1

2

Exploring Disturbance as a Force for Good in Motor Learning

Jack Brookes[†1,2], Faisal Mushtaq[†*1], Earle Jamieson[†1,2], Aaron J. Fath[3],

Geoffrey P. Bingham[3], Peter Culmer[2], Richard M. Wilkie[1] & Mark A. Mon-

Williams[1]

[1]School of Psychology, University of Leeds, Leeds, West Yorkshire, LS2 9JT, United Kingdom

[2]School of Mechanical Engineering, University of Leeds, Leeds, West Yorkshire, LS2 9JT, United Kingdom

[3]Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405-7007, USA

*Correspondence: Faisal Mushtaq, School of Psychology, University of Leeds,

Leeds, West Yorkshire, LS2 9JT, United Kingdom f.mushtaq@leeds.ac.uk

[†]These authors contributed equally to this work

Manuscript Type: Research Article

Abbreviated Title: Disruption and Motor Learning

Number of pages: 36

Number of Figs: 8

Total number of words: 9380

Abstract

27   Disturbance forces facilitate motor learning, but theoretical explanations for this

28   counterintuitive phenomenon are lacking. Smooth arm movements require

29   predictions (inference) about the force-field associated with a workspace. The Free

30   Energy Principle (FEP) suggests that such 'active inference' is driven by 'surprise'.

31   We used these insights to create a formal model that explains why disturbance helps

32   learning. In two experiments, participants undertook a continuous tracking task

33   where they learned how to move their arm in different directions through a novel 3D

34   force field. We compared baseline performance before and after exposure to the

35   novel field to quantify learning. In Experiment 1, the exposure phases (but not the

36   baseline measures) were delivered under three different conditions: (i) robot haptic

37   assistance; (ii) no guidance; (iii) robot haptic disturbance. The disturbance group

38   showed the best learning as our model predicted. Experiment 2 further tested our

39   FEP inspired model. Assistive and/or disturbance forces were applied as a function

40   of performance (low surprise), and compared to a random error manipulation (high

41   surprise). The random group showed the most improvement as predicted by the

42   model. Thus, motor learning can be conceptualised as a process of entropy

43   reduction. Short term motor strategies (e.g. global impedance) can mitigate

44   unexpected perturbations, but continuous movements require active inference about

45   external force-fields in order to create accurate internal models of the external world

46   (motor learning). Our findings reconcile research on the relationship between noise,

47   variability, and motor learning, and show that information is the currency of motor

48   learning.

Introduction

Neonates must determine the complex relationship between perceptual outcomes and motor signals in order to learn how to move their arms effectively. This process is repeated throughout life as humans calibrate to new environments, acquire new skills, experience neuromuscular fatigue or recover from injury. Technological advances have created robotic systems designed to accelerate the acquisition of skilled arm movements in a variety of areas including, amongst others, laparoscopic surgical training and stroke rehabilitation [1]. These devices can provide assistive forces that guide an individual's arm through a desired trajectory or apply disturbance forces that make it more difficult for the individual to move their arm along a given trajectory.

It is now well established that providing assistive forces to neurologically intact individuals can actually impair subsequent learning [2,3]. Conversely, there is growing empirical evidence that providing disturbance forces to impair performance during training of a motor task can have a net positive effect, and lead to improved learning - enhancing performance in the task after the disturbance forces are removed [3–8]. However, formalised theoretical explanations that can account for these counterintuitive phenomena have proven elusive [9]. This is disappointing because it remains unclear how robotic devices might be best optimised in order to enhance learning (beyond this binary observation of differences between assisting and disturbing forces). The lack of a theoretical framework also makes it difficult to explain formally why assistive forces can be beneficial for individuals with neurological impairment [10], and the absence of a framework is hindering the potential utility of robotic technology in motor training. We propose that a 'Shannon' information theory perspective [11,12] could provide a principled approach to understanding why disruptive forces can be beneficial, and such an account could ultimately inform the development of haptic interventions.

2

76   The free energy minimization principle is the leading theoretical explanation of

77   brain and behaviour within the domain of neuroscience, and it accounts for many

78   empirical data within a unifying action, perception and learning framework [13–15].

79   The free-energy principle suggests that biological systems act to minimise free

80   energy (an information theory measure that limits the surprise associated with

81   sampling data). In this conceptualisation, the brain behaves as an active inference

82   machine that formulates predictions about the environment [16]: the better the

83   predictions about the environment, the lower the amount of free energy. Thus, the

84   process of effective motor learning involves the system making increasingly accurate

85   predictions about the perceptual outcome of motor commands given the current state

86   of the system. In other words, the system will minimise entropy (the average amount

87   of surprise) through the development of 'forward models' that act as neural

88   simulators regarding how the current state of the system will respond to a given

89   motor signal [17].

90   Viewed in this way, motor learning requires the system to sample information in

91   order to extract the invariant rules that govern a range of input–output mappings

92   [18,19]. The difficulty faced by the system relates to the large number of internal

93   parameters that connect the sensory input to the motor output i.e. high levels of

94   uncertainty [20]. The example of a neonate learning the mapping between perceptual

95   and motor output illustrates how this problem can be framed from an information

96   theory perspective. The newborn must use information generated from her

97   exchanges with the environment in order to learn the input–output mappings and

98   subsequently refine her predictions, so that she can successfully interact with her

99   new surroundings. The initial reaches will be associated with high levels of

100   uncertainty and will thus have high informational entropy (the average surprise of the

101   outcomes sampled from the probability density). The developmental trajectory,

102   however, will be marked by a reduction in entropy as the certainty of a predictable

103   perceptual outcome following the generation of a motor command will increase.

104   Thus, motor learning can be viewed as a process where entropy (i.e., uncertainty) is

105   reduced through the development of forward models following exposure to

106   information regarding the relationship between perceptual output and motor signal

107   input [16].

108   We propose that this information perspective can account for the previous finding

109   of superior learning outcomes from disturbance haptic force application relative to

110   assistive guidance. Specifically, we suggest that providing assistive forces limits the

111   amount of surprise experienced by the actor and thus constrains the amount of

112   learning. Conversely, disturbance forces expose the individual to more information

113   which facilitates the learning process. Following this logic, a control algorithm that

114   provides an optimal level of surprise should lead to better learning than those that

115   minimise uncertainty. It will be noted that a certain level of motor proficiency is

116   required to sample information within a workspace – if an individual is unable to

117   move their arm through the space then they will be unable to experience the surprise

118   necessary to even start the learning process. This may explain why assistive forces

119   have been found to help individuals with severe neurological impairment [4,21,22] or

120   lesser skilled individuals [3,23] – as these systems allow the individual to sample the

121   requisite information and thereby start the learning process.

122   Our approach is based on the idea that skilful arm movements require accurate

123   predictions about the forces acting on the arm as it moves around the workspace. If

124   these predictions are inaccurate then the system must contend with unexpected

125   perturbations that will force the arm away from its desired trajectory. It has been

126   shown that participants can learn to attenuate the impact of an unexpected

127   perturbation in the short term by developing a 'global impedance' strategy, where

128   joint stiffness rapidly increases in response to the application of a sudden

129   unexpected force[24,25]. The development of a 'global impedance' strategy is a

4

130   useful short term response to environments which contain unpredictable forces.

131   Nevertheless, skilled continuous movements through a workspace require accurate

132   forward models that allow low entropy, suggesting that the system will seek to learn

133   (and thus predict) the underlying force field in which it is operating. On this basis, we

134   predicted that exposure to a complex force field would, over a sufficient period, drive

135   the system to learn how to move skilfully through the workspace (rather than

136   adopting a short term global impedance strategy).

137   To test these ideas, we created a metric that quantified the information sampled

138   as individuals learned to move their hand around an artificial environment containing

139   a complex force field (equivalent to moving the arm through a novel viscous

140   solution). The environment was specifically designed to produce sufficient novelty to

141   limit the possibilities of existing forward models being adapted. These steps allowed

142   us to examine novel motor learning in two experiments whilst providing distinct types

143   of assistive and disturbance forces using an admittance-controlled robotic device. In

144   our second experiment, we created a condition that would enhance learning if the

145   Free Energy Principle inspired model has merit but would not be expected to benefit

146   learning if the system were simply adopting a short term global impedance strategy

147   to cope with the force field.

148   In our experience, there are two points worth highlighting with regard to the

149   reported experiments. First, the experiments appear to have a similarity with a study

150   run within Kawato's laboratories [25]. The method section below should make it clear

151   that the similarity is superficial. In the Kawato study, participants moved their arm

152   along a prescribed path through a normal force field but were exposed to an

153   unexpected perturbation when the arm diverged from the desired spatial path

154   (resulting in participants learning to stiffen their arm in response to such

155   perturbations). In our experiments, participants had to make continual movements

156   through a workspace comprising a completely novel force field. This arrangement

5

157   meant that our participants had to learn the underlying structure of the force field –

158   the experiments were not about the participants moving normally and then suddenly

159   experiencing a perturbation of an unpredictable nature. Second, our experiments

160   included baseline measurements of how well the participants could move their arms

161   in the novel force field. These measurements were taken before and after the

162   participants were given the opportunity to learn the task. The baseline measures did

163   not involve the experimental manipulations (where the robot provided assistive or

164   disruptive forces during the learning process). Thus, the baseline measures provided

165   an index of the motor learning that occurred throughout the experimental sessions.

166   These measures provided the data that we needed to test the predictions of our

167   model.

Materials and Methods

Participants

In Experiment 1, forty-eight right-handed participants (26 male) (M = 29.4 years, SD = 9.34 years, range 20–59 years) were recruited and randomly allocated to one of three training groups: Assistance (n = 15), Active-Control (n = 16) and Disruption (n = 17). One participant from the Active-Control group voluntarily withdrew from the experiment and their data were excluded from further analysis.

In Experiment 2, forty-six right-handed participants (25 male, aged 19 - 56 years, M = 24.93 years, SD = 6.36 years) were randomly allocated to the Adaptive Algorithm (n = 13), Adaptive Disruptive (n=17) and Random (n = 16) conditions. One participant withdrew voluntarily from the Random group after the first session and their data were not included for statistical analysis. The Psychology Research Ethics Committee at the University of Leeds approved the research.

Procedure

In the two reported experiments using a task that required continuous tracking through a complex novel three dimensional force field, participants stood in front of a haptic robot system (HapticMASTER, see Materials) and visual stimuli were displayed on a monitor located behind the device, approximately at eye level (see Fig 1A). Two cursors were used to visually represent the actual hand and the target position of the device end-effector within the workspace on the visual display (see Fig 1D). Upon reaching the start position, the cursor started moving immediately along the first component (sub-path) for that trajectory at a constant speed of 0.1 m/s. Participants were instructed to use their preferred (right) hand to align the end-effector with a moving target as accurately as possible along pre-specified

7

193   trajectories. Movement was in the Y-Z plane of the HapticMASTER system (Z –

194   vertically upwards, Y – horizontally right relative to participant). The target cursor

195   waited until the end of the component was reached by the participant before the next

196   component began.

197   Participants were required to attend five sessions (one per day for 5 consecutive

198   days) of approximately 15 minutes each. In sessions 1 and 5, participants followed a

199   pentagram trajectory for three blocks of ten trials. Participants moved within a

200   workspace force field, but had no error manipulation forces. This trajectory was

201   based on 2D aiming tasks that have previously been used in the assessment of

202   manual dexterity [26]. The pentagram contains five straight line components of equal

203   length (the five edges). In Experiment 1, sessions 2 to 4 (Training) each consisted of

204   four blocks of ten trials, with either assistive (error reducing), no or disruptive (error

205   enhancing) forces (depending on the allocated group) superimposed over the

206   workspace force field, following an inverted pentagram trajectory.

207   The target cursor was a hollow red circle, and the 'current position' cursor was a

208   filled blue circle. A dotted black line was used to indicate the magnitude of the error

209   between the current position and target cursors. To minimize fatigue, self-paced

210   breaks with a minimum of 30 seconds rest (whilst standing or seated) were provided

211   after each block of trials. Each session lasted approximately fifteen minutes.

212   Experiment 2 followed same trial structure as Experiment 1, with the exception of the

213   levels of assistance/disruption, which change trial by trial using various algorithms

214   depending on group.

215   **Materials**

216   The experiments reported here were designed to examine how error

217   manipulation forces affect the learning of a novel workspace force field. The

8

218   HapticMASTER, an admittance-controlled haptic device with a large workspace [27],

219   was used to generate the forces and record kinematics at a rate of 1 kHz.

220   To simulate a novel environment, we created a workspace force field which was

221   a function of position and calculated from the following equations:

$$f_y = 1\sin\left(\frac{2\pi}{0.1}z\right) \qquad\qquad (1)$$

$$f_z = 1\sin\left(\frac{2\pi}{0.1}y\right) \qquad\qquad (2)$$

222   The force from the workspace force field (newtons) was a function of position

223   (y and z, measured in meters) only. From this emerged a relatively novel

224   environment (Fig 1C) where, in order to perform well in the task, participants needed

225   to learn to predict the consequences of motor commands sent to the arm. Error

226   manipulation forces (those that acted to reduce or augment execution error) were

227   subsequently implemented using a mass-spring-damper model, as described in

228   Equation (3):

$$F = m\ddot{x} + c\dot{x} + kx \qquad\qquad (3)$$

229   where (x) is displacement between the end effector and target positions and

230   force is computed as a function of the distance between the actual and target

231   positions of the end-effector. The simulation was implemented in a virtual null-gravity

232   environment, and the end-effector mass, m, set to 3 kg and the damping, c, was set

233   to 10 Ns/m to generate an inertial effect.

234   In Experiment 1, for the Active-Control condition, the stiffness k was set to 0

235   N/m and therefore no forces directly related to the positional error. The assistance

236   group were provided with an assistive force implemented using k = 100 N/m, thereby

237   providing full assistance, and minimizing workspace information sampling. The

238   Disruption group had a disturbance force generated using coefficients k = -100 N/m,

239   thereby providing a large prediction error for initial interactions in this condition and

240   subsequently facilitation a larger range of movement around the workspace and

241   information sampling.

242   In Experiment 2, we varied workspace information acquisition whilst also

243   manipulating the possibility of developing a short term global impedance strategy.

244   Specifically, we created three new training algorithms. In the Adaptive Algorithm (AA)

245   - the virtual spring stiffness (k) varied as a function of task performance (i.e.

246   participants had increased disturbance when performance improved and increased

247   assistance when performance declined). The first trial of the Adaptive-Algorithm

248   condition was always set to no intervention (k = 0 N/m and c = 0 Ns/m) in order to

249   obtain a common benchmark measure of performance at the start of each session.

250   The value of the stiffness coefficient at each trial was adjusted as a function of

251   performance in previous trials, as described by Equation (4). This algorithm has been

252   used previously as a computational model of motor adaptation to predict the force

253   required to minimize adaptation time to a viscous environment during treadmill

254   walking tasks [1]. In our experiment, we used the model to adjust the value of the

255   stiffness coefficient in the current trial as a function of performance in previous trials.

256   This allowed us to consistently keep the amount of error experienced by a participant

257   within a small window:

$$k_{i+1} = f.k_i - g(x_i - x_d) \tag{4}$$

258   The stiffness, k, of the force field for the next trial is a function of the stiffness in

259   the current trial, i, multiplied by a 'forgetting factor', f, and the difference between the

260   demand error and actual error (x_d and x_i, respectively), multiplied by a gain value, g.

261   The values of f and g dictate the relative sensitivity of the algorithm to previous

262   performance (captured by k_i) and error. The sensitivity of the controller to

263   performances obtained in previous trials is controlled by adjusting f: A larger

264   forgetting factor weights the previous trials more heavily, whereas a smaller

10

265   forgetting factor results in more influence from the current trial's force field

266   magnitude. Pilot testing informed the values of *f* and *g* to be used in the experiment

267   and these were subsequently set at 0.5 each.

268   This approach allowed us to constrain the amount of information, as the level of

269   stiffness was tuned to individual performance, constraining information by means of

270   reducing workspace exploration since forces were always at a manageable level.

271   The *Adaptive Disturbance* (AD) condition was identical to the AA condition, but

272   stiffness could only decrease or stay the same between trials (i.e., the change in

273   stiffness' upper limit was 0). This similarly constrained information, but provided

274   increasingly disruptive forces and therefore facilitated development of a global

275   impedance strategy. Finally, performances in these conditions were compared

276   against a *Random* (RAN) group - where an unpredictable stiffness value was

277   provided (disturbance or assistance) across trials. The range of the value of k was

278   clamped in the range -100 and 100 N/m in all 3 algorithms.

279

280   **Fig 1 – Experiment Design** (a) Plan view of the experimental setup showing the relative

281   positions of the participant (bottom), haptic robot arm (middle) and monitor (top); (b) The

282   target trajectories across sessions. The pre- and post-training sessions comprised 3 blocks of

283   10 trials following a pentagram trajectory (with no error manipulation forces). Training (across

284   three sessions with 4 blocks of 10 trials) included error manipulation forces whilst participants

285   navigated across a vertically rotated pentagram trajectory. (c) Quiver plot of the novel

286   workspace force field used across all training sessions and conditions (discretized for

287   illustrative purposes). Inset shows magnified section (approximate size 6cm x 6cm). Arrows

288   indicate the direction and proportional magnitude of the force vector at discrete locations

289   within the workspace. Relative magnitude is shown from white (no force) through to red (high

290   force). (d) Blue cursor indicates the cursor (hand) position during a trial, the red circle

291   indicates the target, the dotted black line shows the participant's current positional error. A

292   virtual spring sits between the cursor and the target and provides assistance, disruption, or no

293   intervention depending on the value of *k*. N.B. Trajectory path and workspace force field

294   remained invisible to participants throughout the experiment.

295   **Metrics**

296   *Motor Learning*

297   Assessment (pre- and post-training) was performed without a spring stiffness (k=0),

298   but with the same workspace force field shown in Fig 1(c). Thus, 'learning' can

299   conceptually be defined as the participant's ability to predict, and counteract, the

300   forces arising from the workspace force field in order to minimize error. To capture

301   how much learning occurred following training in each condition, we calculated the

302   difference in performance in the pre- and post-training training trials. Specifically, we

303   calculated the mean average path error scores for the three pre-test blocks and

304   subtracted this value from the mean average path error scores from the post-test

305   trials. Path error (E_P) was computed as the mean Euclidian straight line distance

306   between the end effector and the current component (sub-path) of the target

307   trajectory. The position of the end effector was subject to a low-pass Butterworth filter

308   (cut-off 250Hz) to remove noise in analysis of movements.

309

310   *Analysis of Training Data*

311   To study changes in performance as a function of training trial, we fitted a first order

312   exponential equation to the training data using the 1st order exponential fit function in

313   the Curve Fitting Toolbox implemented in MATLAB (MathWorks Inc., Natick, MA).

314   Training block number was used as the x value (x = 1 being the first block in the first

315   training session), and average path error during training was used as the y value.

316   The function uses the method of least squares to produce the most probable values

317   of a and b in the function. The values derived from this model for each individual

318   were subjected to group-level analysis to examine differences during training. In

319   other words, we used the parameters of the learning function as summary statistics

320   for random effects analysis using classical inference (i.e. ANOVA).

321   **Quantifying Information**

322   To obtain a metric of information, we first parsed the workspace into discrete,

323   independent voxels of 1 cm x 1 cm (see **Fig 2**; total size 40 cm x 40 cm). For the

324   purposes of analysis, we created a model that assumed participants acquire

325   information about the force output of discrete voxels, and any information acquired

326   when the cursor was located inside a particular voxel was 'assigned' to that voxel. As

327   information is accumulated for a particular voxel, newly acquired information for that

328   voxel is discounted in value according to a weighting function. Weighting the

329   information in this way ensures that initial "inaccurate" estimates about the expected

330   change in force results in high amounts of surprise, and as more information is

331   acquired, lower amounts of surprise. Effectively, the system logarithmically scales

332   ("weights") information in each voxel. The result of this is a metric which captures

333   information acquired through exploration of a workspace – a higher value will result

334   from visiting a large number of independent voxels across the workspace. The voxel

13

335   size of 1cm x 1cm was a largely arbitrary selection; modelling with different voxel

336   sizes in the range 0.25cm – 4cm shows the same pattern of results. Total weighted

337   information gained during training can be conceptualised of as a measure of entropy.

338   Participants were not informed about the underlying workspace force field

339   and it remained invisible throughout the experiment. Thus, without the presence of

340   visual information, we assumed that the sensorimotor system would have no reason

341   to predict a change in force as a function of cursor position (at least at the outset of

342   training). This heuristic leads to a context where the magnitude of the change in force

343   due to the workspace force field at that point in time corresponds to a force prediction

344   error (i.e. the difference between the experienced and predicted force). Thus, new

345   information presented about an individual voxel was approximated as the change in

346   force at a point in time for the voxel at the cursor position (**Fig 2b**). That is, the

347   magnitude of change of the force vector as calculated by the workspace force field

348   equations, Equations (**1**) and (**2**).

349   The information (I) related to a particular voxel (i,j) acquired throughout

350   training up to a time T (total time cursor was positioned inside the voxel) was

351   therefore:

$$I_{ij} = \int_0^T \Delta f(t)dt \tag{5}$$

352   Here, information is 'binned' into the voxel where the end effector position is

353   currently located (i, j). A value of I was computed for every voxel in the workspace

354   under the assumption that information presented for a particular voxel is the

355   magnitude of the change in force, numerically integrated over time for all points in

356   time where the cursor position was inside that voxel (**Fig 2b**). We assumed that new

357   information becomes less valuable as a function of the amount of information already

358   acquired about an individual voxel as learning occurs (where models about the

359   expected force arising from a particular voxel are updated to minimize free energy).

360   This means that observations of changes in force have a higher probability, and

361   therefore less surprise. Instead of using probability of sensory input estimates for

362   each observed change in force, we opted for a more parsimonious solution by

363   approximating surprise with a weighting function - scaling the amount of information

364   presented to an associated information 'value'.

365   The weighting method used has the desired effect for scaling information –

366   the gradient of the weighting function $= 1$ when information $= 0$ and gradually

367   decreases. Weighting the information in this way ensures that initial inaccurate

368   estimates about the expected change in force results in high amounts of surprise

369   and, as more information is acquired, the surprise is lower. The weighting formula, as

370   a function of information presented, was:

$$w(I_{ij}) = \frac{1}{\lambda} \cdot \log \left(\lambda\, I_{ij} + 1 \right) \tag{6}$$

371   where log is the natural logarithm and λ corresponds to a weighting parameter.

372   Higher values of λ lead to lower values of information relative to the amount of

373   cumulative information presented, and thus faster learning about a voxel. The

374   reported results have the value $\lambda = 0.05$, but we tested the model under different

375   assumptions of λ (through values ranging from 0.01 to 1.00) and the pattern

376   remained consistent.

377   We also assumed that the total weighted information (TWI) acquired was

378   equal to the sum of the value weighted information received from each voxel of the

379   workspace. If the workspace consists of N_x cells horizontally, and N_y cells vertically,

380   the information value for the whole workspace at time T can be calculated as:

15

$$TWI = \sum_{i=0}^{N_x} \sum_{j=0}^{N_y} w(I_{ij}) \tag{7}$$

381   In this case the total weighted information assumes that information sampling

382   starts at the beginning of the first training session (Session = 2) and completes at the

383   end of the last training session (Session = 4). The total weighted information was

384   computed per participant and is used in subsequent analyses.

385

386   **Figure 2 – Information quantification** (a) Example simulated cursor movement across a

387   sub-section of the workspace (10cm x 10cm). Workspace force field shown as a quiver plot,

388   where higher force magnitude is represented by darker red shading and arrow size, and force

389   direction indicated by arrow orientation. Workspace separated into 1cm x 1cm voxels. (b)

390   Magnitude of change in force measured when moving along the path shown in (*a*) at a

391   constant velocity over 1 second. Vertical black lines indicate the voxel boundary. Shaded

392   regions under the curve separated by the vertical lines represent the information presented

393   which is attributed to the current voxel. (c) Graphical representation of the weighting function

394   for different values of lambda. Note that at higher values of information (in a voxel), the

395   weighted information becomes relatively lower.

396

397   It is worth noting that we could have quantified information in alternative ways

398   to the approach described above. For example, one could model information

399   acquisition and parameter estimation as a Kalman filter, or using Bayesian inference.

400   However, unlike the participants in our experiments, such models would rapidly

401   converge to the true force in a given area in only a limited number of observations.

402   To circumvent this, we would need to make assumptions that involve including

403   parameters estimating sensory and processing noise to slow the rate of learning.

404   This would provide comparable results to our information scaling method if these

405   approaches were implemented in a discrete voxel based manner (as calculated here

406   - with exploration being rewarded as a means of sampling information and exposure

407   to new areas of the workspace providing more information). More sophisticated

408   models could capture the idea that repeated exposure to forces in a workspace is not

409   sufficient for learning per se- but these also require additional assumptions e.g. an

410   understanding (and model) of how an action (set of muscle contractions) is executed

411   to deal with the force to maintain low positional error. Given that our aim was

412   restricted to capturing the relationship between workspace exploration and

413   information acquisition, we settled on a solution that provided the most parsimonious

414   model of behavior in this task.

415

Statistical Analysis

417   One-way between subject ANOVAs were performed to examine differences

418   between the groups for each of the metrics described above, and Tukey's post-hoc

419   comparison corrected p values are reported where relevant. Partial eta squared (η^2_p)

420   values are reported for effect size. We tested for, but did not find any, violations of

421   the assumption of homogeneity of variance using Levene's test [28]. Error bars on all

422   Figures represent +/- 1 SEM.

423

Experiment 1 – Disturbance Leads to Increased Information Sampling

We first tested the prediction that learning rates could be accelerated through the increased information provided via disturbance forces. We examined training with partially assistive (Assistance group), disturbance (Disturbance group) and no guidance (Active-Control group) forces.

In the training period, the 'Disturbance group' were presented with an additional force vector, whose force was generated using a negative value of k in the mass-spring-damper simulation. We predicted that disturbance forces would lead to (i) more surprise (as indexed by our model of information); (ii) more errors at the outset of training – indexed by a in the fitted function $y = ae^{bx}$ and (iii) increased rate of error reduction over the training period (indexed by b); and finally, as a corollary of the above, (iv) superior motor learning compared (pre- post- error improvement) to the groups with lower information.

The differences in information at the early and late stages for each condition can be seen in **Fig 3**. Formal analysis of the cumulative amount of information for each group at the end of the training block revealed statistically significant differences ($F_{(2, 44)} = 34.21$, $p < .0001$, $\eta^2_p = .609$). This effect was driven by the Disturbance group gathering more information about the workspace relative to the Active-Control ($p < .0001$) and Assistance ($p < .0001$) groups, but there was no difference between the Assistance and Active-Control groups ($p = .876$).

Fig 3 - Information as a by-product of disruption. (a) The Disturbance group had more information over time at a group level; (b) Example heat maps showing the amount of information gathered across the workspace at the outset and end of training for randomly selected individual participants.

451   We next performed an ANOVA on the values for the exponential fit to examine

452   differences at the outset of training. The ANOVA revealed group differences (F (2,

453   44) = 7.623, p = .0014, η^2_p = .257), with the Disturbance group performing worse than

454   the Assistance group (p = .0009), although following correction for multiple

455   comparisons, this was not significantly different to the Active-Control group (p =

456   .1162). When comparing performance across training trials (F (2, 44) = 26.37, p <

457   .0001, η^2_p = .545), we found that the disturbance group showed a steeper decay in

458   error in comparison to the Active-Control (p < .0001) and Assistance Groups (p <

459   .0001). There was no difference between learning for the Assistance and Active-

460   Control conditions (p = .2589).

461   The amount of motor learning was quantified as the error improvement

462   between the mean pre- and post- path error score (both of which were performed

463   without any stiffness intervention [k = 0] and with the upright pentagram shape). We

464   found significant differences in the amount of motor learning between groups (F (2,

465   44) = 5.655, p = .0065, η^2_p = .204). Specifically, the group exposed to Disturbance

466   forces during training on the inverted pentagram trajectory had improved significantly

467   more than the Assistance (p = .0136) and the Active-Control (p = .0202) groups (**Fig**

468   **4**). These results are consistent with our model.

469

470   **Fig 4 - Disturbance accelerates skill acquisition.** (a) Disturbance force training produced

471   a steeper exponential performance curve during the training blocks. (b) The Disturbance

472   training group were able to generalize their learning better than Assistance and Active Control

473   groups, as measured by reduction in mean path error between pre- and post-tests

Experiment 2 – Manipulating Information Sampling Without Facilitating a Short Term Impedance Strategy

The results from Experiment 1 indicate that disturbance results in faster learning in a manner consistent with the hypothesised information-driven process. However, these results do not rule out the possibility that it is disturbance forces per se that facilitate learning. For example, in Experiment 1, the adoption of a short term global impedance strategy (e.g. stiffening arm in all directions when an unexpected force was encountered) in response to disturbance forces could not be ruled out (see [25]). In Experiment 2, we therefore created algorithms that varied the amount of stiffness between trials to facilitate or constrain workspace information acquisition, and importantly make it improbable that the adoption of a global impedance strategy could yield better performance (**Fig 5A-C** and **Fig 7A-C**). The Random training condition exposed participants to an environment with a large degree of uncertainty (i.e. larger magnitude of changes in stiffness and more frequent switches between positive and negative stiffness on a trial-by-trial basis), but with an average level of overall stiffness that was close to zero. This means development of a global impedance strategy would hinder performance under the random condition (as 50% of participants' trials were assisted with the virtual spring on average). It follows that a global impedance explanation would not account for improved performance, but the unpredictability of the stiffness between trials would induce a greater range of workspace sampling and provide the most amount of information. Thus, improved performance could be attributed to the increased exposure to information rather than the adoption of global impedance. In summary, if our hypothesis has merit then it would predict that the Random condition should lead to the best learning, whilst AA and ADA would impair learning (as they constrain information sampling).

In line with our experimental aims, the algorithms produced significantly different mean values of stiffness throughout training ($F_{(2, 41)} = 12.40$, $p < .0001$, $\eta^2_p = .377$),

20

501   mean trial-on-trial stiffness change (F (2, 41) = 931.9, p < .0001, η^2_p = .986), and

502   number of times the task switched from assistive to disruptive (or vice versa) (F (2,

503   41) = 67.25, p < .0001, η^2_p = .7664).

504 **Fig 5 - Emergent properties of the training algorithms**. The level of assistance (positive
505 stiffness) or error enhancement (negative stiffness) during training was varied on a trial-by-
506 trial basis per the participant's allocated group. We reasoned that the increased changes in
507 stiffness (panel a shows magnitude of mean stiffness change between trials plotted) and
508 switching between positive to negative stiffness values (panel b shows group average
509 number of switches throughout training plotted) afforded to the random group would result in
510 increased workspace information sampling and therefore greater surprise , through means
511 other than the provision of a high negative stiffness (panel c shows average stiffness per
512 condition).
513
514
515   Our predictions regarding information differences were borne out with statistically

516   reliable group differences in the cumulative amount of workspace information at the

517   end of training (F (2, 42) = 20.06, p < .0001, η^2_p = .489; **Fig 6D**). The Random group

518   experienced more information relative to the Adaptive-Algorithm (p < .0001) and

519   Adaptive- Disturbance (p < .0001) conditions, but there was no difference between

520   the latter two groups (p = .806).

521

522 **Fig 6** - **Workspace information and surprise.** The stiffness coefficient K (N/m)
523 demonstrates the degree of assistance (positive values/error reduction) and disturbance
524 (negative values/error amplification) on a trial-by-trial basis for example subjects in the (a)
525 Adaptive Algorithm; (b) Adaptive Disturbance Algorithm and (c) Random conditions; (d) The
526 manipulation led to the Random group having more information over time; and (e) Heat maps
527 of the amount of information across the workspace provide a visualization of difference effect
528 for example participants, after the first and last training session.

529

530   From the curve fitting results, there were no reliable differences in task

531   difficulty level as indexed by individual values ($F_{(2, 42)} = 1.491$, $p = 0.2368$, $\eta^2_p =$

532   .066), but the groups did show differences in performance improvement across

533   training ($F_{(2, 42)} = 5.058$, $p = .0108$, $\eta^2_p = .194$). This effect was driven by the

534   Random group showing a steeper curve in training performance compared to the

535   Adaptive Algorithm ($p = .0112$), though it did not reach the statistical significance

536   threshold when compared against the Adaptive Disturbance Algorithm ($p = .0624$).

537   There were no differences between the Adaptive Algorithm and the Adaptive

538   Disturbance conditions ($p = .8613$).

539   We also found group differences in the amount of motor learning from pre- to

540   post-training with no stiffness intervention ($F_{(2, 42)} = 4.541$, $p = .0164$, $\eta^2_p = .178$;

541   **Fig 7B**). There was no statistically reliable difference in learning between the

542   Adaptive Algorithm and Adaptive-Disturbance Algorithm ($p = .914$). Instead, this

543   effect was driven by improvements following exposure to Random levels of

544   assistance/disruption relative to the Adaptive ($p = .018$) and Adaptive- Disturbance

545   algorithms ($p = .009$).

546

547   **Fig 7- Performance on training and learning generalization.** (a) Error reduction rates
548   during training. Abscissa represents block number; (b): Random levels of
549   assistance/disturbance demonstrated better learning, as indexed by the amount of error
550   reduction post training relative to pre-training in a novel workspace. Pre- and post- training
551   assessments are always performed without any stiffness intervention (k=0).

552

553   Finally, given our hypothesis that the amount of information predicts learning,

554   we reasoned that there should be a positive correlation between the amount of

555   information that participants are exposed to during training and the amount of

556   learning (i.e. difference in performance between pre- and post-training sessions).

557   Conducting correlation analyses at a condition-level would have been confounded by

558   our manipulations of information across training groups and would have relatively

559   weak statistical power to detect an underlying relationship (sample sizes varying from

560   13 to 17 per group). Thus, we pooled data across both experiments (n = 86) and

561   performed a simple linear regression to predict learning based on cumulative

562   information exposure during training. Consistent with our hypothesis, we found a

563   statistically significant relationship ($F_{(1, 82)}$ = 10.45, p = .0011), with the information

564   metric explaining 11.2% in variation in learning across all conditions (R^2 = 0.112;

565   **Table 1**; **Fig 8**).

566

567   **Fig 8 - Information exposure predicts learning.** Learning (mean path error reduction
568   between pre- and post- training) as a function of cumulative information acquired during
569   training (total entropy), for all participants in both experiments (R^2 = 0.122).

570

571   Recent evidence from Wu and colleagues [29] demonstrates that the intrinsic

572   movement variability associated with motor commands (from Z_n to Z_{n+1} to Z_{n+2} …)

573   predicts individual rates of motor learning. Indeed, it is possible that increased error

574   variability may be the mechanism by which information about the workspace is

575   acquired. To contextualise and compare the predictive value of the information metric

576   against a more parsimonious model of movement variability, we ran a second

577   regression analysis where we included the standard deviation of path error (per

578   component/sub-path; and averaged across training trials; **Table 1** Model 2).

579   Interestingly, we found that this measure of variability was unable to predict learning

580   in these data (p = .292, R^2 = 0.01) and a direct comparison between a two-parameter

581   model (Model 3; R^2 = 0.116) and Model 1 showed no statistically significant reliable

582   differences (p = .529).

583

584

585 **Table 1 - Information exposure predicts learning**

Model	t	p	β	F	df	p	mult. R^2	adj. R^2
Model 1				10.45	83	0.002	0.112	0.101
Cumulative information	3.232	0.002	2.31×10^{-4}					
Model 2				1.125	83	0.292	0.014	0.001
Path Error Mean SD	1.061	0.292	6.45×10^{-2}					
Model 3				5.34	82	0.006	0.116	0.094
Cumulative information	3.087	0.003	2.59×10^{-4}					
Path Error Mean SD	-0.633	0.529	-4.27×10^{-2}					

586

587

Table 1 - Information exposure predicts learning

24

Discussion

588

589   To date, there have been no principled explanations as to why motor learning can

590   be impaired by haptic assistance and facilitated by disturbance force application [9].

591   The current results support the hypothesis that the underlying mechanism relates to

592   the availability of information, and show that haptic forces that provide more 'surprise'

593   will lead to better learning in novel environments.

594   We created a model (inspired by the Free Energy Principle) to quantify the

595   amount of information available to learners during a task. Experiment 1 showed that

596   disturbance forces led to the accumulation of significantly more information across

597   the training period. These results aligned with our analysis of the amount of motor

598   learning following training, whereby the group that sampled more information showed

599   superior performance relative to a group provided with assistance and to an active-

600   control group. In Experiment 2, we demonstrated that the manipulation of information

601   (created by training individuals on a series of random assistive and disturbance

602   forces) yielded better learning compared to providing predictable levels of

603   assistance/ disturbance tuned to individual performance. It should be noted that the

604   results from Experiment 2 cannot be explained by the adoption of a short term global

605   impedance strategy (without much special pleading).

606   Our findings are consistent with previous results suggesting that disturbance

607   forces might be beneficial for motor learning [4–7]. Importantly, the current work

608   advances these reports by providing, and testing, a theoretical account of why

609   disturbance might accelerate learning. Specifically, we show that these results are

610   predicted by the free energy principle - which proposes that human learning can be

611   conceptualised as a process of free-energy minimization [14]. Here, motor learning is

612   seen as a process of entropy reduction where the average surprise of perceptual

613   outcomes sampled from a probability distribution relating to a motor command is

614   decreased through the development of forward models. The decrease in surprise

615 relates to improved inferences created by the system through exposure to

616 information that relates perceptual output to motor signal input. In line with this,

617 through pooling the data across both experiments, we found that the amount of

618 workspace information participants were exposed to during training could predict a

619 statistically significant amount of variance in learning. Given the plethora of variables

620 that could also have influenced learning across these different manipulations (six

621 experimental conditions in two experiments), it is notable that this relationship

622 between information and learning could be detected.

623 Moreover, we provide evidence that the improved information sampling created by

624 disturbance enables generalisation rather than simple performance facilitation [1,30].

625 Our work thus complements and advances previous observations about the potential

626 benefits of disturbance. For example, an earlier study showed that performance on a

627 tracking task could be improved through delivery of haptic disturbance [5]. This finding

628 could be explained, however, by the participants being trained to become more

629 proficient in deploying feedback control and, indeed, the authors of the study explained

630 their results in terms of a general training improvement in the 'attentional' capabilities

631 of their participants. The problem with such explanations relates to the difficulty in

632 defining and quantifying the term 'attention' when used in this manner. It is therefore

633 interesting to note that the improved tracking performance is predicted within the FEP

634 framework. The presence of haptic disturbance when tracking will generate surprise

635 and thus force the system to act to reduce the entropy (i.e. learn to make effective

636 feedback corrections). Indeed, the Random training condition in our experiment

637 exploited this mechanism in a principled manner by exposing participants to frequent

638 movement-by-movement switches between positive and negative stiffness. Together,

639 these results illustrate the fundamental links between attention and uncertainty (see

640 [31,32]), and suggest that the effects of haptic disturbance can be quantified in a range

641 of different settings through information theory.

642   Our results also build on previous work showing a relationship between variability

643   and motor learning. For example, Van Beers [33] showed that the random effects of

644   planning noise accumulate, in contrast to task-relevant errors which show close to

645   zero accumulation (explained by effective trial-by-trial corrections), whilst Wu et al's

646   experiments [29] (results described earlier), have shown that task-relevant motor

647   variability facilitates faster learning rates. On these grounds, it has been argued that

648   intrinsic movement variability leads to motor exploration, which sub-serves motor

649   learning and performance optimization. Indeed, the idea that action exploration can

650   drive learning has long been mooted in theories of operant behaviour [34] and

651   human development [35–37]. Recent experiments have shown that (a) artificially

652   manipulating the relationship between movements and visuomotor noise can be

653   used to teach people specific control policies [38] and (b) the variability in task-

654   redundant parameters can predict motor adaptation rates [39]. The current findings

655   demonstrate that extrinsic variability delivered through haptic disturbance can, in the

656   same vein, augment learning by increasing the amount of information sampled by the

657   learner. The general notion that increased exposure to information can lead to faster

658   learning is well explained by theories of structural learning and has good support

659   from a range of empirical studies [18,19,40–43] including investigations of

660   laparoscopic surgical training [44]. Our extension to these ideas is that learning of the

661   structure can be directly related to the amount of information available to the learner.

662   Indeed, regression analyses for our data shows that the amount of information

663   accumulated over training (as indexed by our model) provided greater explanatory

664   power compared to a measure of motor variability alone in this task.

665   These findings raise the issue of which neural substrates underpin these learning

666   processes. The neural processes that implement the computational algorithms

667   exploited by the human nervous system remain to be discovered [45,46]. Likewise,

668   the underlying control mechanisms supporting skilled arm movements are poorly

669   understood and, as such, it is difficult to speculate on how the individuals learned to

670   compensate for the complex force field, but we suggest that the learning was likely to

671   involve processes related to optimal feedback control as well as predictive

672   mechanisms [47–49].

673   Our findings suggest that the participants developed forward or inverse models

674   that allowed them to predict (and thus compensate for) the novel force field through

675   which they needed to move. It has been shown previously that participants can learn

676   a short term strategy of stiffening their arm to resist the effects of sudden unexpected

677   force perturbations [24,25]. This work has demonstrated that humans learn to use

678   selective control of impedance geometry in order to stabilise unstable dynamics in a

679   skilful and energy efficient manner. It is probable that participants in the current

680   experiments adopted such a strategy when they were first exposed to the novel

681   workspace (as they were unable to predict the forces that were applied as they

682   moved through the space). Importantly, there was a regular (lawful) structure to the

683   novel workspace, in the same way that the world provides a lawful force field through

684   which the neonate must learn to move their arm. We hypothesised that the system

685   would learn the underlying force field so that the arm could move skilfully through the

686   workspace rather than repeatedly contend with unexpected displacement. This

687   hypothesis was based on the free energy minimization principle which suggests

688   human behaviour is marked by continual attempts to reduce entropy (i.e. minimise

689   surprise). Experiment 2 allowed us to test whether participants were learning the

690   force field or adopting a global impedance strategy, by which the arm is stiffened in

691   all directions to counteract external force interventions. As outlined above and

692   demonstrated in previous research, participants are likely to adopt a global

693   impedance strategy when the force intervention is largely disruptive and increases

694   error ($k < 0$). However, in Experiment 2, the random condition consisted of (on

695   average) 50% assistive trials, whereby the force intervention *assisted* movement,

28

696   thus rendering such a strategy sub-optimal. We reasoned that, in contrast to the

697   random forces, the adaptive disturbance algorithm, where participants were provided

698   with a more consistent presentation of disturbance forces would be more likely to

699   adopt an impedance control strategy. Given that we observed improved learning in

700   the random condition, impedance control is unlikely to provide a full account of these

701   data. Instead, these results indicate that participants were learning to skilfully

702   counteract the underlying workspace force field and we propose that this learning

703   was promoted, in part, through the increased information acquired during training.

704   Finally, it is important to note that this study used neurologically intact adults as

705   participants and whilst the force field in the two experiments allowed us to examine

706   novel skill learning, the difficulty was tuned to a level such that all participants could

707   complete the task. We speculate that disrupting the training of individuals with

708   neurological deficits (e.g. cerebral palsy) might not be beneficial, and constraining

709   errors in these populations could speed up learning by helping the individuals sample

710   the necessary information [21]. Consistent with this, there is work with stroke

711   survivors that has shown that error amplification is useful in rehabilitation for mild

712   impairment, but error guidance is necessary for patients with more severe damage

713   [50]. Likewise, haptic guidance has been found to be beneficial for people with

714   relatively low skill levels, but error enhancement is better for highly skilled individuals

715   [3,51]. The current work builds on these observations and provides a theoretical

716   framework for the development of optimized robotic training devices in skill training

717   and rehabilitation.

Author Contributions

719   E.J. & J.B. conducted the observations. E.J. and J.B. reduced the data. F.M., J.B,

720   and M.M-W wrote the manuscript. All authors discussed the results and implications

721   and commented on the manuscript at all stages.

722   **Data Availability**

723   Data files used in analyses are openly available on the Open Science Framework:

724   https://osf.io/7c95b/

725   **Acknowledgements**

References

1. Reinkensmeyer DJ, Patton JL. Can robots help the learning of skilled actions? Exerc Sport Sci Rev. 2009;37: 43–51. doi:10.1097/JES.0b013e3181912108

2. Marchal-Crespo L, Schneider J, Jaeger L, Riener R. Learning a locomotor task: with or without errors? J Neuroeng Rehabil. 2014;11: 25. doi:10.1186/1743-0003-11-25

3. Sigrist R, Rauter G, Riener R, Wolf P. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review. Psychon Bull Rev. 2013;20: 21–53. doi:10.3758/s13423-012-0333-8

4. Cesqui B, Aliboni S, Mazzoleni S, Carrozza MC, Posteraro F, Micera S. On the use of divergent force fields in robot-mediated neurorehabilitation. Proc 2nd Bienn IEEE/RAS-EMBS Int Conf Biomed Robot Biomechatronics, BioRob 2008. IEEE; 2008; 854–861. doi:10.1109/BIOROB.2008.4762927

5. Lee J, Choi S. Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance. 2010 IEEE Haptics Symp HAPTICS 2010. 2010; 335–342. doi:10.1109/HAPTIC.2010.5444635

6. Huang VS, Shadmehr R. Evolution of motor memory during the seconds after observation of motor error. J Neurophysiol. 2007;97: 3976–85. doi:10.1152/jn.01281.2006

7. Emken JL, Reinkensmeyer DJ. Robot-enhanced motor learning: Accelerating internal model formation during locomotion by transient dynamic amplification. IEEE Trans Neural Syst Rehabil Eng. 2005;13: 33–39. doi:10.1109/TNSRE.2004.843173

8. Hasson CJ, Zhang Z, Abe MO, Sternad D. Neuromotor Noise Is Malleable by Amplifying Perceived Errors. PLoS Comput Biol. 2016;12: 1–28. doi:10.1371/journal.pcbi.1005044

9. Heuer H, Lüttgen J. Robot assistance of motor learning: A neuro-cognitive perspective. Neurosci Biobehav Rev. Elsevier Ltd; 2015;56: 222–240. doi:10.1016/j.neubiorev.2015.07.005

10. Hesse S, Schmidt H, Werner C, Bardeleben A. Upper and lower extremity robotic devices for rehabilitation and for studying motor control. Curr Opin Neurol. 2003;16: 705–710. doi:10.1097/00019052-200312000-00010

11. Shannon CE. A mathematical theory of communication. Bell Syst Tech J.

1948;27: 379–423. doi:10.1145/584091.584093

12. Shannon WE. The mathematical theory of communication. the star online. Univ of Illinois Press; 1949.

13. Friston K. Learning and inference in the brain. Neural Networks. 2003;16: 1325–1352. doi:10.1016/j.neunet.2003.06.005

14. Friston K. The free-energy principle: a rough guide to the brain? Trends Cogn Sci. 2009;13: 293–301. doi:10.1016/j.tics.2009.04.005

15. Friston K. The free-energy principle: a unified brain theory? Nat Rev Neurosci. 2010;11: 127–138. doi:10.1038/nrn2787

16. Friston KJ, Daunizeau J, Kilner J, Kiebel SJ, Friston KJ, Daunizeau J, et al. Action and behavior: A free-energy formulation. Biol Cybern. 2010;102: 227–260. doi:10.1007/s00422-010-0364-z

17. Miall RC, Wolpert DM. Forward models for physiological motor control. Neural Networks. 1996;9: 1265–1279. doi:Doi 10.1016/S0893-6080(96)00035-4

18. Braun DA, Mehring C, Wolpert DM. Structure learning in action. Behav Brain Res. 2010;206: 157–65. doi:10.1016/j.bbr.2009.08.031

19. Braun DA, Waldert S, Aertsen A, Wolpert DM, Mehring C. Structure Learning in a Sensorimotor Association Task. PLoS One. 2010;5. doi:Artn E8973Doi 10.1371/Journal.Pone.0008973

20. Bays PM, Wolpert DM. Computational principles of sensorimotor control that minimize uncertainty and variability. J Physiol. 2007;578: 387–96. doi:10.1113/jphysiol.2006.120121

21. Snapp-Childs W, Mon-Williams M, Bingham GP. A sensorimotor approach to the training of manual actions in children with developmental coordination disorder. J Child Neurol. 2013;28: 204–12. doi:10.1177/0883073812461945

22. Lum PS, Burgar CG, Shor PC, Majmundar M, Van der Loos M. Robot-assisted movement training compared with conventional therapy techniques for the rehabilitation of upper-limb motor function after stroke. Arch Phys Med Rehabil. 2002;83: 952–959. doi:10.1053/apmr.2001.33101

23. Bouchard AE, Corriveau H, Milot M-H. Comparison of haptic guidance and error amplification robotic trainings for the learning of a timing-based motor task by healthy seniors. Front Syst Neurosci. 2015;9: 52. doi:10.3389/fnsys.2015.00052
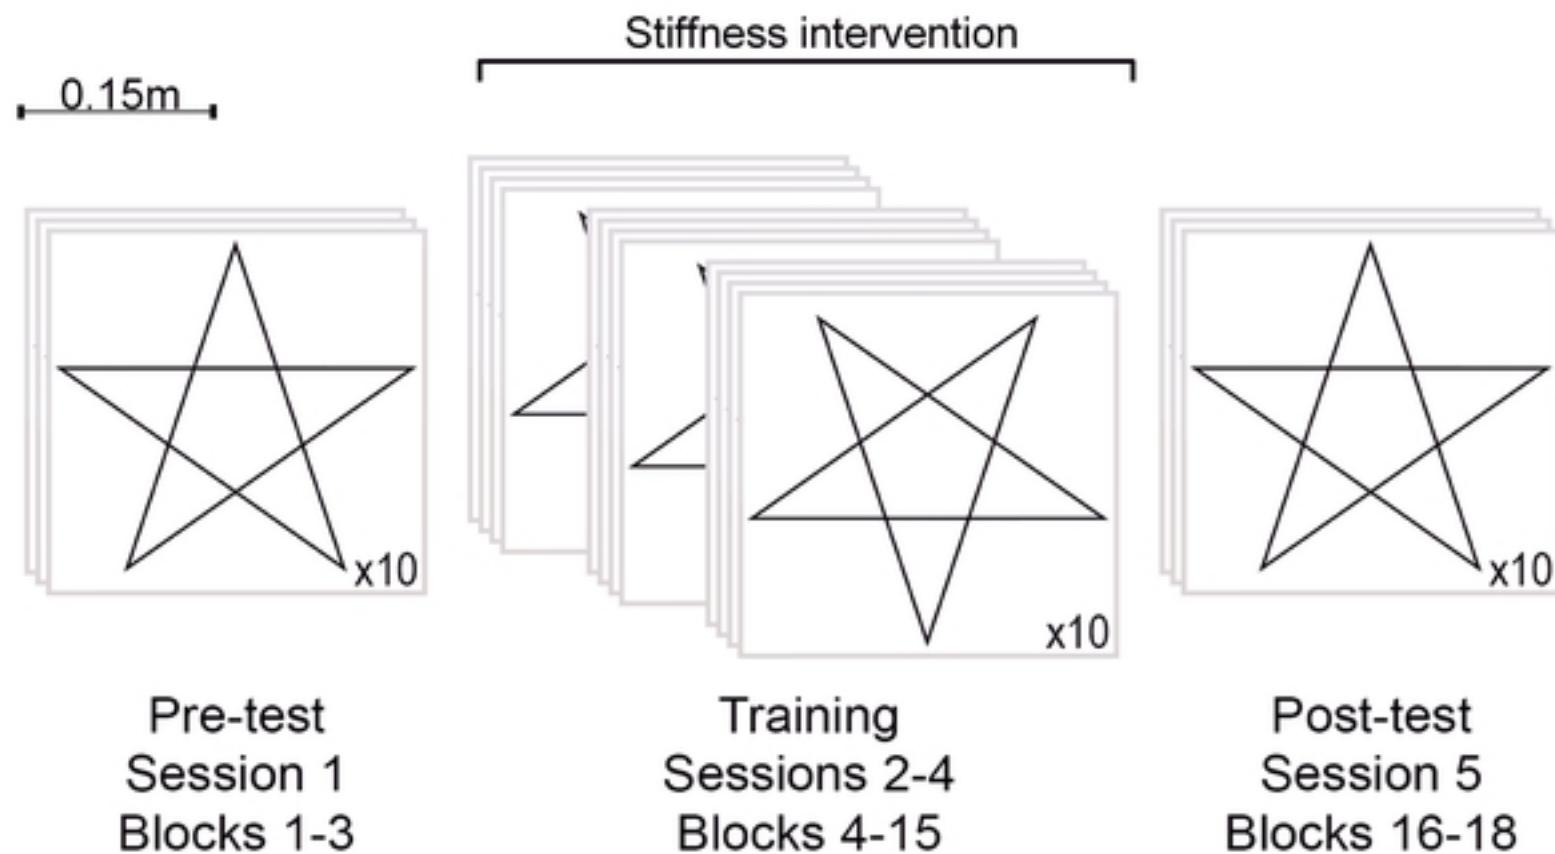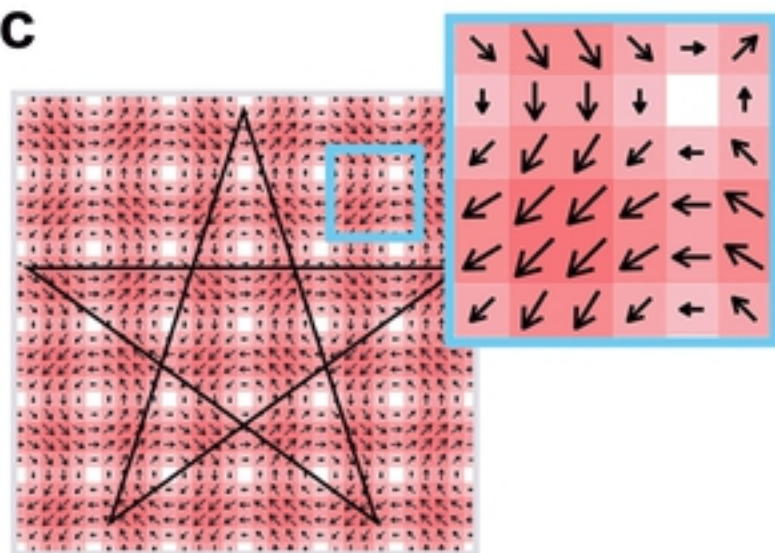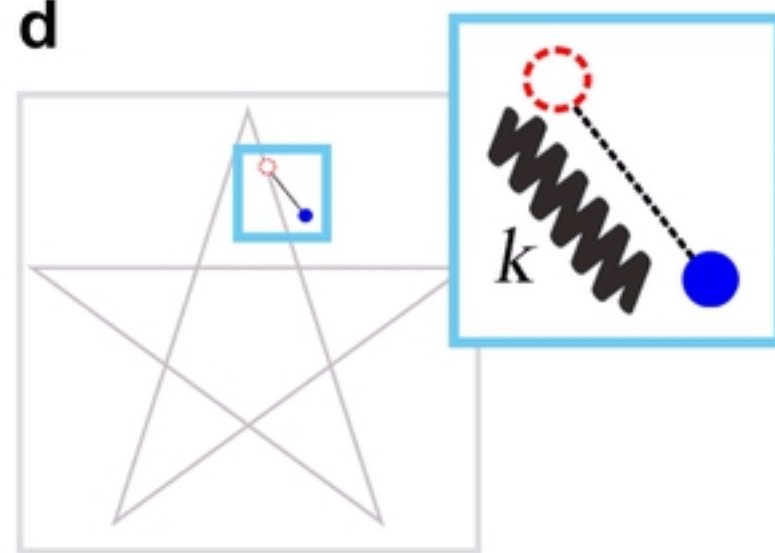
32

24.   Burdet E, Tee KP, Mareels I, Milner TE, Chew CM, Franklin DW, et al. Stability and motor adaptation in human arm movements. Biol Cybern. 2006;94: 20–32. doi:10.1007/s00422-005-0025-9

25.   Burdet E, Osu R, Franklin DW, Milner TE, Kawato M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. Nature. 2001;414: 446–449. doi:10.1038/35106566

26.   Flatters I, Mushtaq F, Hill LJB, Holt RJ, Mon-Williams M. The relationship between postural stability and manual control in children. Exp brain Res. 2014;232: 2907–2917.

27.   Linde R Van Der, Lammertse P, Frederiksen E, Ruiter B, B. The HapticMaster, a new high-performance haptic interface. Proc EuroHaptic, Edinburgh, UK. 2002; 1–5. doi:10.1.1.2.7408

28.   Levene H. Robust testes for equality of variances. In: Olkin I, editor. Contributions to Probability and Statistics. Palo Alto Calif: Stanford University Press; 1960. pp. 278–202.

29.   Wu HG, Miyamoto YR, Castro LNG, Ölveczky BP, Smith MA. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. Nat Neurosci. 2014;17. doi:10.1038/nn.3616

30.   Reisman DS, Wityk R, Silver K, Bastian AJ. Locomotor adaptation on a split-belt treadmill can improve walking symmetry post-stroke. Brain. 2007;130: 1861–1872.

31.   Feldman H, Friston K. Attention, uncertainty, and free-energy. Front Hum Neurosci. 2010;4: 215. doi:10.3389/fnhum.2010.00215

32.   Mushtaq F, Bland AR, Schaefer A. Uncertainty and cognitive control. Front Psychol. 2011;2: 249. doi:10.3389/fpsyg.2011.00249

33.   van Beers RJ. Motor Learning Is Optimally Tuned to the Properties of Motor Noise. Neuron. 2009;63: 406–417. doi:10.1016/j.neuron.2009.06.025

34.   Barto AG, Sutton RS, Barto AG. Reinforcement learning: An introduction. IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council. Cambridge, MA: MIT Press; 1998. doi:10.1109/TNN.1998.712192

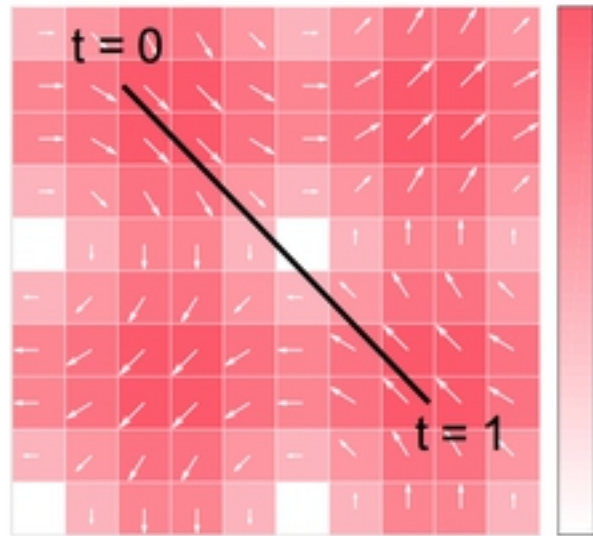35.   Gibson E. Exploratory Behavior In The Development Of Perceiving, Acting, And The Acquiring Of Knowledge. Annu Rev Psychol. 1988;39: 1–41.
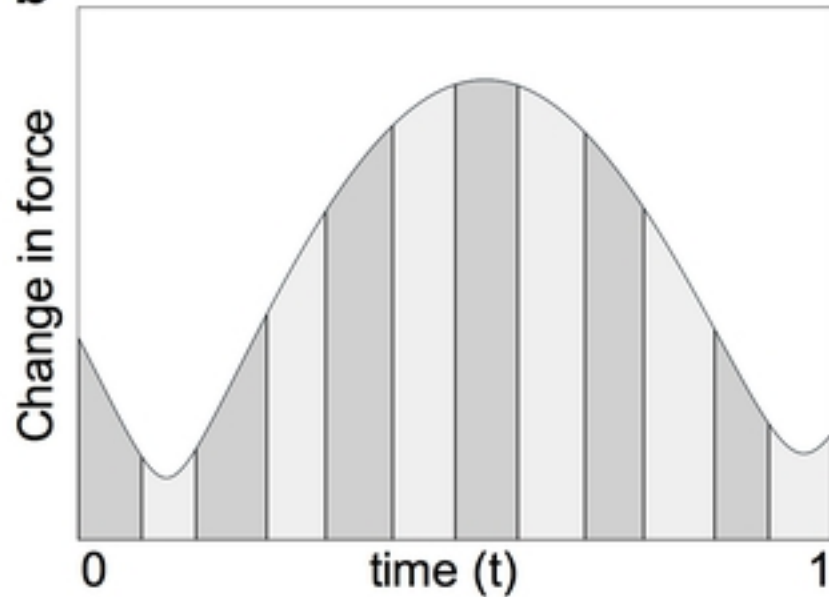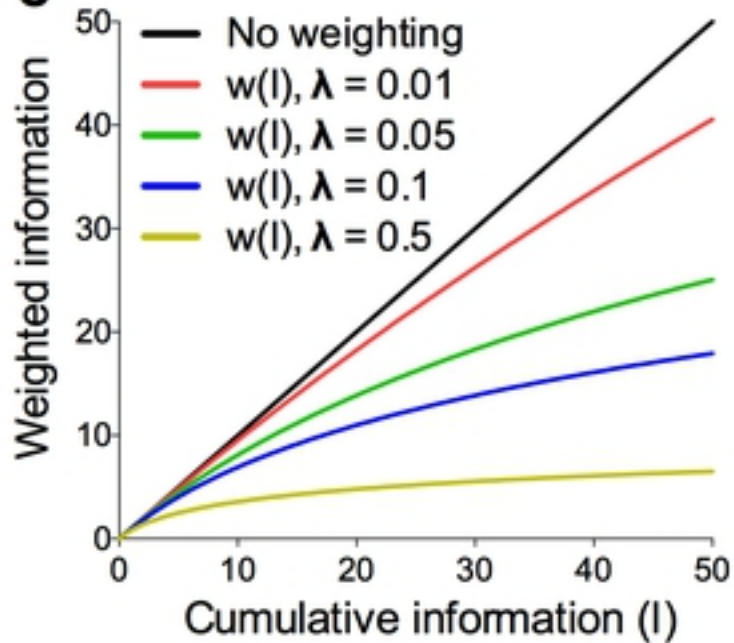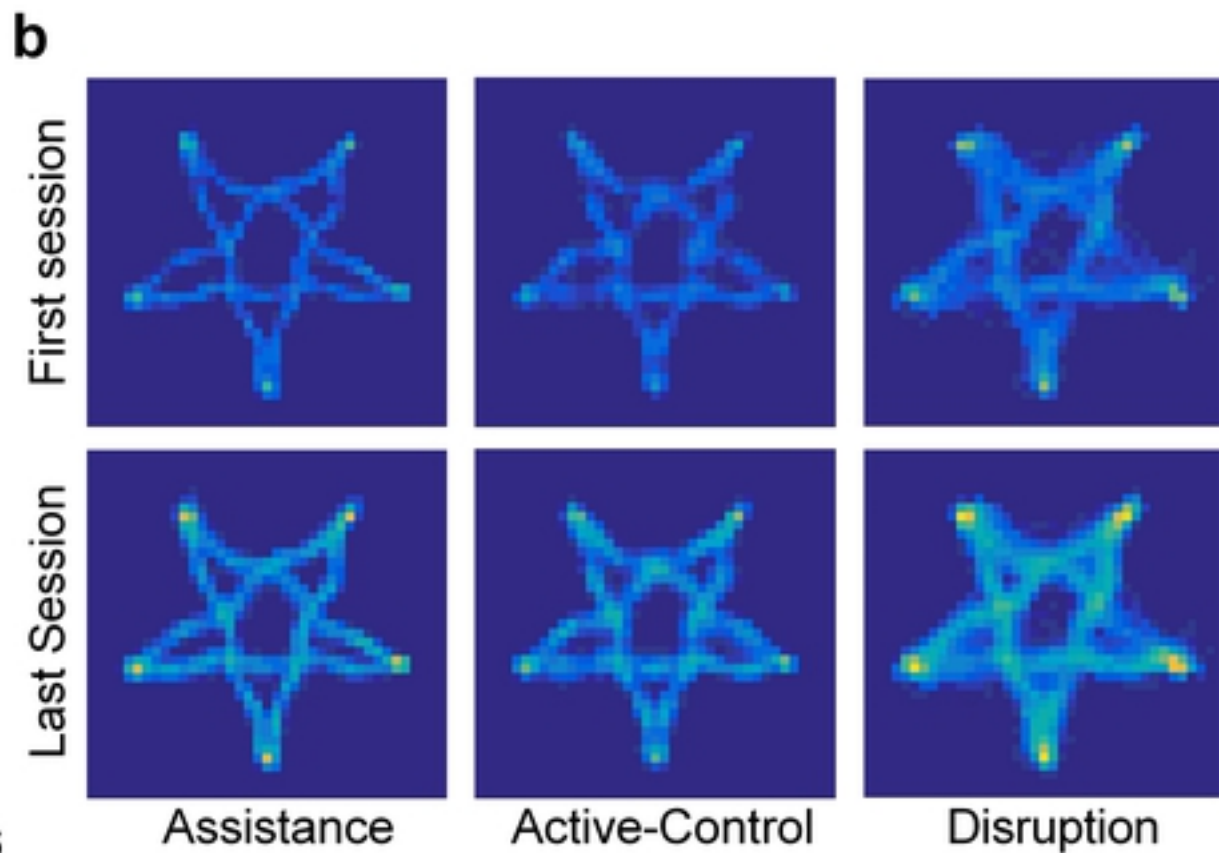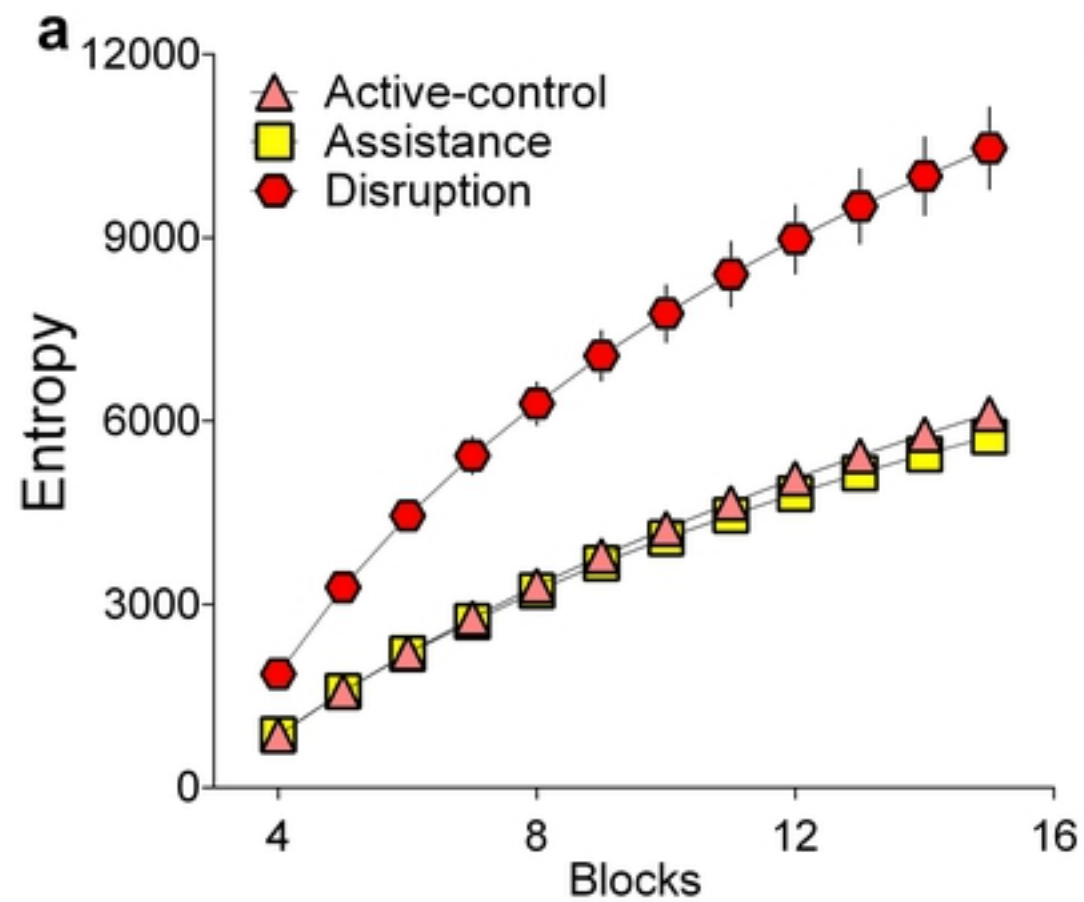
33

doi:10.1146/annurev.psych.39.1.1

36. Thelen E. The (re)discovery of motor development: Learning new things from an old field. Dev Psychol. 1989;25: 946–949. doi:10.1037/0012-1649.25.6.946

37. Bruner JS. Organization of early skilled action. Child Dev. 1973;44: 1–11. doi:10.2307/1127671

38. Thorp EB, Kording KP, Mussa-Ivaldi FA. Using noise to shape motor learning. J Neurophysiol. 2017;117: 728–737. doi:10.1152/jn.00493.2016

39. Singh P, Jana S, Ghosal A, Murthy A. Exploration of joint redundancy but not task space variability facilitates supervised motor learning. Proc Natl Acad Sci U S A. 2016; 201613383. doi:10.1073/pnas.1613383113

40. Turnham EJA, Braun DA, Wolpert DM. Inferring Visuomotor Priors for Sensorimotor Learning. Plos Comput Biol. 2011;7. doi:ARTN e1001112DOI 10.1371/journal.pcbi.1001112

41. Yousif N, Diedrichsen J. Structural learning in feedforward and feedback control. J Neurophysiol. 2012/08/17. 2012;108: 2373–2382. doi:10.1152/jn.00315.2012

42. Johnson RL, Culmer PR, Burke MR, Mon-Williams M, Wilkie RM. Exploring structural learning in handwriting. Exp brain Res. 2010/10/26. 2010;207: 291–5. doi:10.1007/s00221-010-2438-5

43. Braun DA, Aertsen A, Wolpert DM, Mehring C. Motor Task Variation Induces Structural Learning. Curr Biol. Elsevier Ltd; 2009;19: 352–357. doi:10.1016/j.cub.2009.01.036

44. White AD, Giles O, Sutherland RJ, Ziff O, Mon-Williams M, Wilkie RM, et al. Minimally invasive surgery training using multiple port sites to improve performance. Surg Endosc. 2013; 1–6. doi:10.1007/s00464-013-3307-7

45. Wolpert DM, Ghahramani Z, Flanagan JR. Perspectives and problems in motor learning. Trends Cogn Sci. 2001;5: 487–494. doi:Doi 10.1016/S1364-6613(00)01773-3

46. Wolpert DM. Computational approaches to motor control. Trends Cogn Sci. 1997;1: 209–216. doi:Pii S1364-6613(97)01070-XDoi 10.1016/S1364-6613(97)01070-X

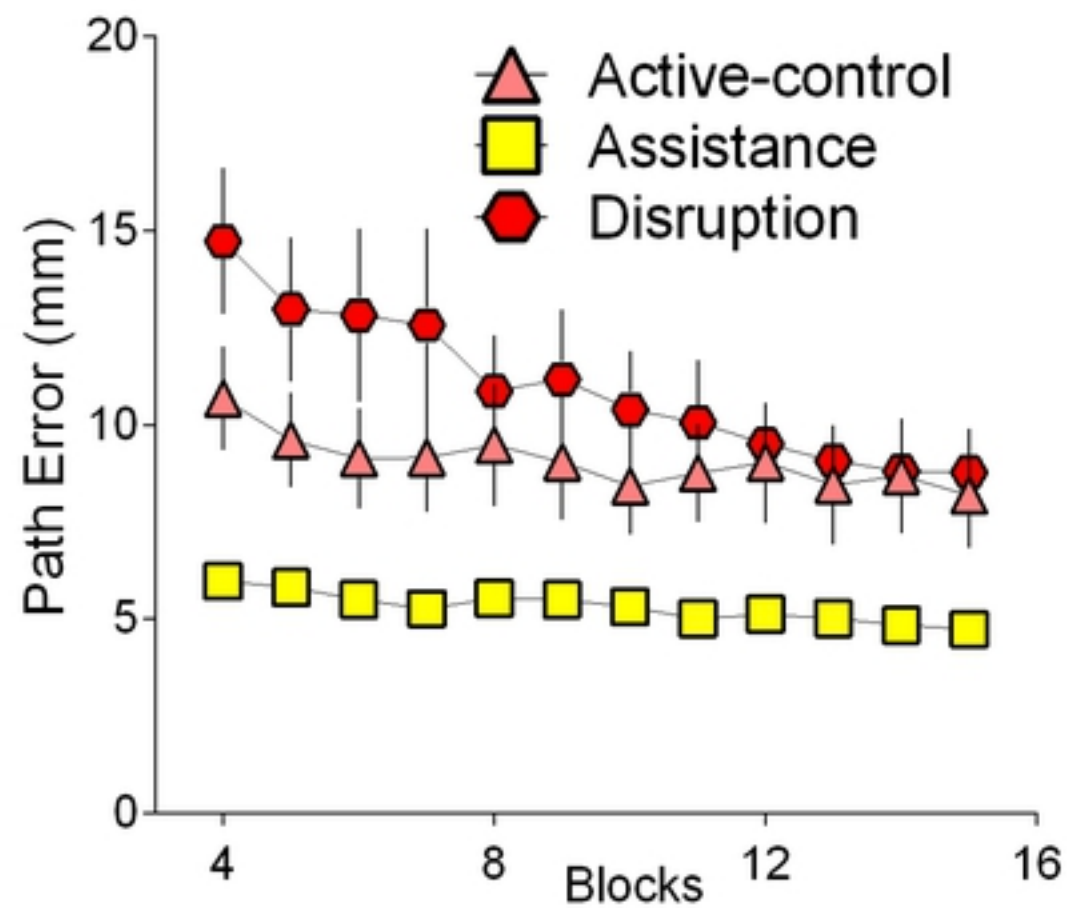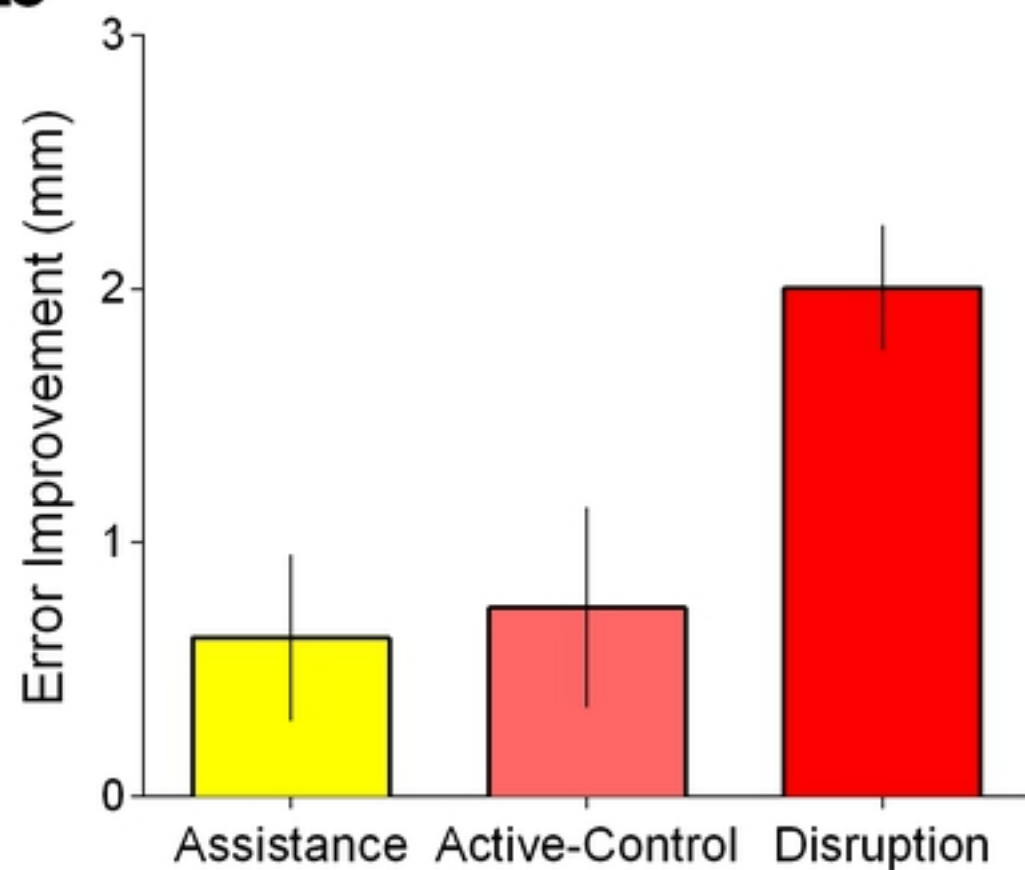47. Todorov E. Optimality principles in sensorimotor control. Nat Neurosci. 2004;7: 907–15. doi:10.1038/nn1309

48.   Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. Nat Neurosci. 2002;5: 1226–35. doi:10.1038/nn963

49.   Franklin DW, Wolpert DM. Computational Mechanisms of Sensorimotor Control. Neuron. 2011;72: 425–442. doi:DOI 10.1016/j.neuron.2011.10.006

50.   Marchal-Crespo L, Reinkensmeyer DJ. Review of control strategies for robotic movement training after neurologic injury. J Neuroeng Rehabil. 2009;6: 20. doi:10.1186/1743-0003-6-20

51.   Milot MH, Marchal-Crespo L, Green CS, Cramer SC, Reinkensmeyer DJ. Comparison of error-amplification and haptic-guidance training techniques for learning of a timing-based motor task by healthy individuals. Exp Brain Res. 2010;201: 119–131. doi:10.1007/s00221-009-2014-z
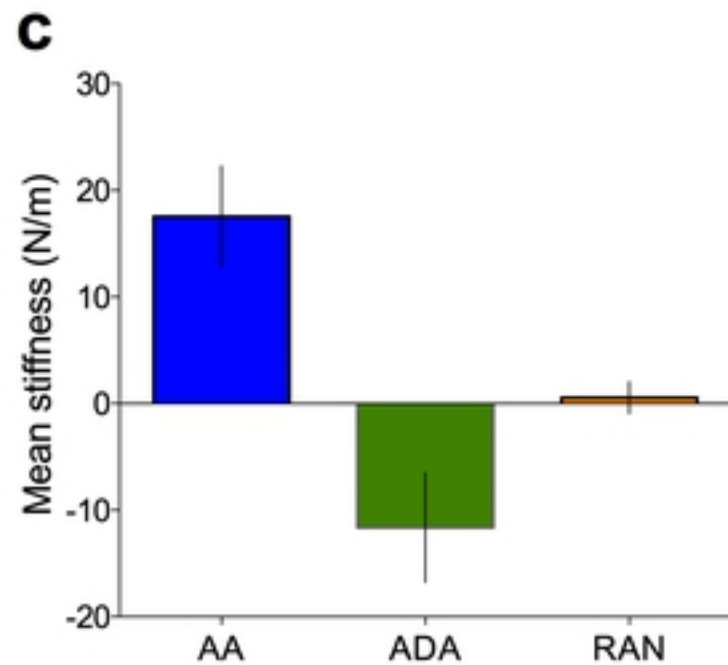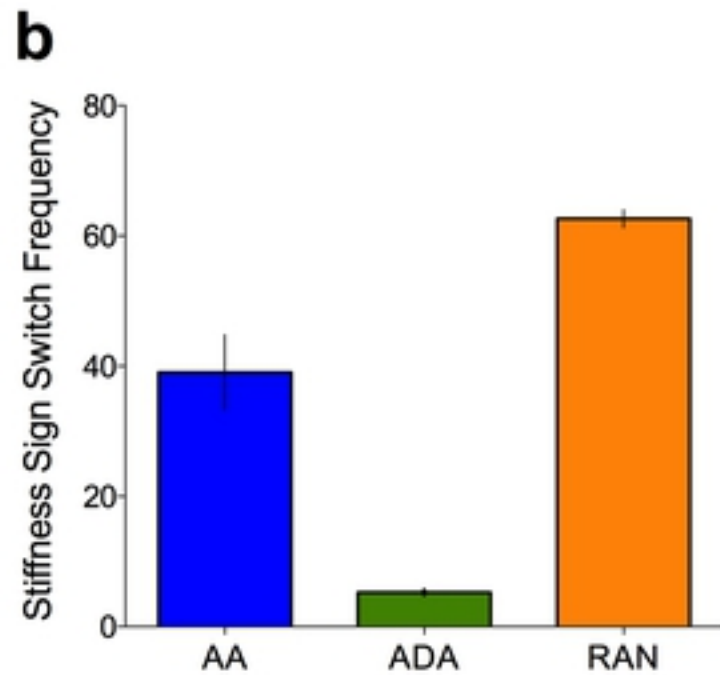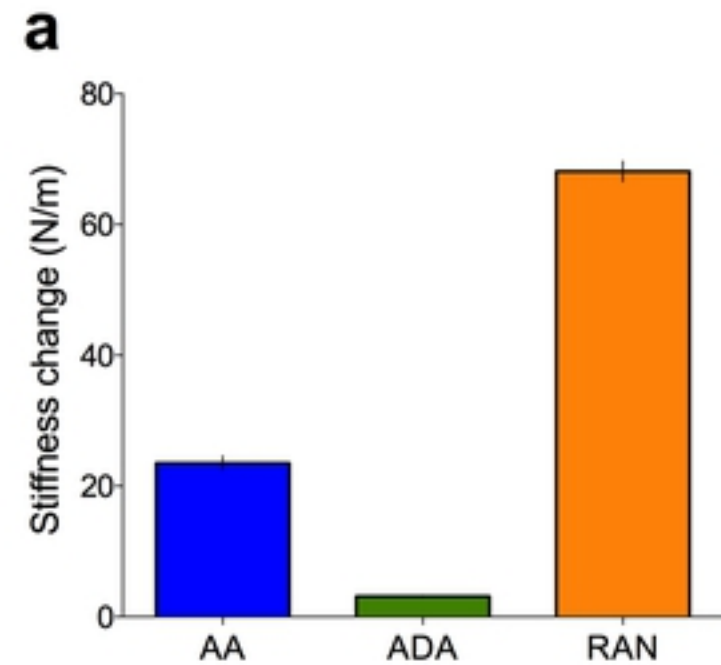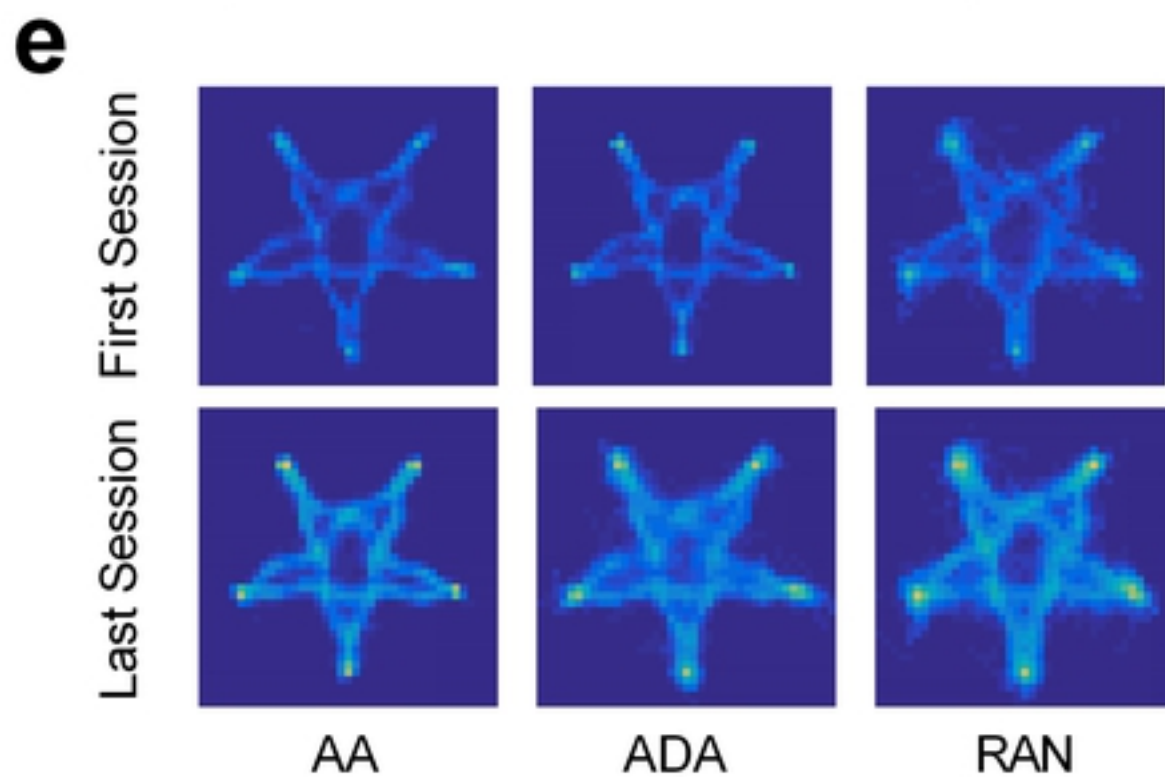
Y = 0.0002454*X - 0.3749

a

b

Stiffness intervention

0.15m

×10

×10

×10

Pre-test
Session 1
Blocks 1-3

Training
Sessions 2-4
Blocks 4-15

Post-test
Session 5
Blocks 16-18

c

d

k

a

t = 0

t = 1

b

Change in force

0   time (t)   1

c

Weighted information

- No weighting
- $w(l)$, $\lambda = 0.01$
- $w(l)$, $\lambda = 0.05$
- $w(l)$, $\lambda = 0.1$
- $w(l)$, $\lambda = 0.5$

Cumulative information (l)

a

Entropy vs Blocks

- ▲ Active-control
- ■ Assistance
- ⬡ Disruption

b

	Assistance	Active-Control	Disruption
First session			
Last Session			

a

Stiffness Coefficient K (N/m)

Trial

b

Trial

c

Trial

d

Entropy

■ Adaptive Algorithm
● Adaptive-Disruptive Algorithm
◆ Random

Block

e

First Session

Last Session

AA   ADA   RAN

a

Path Error (mm) vs Block

- ■ Adaptive Algorithm
- ● Adaptive-Disruptive Algorithm
- ◆ Random

b

Error Improvement (mm) for AA, ADA, RAN