participants' mean performance. Here we present a pilot study that examined whether patients with cerebellar damage can use this reinforcement training to learn a more complex motor task - to decrease the path length of their reaches. We compared binary reinforcement training to a control condition of massed practice without reinforcement feedback. In both conditions, participants made target-directed reaches in 3-dimensional space while vision of their movement was occluded. In the reinforcement training condition, reaches with a path length below participants' mean were reinforced with an auditory stimulus at reach endpoint. We found that patients were able to use reinforcement signaling to significantly reduce their reach paths. Massed practice produced no systematic change in patients' reach performance. Overall, our results suggest that binary reinforcement training can improve reaching movements in patients with cerebellar damage and the benefit cannot be attributed solely to repetition or reduced visual control.

Keywords

Cerebellum; Ataxia; Reinforcement; Motor learning; Reaching

Terms of use and reuse: academic research for non-commercial purposes, see here for full terms. <http://www.springer.com/gb/open-access/authors-rights/aam-terms-v1>

Please address correspondence to: Amanda S. Therrien, PhD, Moss Rehabilitation Research Institute, 50 Township Line Rd., Elkins Park, PA, 19027, therriea@einstein.edu.

Publisher's Disclaimer: This Author Accepted Manuscript is a PDF file of a an unedited peer-reviewed manuscript that has been accepted for publication but has not been copyedited or corrected. The official version of record that is published in the journal is kept up to date and so may therefore differ from this version.

Conflict of Interest Statement

The authors declare that they have no conflict of interest.

Introduction

Damage to the cerebellum impairs motor coordination, causing a disabling movement disorder called ataxia. During reaching, ataxia is characterized by dysmetria (over- and undershooting of target locations) and multiple feedback driven corrections causing irregular and prolonged hand trajectory paths [1, 2]. As a result, people with cerebellar damage can experience significant difficulty with many activities of daily life such as reaching for a glass of water, eating and dressing.

Reaching ataxia is thought to stem from impaired sensorimotor prediction [3–5], which involves computing the expected sensory outcome of motor commands based on an internal model of limb dynamics. Prediction signals keep movement well calibrated in the face of time-delayed peripheral feedback and adaptive motor learning (or adaptation) is thought to keep sensory predictions up to date. Adaptation relies on access to the direction and magnitude of movement errors (i.e. vector errors). Unexpected vector errors indicate that the predicted consequences of a movement did not match the actual consequences - there was an error in sensory prediction. Sensory prediction errors drive adaptation mechanisms to update movement commands as well as internal predictions on a trial-by-trial basis [6, 7]. Many studies have shown that cerebellar damage impairs error-driven adaptation in response to visuomotor and force-field perturbations during reaching [8–14]. Importantly, cerebellar patients do not show the characteristic motor after-effects that are hallmarks of adaptation in these paradigms [10,12,14,15].

Unlike adaptation, reinforcement learning relies on scalar measures of action outcome such as success or failure (for reviews see [16,17]). This mechanism requires exploring different task solutions to determine which actions lead to successful outcomes. Over time, the learner can estimate the probability that each solution will result in a rewarding outcome and shift action choice toward those solutions with the highest probability of reward. When learning a simple reaching task, removing visual feedback of vector error and providing only binary feedback of movement outcome may bias the motor system to favor reinforcement mechanisms over cerebellum-dependent adaptation [18].

We previously studied how a group of cerebellar patients learned to alter their reach direction in a simple 2-dimensional task with either vector error or binary reinforcement feedback [14]. Cerebellar patients showed no learning with visual feedback of vector errors - marked by an absence of motor after-effects. Conversely, the same patients could learn to varying degrees with reinforcement, and fully retained what was learned. This result suggested that binary reinforcement paradigms may represent a favorable approach to motor training in patients with cerebellar damage. However, it remained unclear whether binary reinforcement training would benefit cerebellar patients in more complex learning tasks and whether it could be used to improve movement characteristics directly related to their ataxia.

Here, we present the results of a pilot study that examined whether cerebellar patients could use binary reinforcement signaling to improve a feature of their reaching ataxia - the prolonged and irregular trajectory path travelled by the hand - in an unconstrained 3-dimensional reaching task. We show that cerebellar patients were able to learn to reduce

their path length and straighten their reaches to a single target with reinforcement feedback. A control intervention of massed practice without reinforcement yielded no change in patients' reach paths. Our results provide evidence that binary reinforcement training can reduce a feature of reaching ataxia and the mechanism of benefit is not attributable to simple repetition or reduced visual feedback alone.

Method

Participants

The study was approved by the Johns Hopkins Institutional Review Board and all subjects gave written informed consent prior to participation. 10 patients with ataxia due to cerebellar damage (6 males, 4 females, mean age \pm SD: 59.4 ± 15.0 years) were recruited. A sample of 10 neurologically healthy control participants, matched for mean age and handedness (3 males, 7 females, mean age \pm SD: 54.0 ± 11.5 years) were also recruited for baseline comparison of reach paths. Further details about the participants' characteristics are shown in Table 1. The severity of patients' ataxia was characterized using the International Cooperative Ataxia Rating Scale (ICARS) [19].

Apparatus

Figure 1a shows the experimental set-up. Participants sat in a chair and made reaching movements with their dominant arm in 3-dimensional space while wearing a virtual reality headset (Oculus Rift). Movement kinematics were collected at 100 Hz using Optotrak Certus motion capture hardware (Northern Digital, Waterloo, Ontario, Canada). Infrared-emitting active markers were placed on the shoulder (acromion process), elbow (lateral humeral epicondyle), wrist (ulnar styloid process), and index fingertip (distal phalange) of the dominant arm. A tensor bandage was used to wrap the thumb, middle, third and fourth fingers under the palm to help participants maintain a pointing posture with their dominant hand. Movement kinematics were integrated with the virtual reality environment using custom-written software (Vizard, World Viz Inc., Santa Barbara, California, USA).

Procedure

Target position specification—Participants performed reaching movements from a home position to a single target. The center of the home position was defined as the position of the fingertip when the elbow angle was at 135-degrees of flexion, and the shoulder angle was at 0-degrees of flexion and 75-degrees of abduction. A physical armrest was placed under the forearm in the home position to give participants an additional tactile cue as well as allow them to rest their arm between trials. The center for the target was defined as the position of the first metacarpal-phalangeal joint (fingertip marker position minus the finger length) when the elbow angle was at 0-degrees of flexion and shoulder angle was at 90-degrees of flexion and 0-degrees of abduction. Spherical radii of 25 mm and 38 mm were then applied to the home and target positions, respectively, to define each as a sphere in the 3-dimensional virtual environment. The addition of these radii produced an average minimal straight-line reach distance of 235.30 ± 60.27 mm (mean \pm SD) across patients and 312.77 ± 33.95 mm (mean \pm SD) across controls.

Experimental tasks for cerebellar patients—Cerebellar patients completed two reaching tasks, one reinforcement training task and one massed practice control task. The order of task completion was counterbalanced across participants. Both tasks began with an 80-trial baseline phase, that was divided into 40 trials where participants were shown the visual cursor representing their index fingertip position and 40 trials where no fingertip cursor was provided. Following the baseline phase, participants completed a 160-trial intervention phase where either the reinforcement training task or massed practice control task were performed. The visual cursor was not provided throughout this phase. Following the intervention phase, participants completed a 120-trial retention phase that was divided into 60 trials where the visual cursor remained invisible and 60 trials where the visual cursor was shown (Figure 1b).

A trial began with the participant's finger in the home position. Once the finger was held in the home position for 500 ms, the target sphere appeared (Figure 1a). Participants were instructed to reach straight to the target as quickly as possible and then, once the index finger entered the target sphere, hold their finger in that position until the target disappeared (500 ms). The target disappearing signaled the end of the current trial, at which time participants could return their hand to the home position to begin the next trial.

The visual cursor representing the index finger position was always shown before the start of each trial to help participants navigate to the home position. If the trial was one where visual feedback was removed (Figure 1b), the visual cursor became invisible at the same time that the target appeared, only reappearing when the index finger entered the target sphere.

Participants were encouraged to move as quickly as possible during their reaches, though no specific movement speed requirements were put in place. Instead, participants were given a maximum reach time of 5 s. If reaches exceeded 5 seconds (i.e. the participant had not reached the target sphere within that time), the trial was stopped and repeated. Prior to beginning the experiment, participants performed a 20-trial practice block with the visual cursor present to familiarize themselves with the task and virtual environment.

In the intervention phase of the reinforcement training task, participants were required to learn to reduce the path length of their reach trajectories. To implement this training, we provided binary reinforcement feedback about participants' path length on each trial following a closed-loop reinforcement schedule [14]. That is, we calculated a moving average of the path lengths of participants' previous 10 reaches. Reaches with a path length less than or equal to the moving average (i.e. the current reach was the same or better than the mean of the previous 10 reaches) were rewarded at the end of the trial with a pleasant tone and a point added to a tracker in the visual display. Reaches with a path length greater than the moving average (i.e. the current reach was worse than the mean of the previous 10 reaches) resulted in an unpleasant tone and no points added. Using this type of closed-loop reinforcement schedule, participants are continuously driven to straighten their reaches.

Participants were not informed that the reinforcement signaling was related to the path length of their reaches. They were simply told they would be rewarded for reaching in a particular way and it was up to them to find the successful reaching pattern. Binary

reinforcement was delivered at the end of each trial to inform participants if they had performed correctly or incorrectly. Participants' goal was to obtain as many points as possible in the intervention phase. Both the pleasant and unpleasant tones were presented to participants prior to the task to familiarize them with the feedback signaling.

In contrast to the reinforcement training task, the intervention phase of the massed practice control task consisted of simple repetition. Participants were instructed to continue reaching straight to the target, as quickly as possible. They received no feedback about the path length of their reaches.

Experimental task for control participants—Control participants were recruited for a single baseline reaching task. The reach path lengths exhibited by control participants in this task served as a standard against which to compare cerebellar patients' baseline reach path lengths to demonstrate their reaching ataxia. Control participants did not perform either of the training interventions tested in the cerebellar patient group because those tasks were designed to promote the reduction of irregular and prolonged movement paths that characterize reaching ataxia. Control participants do not show irregular, highly curved or prolonged movement paths when performing target directed reaches (e.g. [2, 3]).

The baseline task comprised a total of 200 trials, divided into 100 trials where participants were shown a visual cursor representing their index finger movement throughout the reach and 100 trials where the cursor was invisible. Individual trials followed the same procedure as outlined in the description of the patient tasks. The order in which the 2 visual feedback conditions were performed was counterbalanced across participants.

Analysis

The path length of each reach trajectory was converted to a percentage of the minimal straight-line distance for each subject. Shapiro-Wilk tests revealed that the patient data obtained was not normally distributed, so nonparametric statistical analysis was performed on group means of subjects' median path length, path length interquartile range (IQR, the range of values that lie between the 25th and 75th percentiles) and median peak velocity for each subject.

For control participants, subject medians were calculated for the final 40 trials of each visual feedback condition. For cerebellar patients, subject medians were calculated for the following experiment phases: baseline with visual feedback (BL-V, trials 1:40), baseline without visual feedback (BL-NV, trials 41:80), end of intervention (trials 201–240), retention without visual feedback (Ret-NV, trials 241–280), and retention with visual feedback (Ret-V, trials 301–340). To quantify individual patient performance in each task, we computed the total learning as the change in median path length, in mm, between the baseline phase without visual feedback and the end of the intervention.

To demonstrate patients' reaching ataxia, the median path length in the baseline with visual feedback and baseline without visual feedback phases was compared to control participants' medians in the corresponding visual feedback conditions using Mann-Whitney *U* tests. Within the patient tasks, statistical analysis comprised a series of planned comparisons,

tested with matched pairs Wilcoxon signed rank tests. Comparisons were made between the reinforcement training and massed practice control tasks over the two baseline phases. Within each task, comparisons were made between the baseline with visual feedback and retention with visual feedback phases, between the baseline without visual feedback and end of intervention phases, and between the baseline without visual feedback and retention without visual feedback phases. All 10 participants recruited for this experiment completed the reinforcement training task. However, due to scheduling constraints, only 8 participants completed the massed practice control task. As a result, comparisons made across tasks were performed using data only from the 8 participants who completed both tasks. For all statistical analysis the alpha value was set at 0.05.

Results

Figure 2 shows that the cerebellar patients tested here exhibited increased path lengths compared to a group of age-matched neurologically healthy control participants. Statistical results are shown in Table 2. When provided with visual feedback of their hand (Figure 2a) cerebellar patients showed significantly greater path lengths than controls in both the reinforcement training and massed practice control tasks. Cerebellar patients also showed significantly greater baseline path lengths compared to controls when visual feedback of the hand was removed (Figure 2b).

Cerebellar patients completed two reaching tasks: a reinforcement training task and a massed practice control task. Figure 3 shows the time series of reach path lengths for the cerebellar group in each task. Note that the patients exhibited similar path lengths across the two tasks in the baseline phases. In the reinforcement training task (Figure 3a), patients showed highly variable reach paths early in the intervention phase, but were able to reduce their path length by the end of the intervention. The reduction in path length was also maintained in the first retention phase where visual feedback of hand position remained invisible. Path lengths increased when visual feedback of the hand position was restored in the second retention phase, though still remained below baseline. In contrast, when reaching within the massed practice control task (Figure 3b), patients showed no change in path length from baseline to the end of the intervention. Accordingly, there was also no change in path length in the two retention phases of the task.

The total learning exhibited by each patient in the two experimental tasks is shown in Figure 3c–d. In the reinforcement training task, 8 of the 10 patients tested learned to reduce their path length from the baseline phase without visual feedback to the end of the intervention (Figure 3c). However, only 3 of the 8 patients who completed the massed practice control task showed reduced path lengths from baseline (Figure 3d).

Planned comparisons of median path length between the two tasks showed no significant differences at baseline, with or without visual feedback of hand position (Figure 4a, Table 3). In the reinforcement training task, patients showed a significant reduction in median path length from the baseline phase without visual feedback to the end of the intervention. Median path length in the retention phase without visual feedback was also significantly lower than at baseline. There was no significant difference in path length between the

baseline and retention phases where visual feedback was provided. There were also no significant differences found among any of the planned comparisons performed across phases of the massed practice control task. All statistical comparisons across phases of the reinforcement and massed practice control tasks are shown in Table 4.

To analyze path length variability, we computed the interquartile range (IQR) for each subject and experiment phase. Planned comparisons across tasks showed no difference in variability in either baseline phase (Figure 4b, Table 3). Within the reinforcement training task, there was a significant reduction in IQR from the baseline phase without visual feedback to the end of the intervention. There was also a reduction in IQR from the baseline phase without visual feedback to the retention phase without visual feedback. There was no significant difference in IQR between the baseline with visual feedback and retention with visual feedback phases. Comparisons of IQR across phases of the massed practice control task yielded no significant results (Table 4)

Finally, comparisons of median peak velocity across tasks showed no significant differences at baseline (Figure 4c, Table 3). Within the reinforcement training task, there was a significant difference in velocity between the baseline with visual feedback and retention with visual feedback phases. However, no differences were found between the baseline phase without visual feedback, the end of the intervention or retention phase without visual feedback. No significant differences were found within the massed practice control task (Table 4).

There are multiple ways by which cerebellar patients could have reduced their path length in the reinforcement training task. To visualize whether there was a systematic pattern in how path lengths were altered, we plotted all reach trajectories in the reinforcement training and massed practice control tasks across the experiment phases where visual feedback was removed (Figures 5 and 6). 3-dimensional reach trajectories were normalized by trajectory length and plotted as probability distributions in each of the 3 planes. We performed 3-dimensional axis rotations to align the start and end points of each trajectory (i.e. the straight-line reach path) along the x-axis. The starting point of the reach was aligned to the origin and the target was located along the x-axis. We then created probability distributions to examine the density of reach trajectories across the group in 3 planes. An increase in density (hot colors) indicated that more reaches crossed through a given point in cartesian space.

In the reinforcement training task, viewing reach trajectories in the Y-Z plane (i.e. looking through the reach, Figure 5a) shows that at baseline reaches tended to deviate from a straight line laterally in the rightward direction. Viewing the X-Y plane (i.e. a birdseye view, Figure 5b) shows these rightward deviations occurred through the middle portion of the reach - marked by a spreading of the probability distribution (cooler colors mid-reach). Examining the X-Z plane (i.e. a sideline view of the reach, Figure 5c) shows that at baseline, vertical path deviations were spread equally above and below the straight line. After the reinforcement intervention, rightward lateral deviations in the X-Y plane decreased in the middle portions of the reach path (marked by a tightening of the probability distribution, i.e. hotter colors, Figure 5b middle). Vertical deviations from straight in the early portion of the

reach also decreased after the reinforcement intervention (Figure 5c, middle). Both of these changes were maintained in the retention phase.

At baseline, reach trajectories in the massed practice control task showed a similar pattern to that observed in the reinforcement task (Figure 6). However, from baseline to the end of the intervention phase, trajectories continued to show lateral deviations from straight in the rightward direction (Figure 6b). Deviations in the vertical plane (Figure 6c) also continued from baseline to the end of the intervention, even appearing to worsen slightly - marked by an increase in density below a straight line.

Overall, examination of reach trajectories indicates that the reinforcement training task induced a systematic straightening of reaching movements across patients in the early and middle portions of the reach. Conversely, the massed practice control intervention produced no systematic change in patients' reach trajectories, which is consistent with no change in reach path length.

Discussion

In this pilot study, we examined whether patients with cerebellar damage could use binary reinforcement signaling to learn to improve a feature of their reaching ataxia. We used a closed-loop reinforcement schedule, where reward is contingent upon prior performance, to reinforce reaching movements with shorter path lengths. Reinforcement training produced a significant reduction in patients' path length from baseline that was retained after removing reinforcement feedback. A control intervention of massed practice without reinforcement produced no systematic change in patients' reach paths. Our results suggest that binary reinforcement training may represent a reasonable approach to motor training for cerebellar patients.

Our study asked the important question of whether binary reinforcement signaling could be used to train a movement parameter that is directly related to cerebellar reaching ataxia. Due to the combination of dysmetria and subsequent corrective movements, reach trajectories in cerebellar patients can show irregularly curved, variable and prolonged movement paths compared to neurologically healthy controls [2–4, 20–21]. In line with this, our patient sample showed reach path lengths at baseline that were significantly longer than a group of age-matched control participants. Motor rehabilitation interventions in ataxia typically yield variable outcomes, which has been suggested to stem, in part, from cerebellar damage impairing adaptive motor learning [22]. That patients learned to reduce their prolonged path length in our task suggests that binary reinforcement may hold promise as a training intervention for cerebellar ataxia.

Although patients learned to reduce their path length in the reinforcement training task, their performance appeared to worsen early in the intervention phase. That is, reach paths became highly variable – showing a sharp increase followed by a sharp decrease – at the onset of the reinforcement intervention before stabilizing into a more gradual learning curve. Such increases in variability early in learning may reflect exploration behavior as patients attempt various possible task solutions to determine those that yield reward [23]. In the

reinforcement training task subjects were instructed that they would be rewarded for reaching in a particular way, but they would need to determine the correct pattern themselves. They were also cued to the beginning of the intervention phase of the task. Given these instructions, it is reasonable that patients may have engaged in some strategic exploration early in learning.

We compared binary reinforcement training to a control intervention of massed practice in which patients made repeated reaches to the end target. In this condition, they were instructed to try to reach straight to the target, but they received no reinforcement feedback about their path length. The massed practice intervention produced no systematic change in patients' reach paths. This result indicates that learning observed in the binary reinforcement task was not simply the result of repetition (or practice) of the movement. However, the massed practice intervention used in our study was not identical to protocols designed to leverage use-dependent plasticity. Therefore, it is important to note that our results do not exclude the possibility that repetition of a more constrained movement might yield some benefit for cerebellar patients.

Patients learned to reduce their path length in the binary reinforcement task, yet learning did not robustly transfer to a second retention phase where visual feedback of hand position was restored. Reinstating visual feedback of hand position may have shifted the sensorimotor system towards problematic vector error based or visual control mechanisms that were not readily available when visual feedback was removed, but may dominate when it is present [24, 25]. Alternatively, simple decay processes may have driven the effect - we did not counterbalance the order of retention phases across participants with visual retention always occurring last in the experiment. Nevertheless, restoring visual feedback did not result in path length immediately returning to baseline levels suggesting that some portion of learned reaching movement remained accessible.

Although our present and previous work has shown that binary reinforcement signaling can improve motor learning after cerebellar damage, the precise nature of this benefit remains unclear. One possibility is that binary reinforcement decreases central reliance on cerebellum-dependent learning mechanisms driven by sensory prediction errors. Izawa and Shadmehr [18] posited a mechanistic model in which the outputs of an adaptive and a reinforcement-based learning system were combined to determine motor output. In their model, reducing vector error feedback reduced the reliability of sensory prediction error signals, which biased the learning system toward the output of reinforcement-based mechanisms. In support of their model, they found that learning to counter a visuomotor rotation with binary reinforcement occurred without the sensory recalibration normally seen with adaptive, sensory prediction error driven learning.

An alternative explanation lies in the fact that our binary reinforcement task specifically removed visual feedback of vector error. By removing visual feedback, our task also eliminated patients' capacity to guide their movement using visual control mechanisms. Cerebellar patients may rely more heavily on visual feedback to control their movement, possibly to compensate for impaired sensory prediction [21, 26, 27]. Indeed, a reliance on time-delayed feedback to make movement corrections has been posited as a substantial

contributor to the oscillatory movement patterns that characterize ataxia [14, 27, 28]. Therefore, removing visual feedback in our task may have helped patients decrease their path length simply by reducing feedback-driven movement oscillations. However, visual feedback of hand position was also removed during the massed practice control intervention. That no change in path length was observed in the control task argues against a benefit driven solely by reduced visual control.

Finally, a third possibility is that cerebellar patients employed an explicit strategy to learn to alter their reaches. The use of explicit strategies can contribute to the behavior observed in motor learning tasks [29–33] and motor learning with binary reinforcement has been found to rely heavily on the use of explicit strategies in healthy young individuals [34, 35]. Indeed, we did find evidence supporting explicit exploration at the beginning of the reinforcement intervention. However, recent work has found that patients with cerebellar damage are impaired in adopting explicit strategies to compensate for their learning deficits [25, 36]. Impaired explicit learning could result from cerebellar damage specifically disrupting the computation of a strategy (e.g. [37]) or more general age-related declines in spatial and working memory [25]. Regardless, we feel this literature suggests that adoption of an explicit strategy was unlikely to be a main driver of our results; although, we cannot rule it out as a potential contributor.

In summary, we have shown that patients with cerebellar damage can use binary reinforcement signaling to learn to reduce the prolonged path lengths that characterize reaching ataxia and the benefit of this training cannot be attributed solely to simple repetition or reduced visual control. While the results of this preliminary study support further investigations into the efficacy of reinforcement-based training protocols for cerebellar ataxia, there are some limitations of this work. First, our study tested only a small sample of cerebellar patients and the effects observed were subtle. Additionally, we examined only a single learning session with a short retention phase, within which we saw possible evidence of decay. Further work is needed to determine whether the effects observed here are robust to testing a larger patient sample and whether training protocols can be developed to encourage retention. Future studies are also needed to address outstanding questions regarding the precise mechanism underlying the benefit that binary reinforcement may provide for cerebellar patients and the generalizability of learning in these tasks.

Funding

This work was supported by National Institute of Child Health and Human Development HD040289 to AJB and a Johns Hopkins Distinguished Science of Learning Fellowship to AST.

References

1. Holmes G The symptoms of acute cerebellar injuries due to gunshot injuries. *Brain*. 1917;40:461–535.
2. Bastian A, Thach W. Cerebellar outflow lesions: A comparison of movement deficits resulting from lesions at the levels of the cerebellum and thalamus. *Annal Neurol*. 1995;38: 881–892. [PubMed: 8526460]
3. Bastian A, Martin T, Keating J, Thach W. Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J Neurophysiol*. 1996;76:492–509. [PubMed: 8836239]

4. Bhanpuri N, Okamura A, Bastian A. Predicting and correcting ataxia using a model of cerebellar function. *Brain*. 2014;137:1931–1944. [PubMed: 24812203]
5. Therrien A, Bastian A. The cerebellum as a movement sensor. *Neurosci Lett*. 2018;688:37–40. [PubMed: 29966751]
6. Krakauer J, Pine Z, Ghilardi M, Ghez C. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci*. 2000;20:8916–8924. [PubMed: 11102502]
7. Shadmehr R, Smith M, Krakauer J. Error correction, sensory prediction, and adaptation in motor control. *Neuroscience*. 2010;33:89–108.
8. Weiner MJ, Hallett M, Funkenstein HH. Adaptation to lateral displacement of vision in patients with lesions of the central nervous system. *Neurology*. 1983;33:766–72. [PubMed: 6682520]
9. Martin T, Keating J, Goodkin H, Bastian A, Thach W. Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation. *Brain*. 1996;119:1183–98. [PubMed: 8813282]
10. Maschke M, Gomez C, Ebner T, Konczak J. Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. *J Neurophysiol*. 2004;91:230–238. [PubMed: 13679403]
11. Smith M, Shadmehr R. Intact ability to learn internal models of arm dynamics in Huntington's Disease but not cerebellar degeneration. *J Neurophysiol*. 2005;93:2809–2821. [PubMed: 15625094]
12. Tseng Y, Diedrichsen J, Krakauer J, Shadmehr R, Bastian A. Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol*. 2007;98:54–62. [PubMed: 17507504]
13. Schlerf J, Xu J, Klempfuss N, Griffiths T, Ivry R. Individuals with cerebellar degeneration show similar adaptation deficits with large and small visuomotor errors. *J Neurophysiol*. 2013;109:1164–1173. [PubMed: 23197450]
14. Therrien A, Wolpert D, Bastian A. (2016). Effective reinforcement learning following cerebellar damage requires a balance between exploration and motor noise. *Brain*. 2016;139:101–114. [PubMed: 26626368]
15. Izawa J, Criscimagna-Hemminger S, Shadmehr R. Cerebellar contributions to reach adaptation and learning sensory consequences of action. *J Neurosci*. 2012;32:4230–4239. [PubMed: 22442085]
16. Sutton RS, Barto G. An introduction to reinforcement learning. Cambridge, MA: MIT Press; 1998.
17. Lee D, Seo H, Jung M. Neural basis of reinforcement learning and decision making. *Ann Rev Neurosci*. 2012;35:287–308. [PubMed: 22462543]
18. Izawa J, Shadmehr R. Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comp Biol*. 2011;7:e1002012.
19. Trouillas P, Takayanagi T, Hallett M, Currier R, Subramony S, Wessel K, Bryer A, Diener H, Massaquoi S, Gomez C, Coutinho P, Hamida M, Campanella G, Filla A, Schut L, Timmann D, Honnorat J, Nighoghossian N, Manyam B. International Cooperative Ataxia Rating Scale for pharmacological assessment of the cerebellar syndrome. *J Neurol Sci*. 1997;145:205–211. [PubMed: 9094050]
20. Bastian A Learning to predict the future: the cerebellum adapts feedforward movement control. *Curr Opin Neurobiol*. 2006;16:645–649. [PubMed: 17071073]
21. Zimmet A, Bastian A, Cowan N. Cerebellar patients have intact feedback control that can be leveraged to improve reaching. *bioRxiv*. 2019.
22. Ilg W, Bastian A, Boesch S, Burciu R, Celnik P, Claasen J, Feil K, Kalla R, Miyai I, Nachbauer W, Schöls L, Strupp M, Synofzik M, Teufel J, Timmann D. Consensus paper: Management of degenerative cerebellar disorders. *Cerebellum*. 2014;13:248–268. [PubMed: 24222635]
23. Uehara S, Mawase F, Therrien A, Cherry-Allen K, Celnik P. Interactions between motor exploration and reinforcement learning. *J Neurophysiol*. 2019; 122: 797–808. [PubMed: 31242063]
24. Cashaback J, McGregor H, Mohatarem A, Gribble P. Dissociating error-based and reinforcement-based loss functions during sensorimotor learning. *PLOS Comp Biol*. 2017;13:e1005623.
25. Wong A, Marvel C, Taylor J, Krakauer J. Can patients with cerebellar disease switch learning mechanisms to reduce their adaptation deficits? *bioRxiv*. 2018

26. Beppu H, Suda M, Tanaka R. Analysis of cerebellar motor disorders by visually guided elbow tracking movement. *Brain*. 1984;107:787–809. [PubMed: 6478179]
27. Day B, Thompson P, Harding A, Marsden C. Influence of vision on upper limb reaching movements in patients with cerebellar ataxia. *Brain*. 1998;121:357–372. [PubMed: 9549511]
28. Miall R, Christensen L, Cain O, Stanley J. Disruption of state estimation in the human lateral cerebellum. *PLoS Biol*. 2007;5:e316. [PubMed: 18044990]
29. Mazzoni P, Krakauer J. An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci*. 2006;26:3642–3645. [PubMed: 16597717]
30. Benson B, Anguera J, Seidler, R. A spatial explicit strategy reduces error but interferes with sensorimotor adaptation. *J Neurophysiol*. 2011;105:2843–2851. [PubMed: 21451054]
31. Taylor J, Ivry R. Flexible cognitive strategies during motor learning. *PLoS Comp Biol*. 2011;7:e1001096.
32. McDougle S, Bond K, Taylor J. Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *J Neurosci*. 2015;35:9568–9579. [PubMed: 26134640]
33. Schween R, McDougle S, Hegele M, Taylor J. Assessing explicit strategies in force field adaptation. *J Neurophysiol*. 2020;123:1552–1565. [PubMed: 32208878]
34. Codol O, Holland P, Galea J. The relationship between reinforcement and explicit control during visuomotor adaptation. *Sci Rep*. 2018;8:9121. [PubMed: 29904096]
35. Holland P, Codol O, Galea J. Contribution of explicit processes to reinforcement-based motor learning. *J Neurophysiol*. 2018;119:2241–2255. [PubMed: 29537918]
36. Butcher P, Ivry R, Kuo S, Rydz D, Krakauer J, Taylor J. The cerebellum does more than sensory prediction error-based learning in sensorimotor adaptation tasks. *J Neurophysiol*. 2017;118:1622–1636. [PubMed: 28637818]
37. McDougle S, Tsay J, Taylor J, Ivry R. Cerebellar degeneration selectively disrupts continuous mental operations in visual cognition. *bioRxiv*. 2020.
38. Campbell WW. *DeJong's the neurologic examination*. Baltimore: Lippincott Williams and Wilkins. 2005

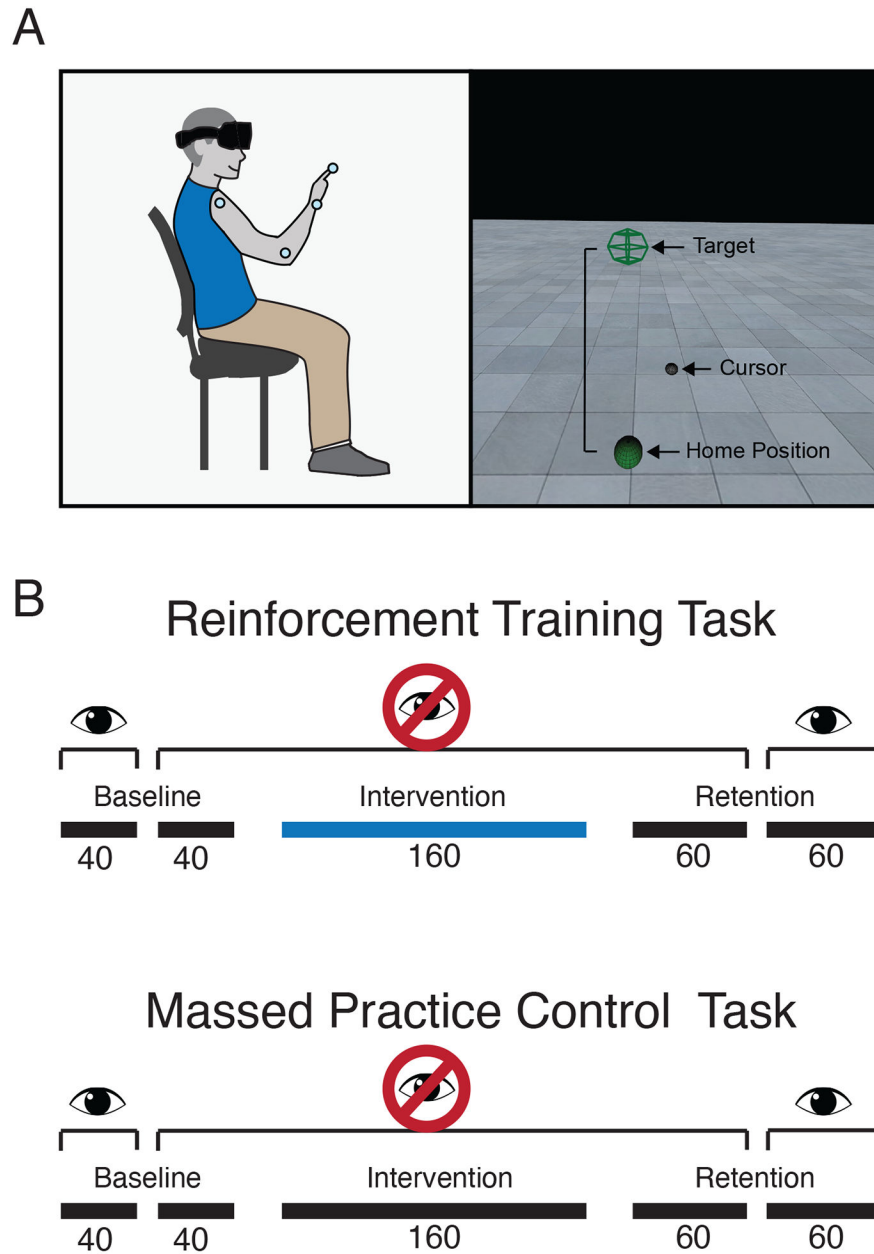


Figure 1. Schematic of experimental set up and task overview.

(A) Participants sat in a chair while wearing a virtual reality headset and motion capture markers on the shoulder, elbow, wrist and index fingertip of the dominant arm. In a virtual environment, they made reaching movements from a home position to a single target. Both the home and target positions were dependent on participants' anthropometrics. (B) Cerebellar patients completed two reaching tasks: a reinforcement training task and a massed practice control task. Both tasks began with a baseline phase of 40 trials where a visual cursor representing the index fingertip position was shown throughout the reach and 40 trials where the cursor was invisible. The baseline phase was followed by a 160-trial intervention phase. In the reinforcement training task, the intervention had participants receive closed-loop reinforcement feedback about the path length of their reaches. Reaches

with a path length at or below a running mean of the preceding 10 reaches were reinforced with a pleasing tone at the end of the movement and a point appearing in a tracker visible in the virtual display. Reaches with a path length greater than the running mean resulted in an unpleasant tone and no points added. Participants were not informed that reinforcement signaling was related to path length. The intervention phase of the massed practice control task had participants repeat reaches to the target without reinforcement. In both tasks, the intervention phase was followed by a retention phase where reinforcement feedback was removed, divided into 60 trials where the visual cursor remained invisible and 60 trials where it was restored.

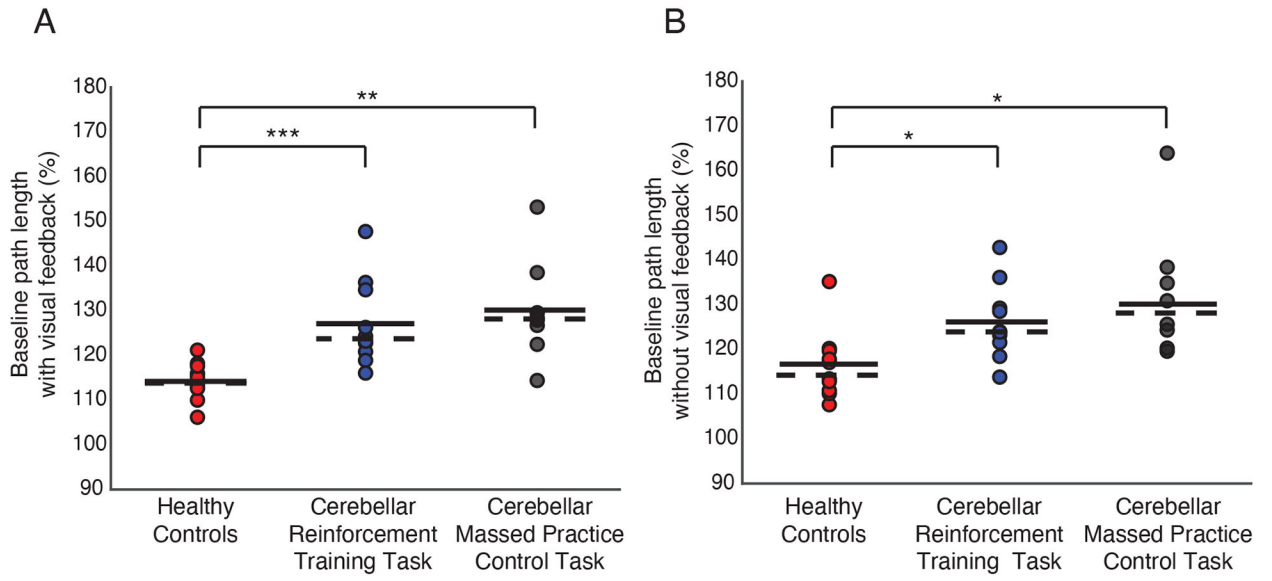


Figure 2. Comparison of baseline path lengths between cerebellar patients and healthy controls. (A) Reach path lengths when visual feedback of index fingertip position was shown. (B) Reach path lengths when visual feedback of index fingertip position was removed. Reach path lengths are expressed as a percentage of the minimal straight-line distance between the home and target positions. Markers represent individual subjects. Solid horizontal lines represent sample mean. Dashed horizontal lines represent sample median. * $p < .05$, ** $p < .01$, *** $p < .001$.

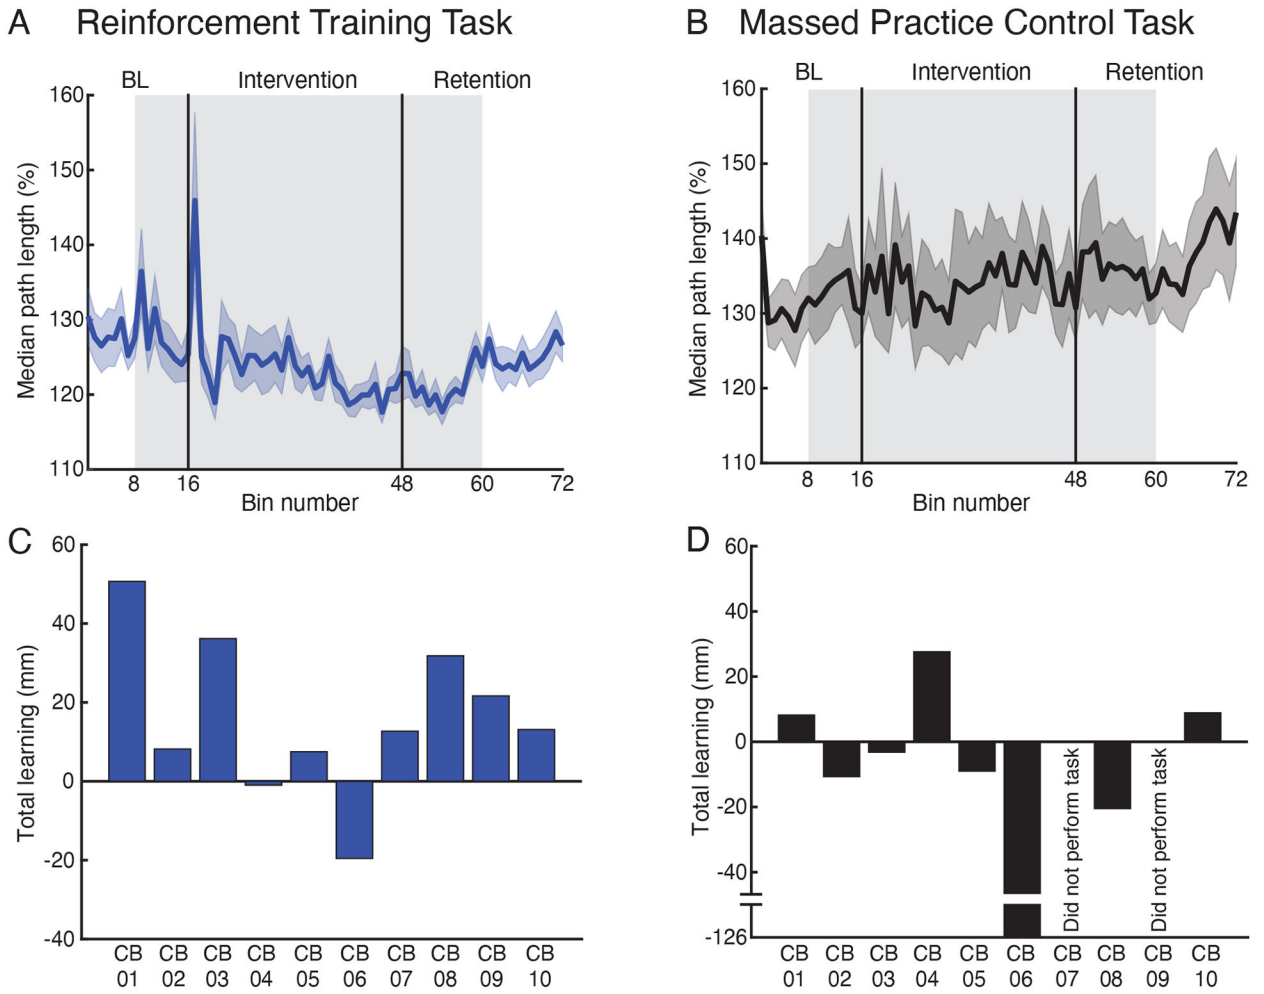


Figure 3. Cerebellar patients' performance in the reinforcement training and massed practice control tasks.

(A-B) Group reach path lengths over trial bins in the reinforcement training task (A) and massed practice control task (B). Each bin comprises the mean of 5 trials. Reach path lengths are expressed as a percentage of the minimal straight-line distance between the home and target positions. The grey shaded region highlights trials where visual feedback was removed. (C-D) The total learning shown by each patient in the reinforcement training task (C) and massed practice control task (D). The total learning was defined as the difference in median path length from the baseline phase without visual feedback (trials 41–80) to the end of the intervention (trials 221–240).

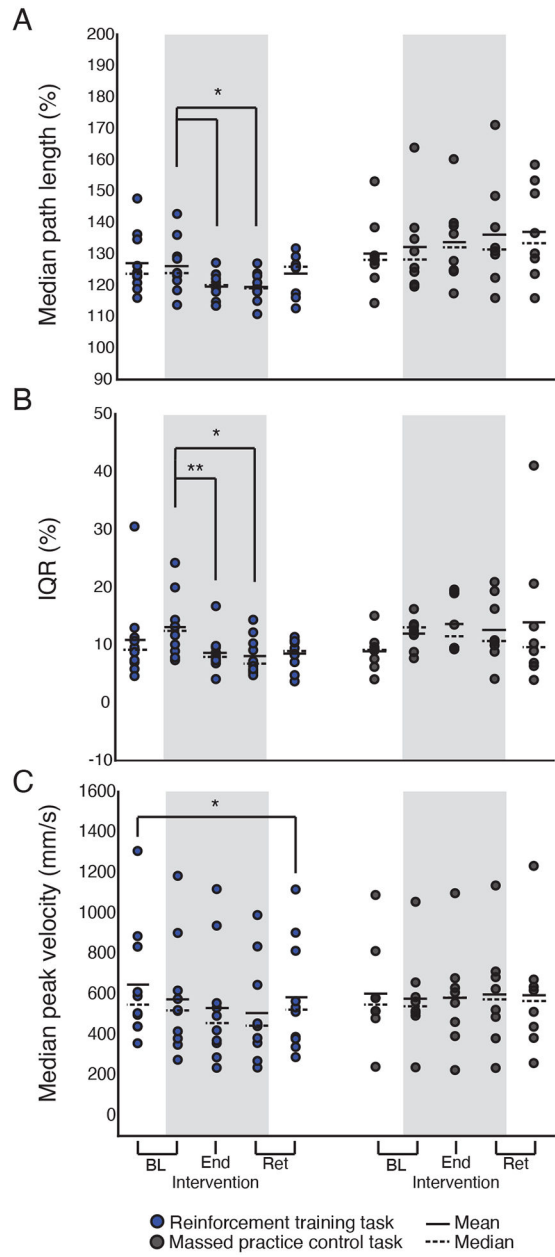


Figure 4. Comparison of reinforcement training task and massed practice control task. Group results for median path length (A), path length interquartile range (B) and median peak velocity (C). Median path length and the path length interquartile range are expressed as a percentage of the minimal straight-line distance between the home and target positions. In each panel, the grey shaded region highlights task phases where visual feedback was removed. Markers represent individual subjects. Solid horizontal lines represent sample mean. Dashed horizontal lines represent sample median. * $p < .05$, ** $p < .01$.

Reinforcement Training Task

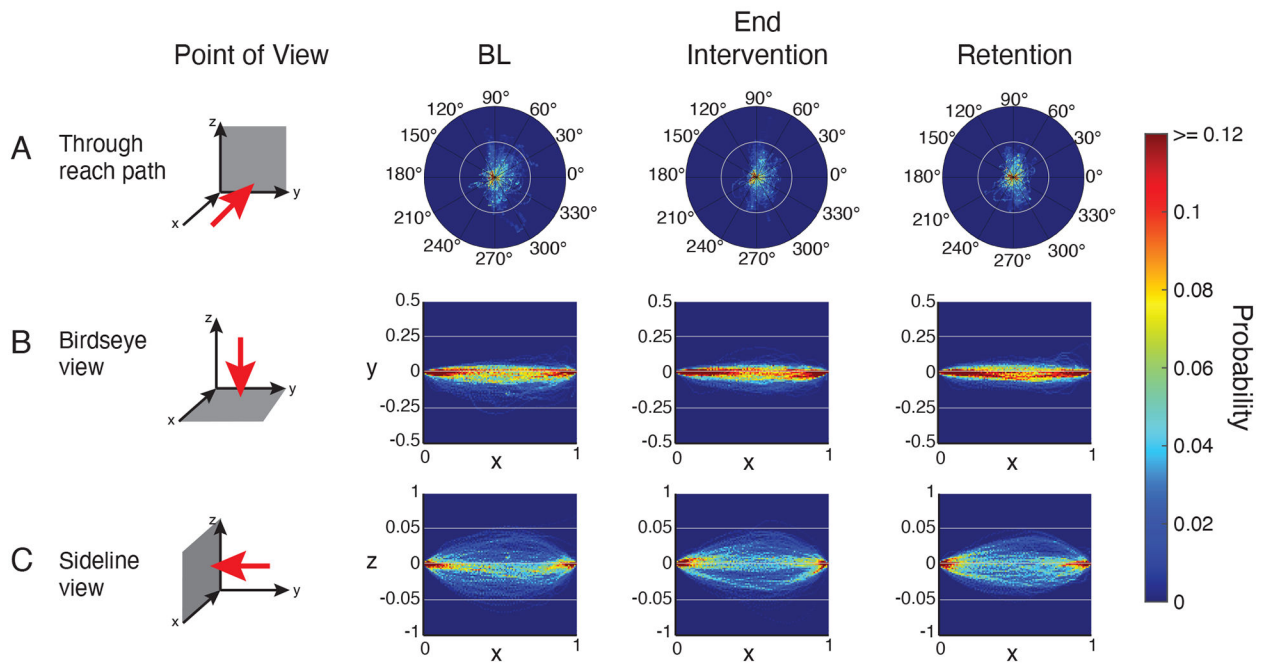


Figure 5. Group reach trajectories in the reinforcement training task phases where visual feedback was removed.

Three-dimensional reach trajectories normalized by trajectory length and plotted as probability distributions in 3 planes: Z-Y plane (A), X-Y plane (B), X-Z plane (C).

Massed Practice Control Task

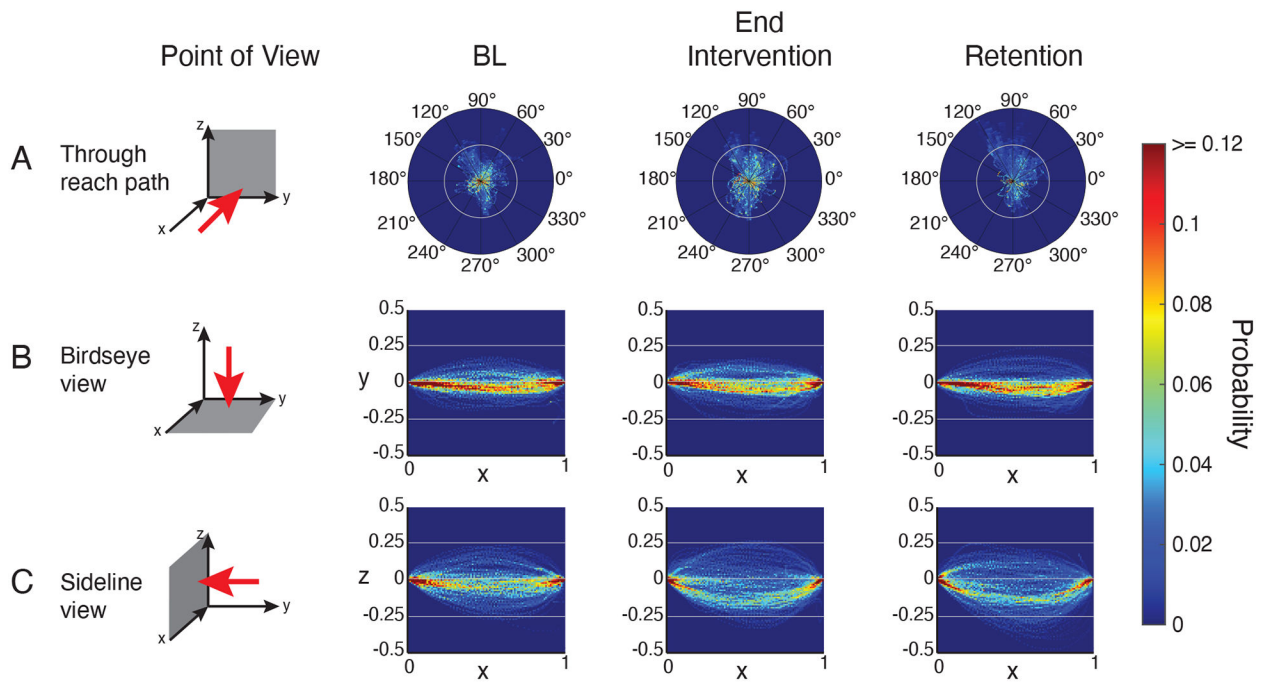


Figure 6. Group reach trajectories in the massed practice control task phases where visual feedback was removed.

Three-dimensional reach trajectories normalized by trajectory length and plotted as probability distributions in 3 planes: Z-Y plane (A), X-Y plane (B), X-Z plane (C).

Table 1.

Subject Demographics

Subjects	Age (years)	Gender	Handedness	Diagnosis	ICARS	
					Total (/100)	Limb (/52)
CB01	67	M	R	ADCAIII	12	4
CB02	73	F	R	SCA 6	49	22
CB03	55	M	R	Sporadic	59	31
CB04	57	F	R	SCA 8	41	23
CB05	69	F	R	ADCAIII	52	21
CB06	65	M	R	SCA 6	42	18
CB07	63	M	R	SCA 6 & 8	62	23
CB08	21	F	R	SCA 8	35	19
CB09	54	M	L	ADCAIII	28	9
CB10	70	M	R	SCA 8	38	16
CB Group	59.4 ± 15.0	M = 6/10	R = 9/10		41.8 ± 15.0	18.6 ± 7.6
Control Group	54 ± 11.5	M = 3/10	R = 9/10			

ICARS: International Cooperative Ataxia Rating Scale. CB: cerebellar patient. Control group age matched to cerebellar patients. F: female. M: male. R: right handed. L: left handed. ACDA III: autosomal dominant cerebellar ataxia type 3. SCA: spinocerebellar ataxia types 6 and 8. Sporadic: sporadic adult-onset cerebellar ataxia. Group mean ± SD. None of the patients or controls had sensory loss in clinical tests of proprioception and monofilament testing for tactile sensation (Campbell, 2005, [38]).

Table 2.

Comparison of baseline reach path lengths between cerebellar patients and healthy controls

Comparison	Group	n	Mean Rank	<i>U</i>	<i>z</i>	Sig (2-tailed)
BL-V	Control	10	6.10	6.00	-3.326	.001
	CB Reinf	10	14.90			
	Control	10	6.00	5.000	-3.110	.002
	CB MP	8	13.88			
BL-NV	Control	10	7.20	17.00	-2.495	.013
	CB Reinf	10	13.80			
	Control	10	6.60	11.00	-2.577	.010
	CB MP	8	13.13			

Means were compared using Mann-Whitney U tests at alpha = .05. BL-V: Baseline with visual feedback, BL-NV: Baseline without visual feedback, Control: control group, CB: Cerebellar group, Reinf: Reinforcement training task, MP: Massed practice control task. Significant results are bolded.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3.

Comparison of baseline values between tasks

Dependent Variable	Condition	Comparison	Mean Ranks (positive)	Sum of Ranks (positive)	z	Sig. (2 tailed)
Median path length	V	Reinforce Mass Prac	4.00	8.00	-1.400	.161
	NV	Reinforce Mass Prac	3.33	10.00	-1.120	.263
Path length IQR	V	Reinforce Mass Prac	4.00	16.00	-0.280	.779
	NV	Reinforce Mass Prac	4.20	21.00	-0.420	.674
Median peak velocity	V	Reinforce Mass Prac	4.50	27.00	-1.260	.208
	NV	Reinforce Mass Prac	3.83	23.00	-0.700	.484

Means were compared using Wilcoxon signed rank tests at alpha = .05. NV: No vision, V: vision, Reinforce: Reinforcement training task, Mass Prac: Massed practice control task. No significant differences were observed.

Table 4.

Comparisons of dependent values across experiment phases within task

Task	Dependent Variable	Comparison	Mean Ranks (positive)	Sum of Ranks (positive)	<i>z</i>	Sig. (2 tailed)
Reinforcement	Median path length	BL-V Ret-V	4.25	17.00	-1.070	.285
		BL-NV End. Int	4.00	8.00	-1.988	.047
		BL-NV Ret-NV	2.67	8.00	-1.988	.047
	Path length IQR	BL-V Ret-V	4.14	29.00	-0.153	.878
		BL-NV End. Int	2.00	2.00	-2.599	.009
		BL-NV Ret-NV	5.00	5.00	-2.293	.022
	Median peak velocity	BL-V Ret-V	2.33	7.00	-2.090	.037
		BL-NV End. Int	3.33	10.00	-1.784	.074
		BL-NV Ret-NV	6.00	12.00	-1.580	.114
Massed Practice	Median path length	BL-V Ret-V	4.57	32.00	-1.960	.050
		BL-NV End. Int	4.60	23.00	-.700	.484
		BL-NV Ret-NV	4.33	26.00	-1.120	.263
	Path length IQR	BL-V Ret-V	5.00	25.00	-0.980	.327
		BL-NV End. Int	7.00	21.00	-0.420	.674
		BL-NV Ret-NV	4.50	18.00	0.000	1.00
	Median peak velocity	BL-V Ret-V	3.60	18.00	0.000	1.00
		BL-NV End. Int	5.00	20.00	-0.280	.779
		BL-NV Ret-NV	4.80	24.00	-0.840	.401

Means were compared using Wilcoxon signed rank tests at alpha = .05. Significant results are bolded.