HUMAN-ROBOT INTERACTION

When a robot teaches humans: Automated feedback selection accelerates motor learning

Georg Rauter[1,2]*, Nicolas Gerig[1,2], Roland Sigrist[1], Robert Riener[1,3], Peter Wolf[1]

A multitude of robotic systems have been developed to foster motor learning. Some of these robotic systems featured augmented visual or haptic feedback, which was automatically adjusted to the trainee's performance. However, selecting the type of feedback to achieve the training goal usually remained up to a human trainer. We automated this feedback selection within a robotic rowing simulator: Four spatial errors and one velocity error were considered, all related to trunk-arm sweep rowing set as the training goal to be learned. In an alternating sequence of assessments without augmented feedback and training sessions with augmented, concurrent feedback, the experimental group received feedback, thus addressing the main shortcoming of the previous assessment. With this approach, each participant of the experimental group received an individual sequence of 10 training sessions with feedback. The training sequences from participants in the experimental group were consecutively applied for participants in the control group. Both groups were able to reduce spatial and velocity errors due to training. The learning rate of the requested velocity profile was significantly higher for the experimental group compared with the control group. Thus, our robotic rowing simulator accelerated motor learning by automated feedback selection. This demonstration of a working, closed-loop selection of types of feedback, i.e., training conditions, could serve as the basis for other robotic trainers incorporating further human expertise and artificial intelligence.

INTRODUCTION

Robot-assisted training in combination with virtual reality is becoming increasingly popular for facilitating motor learning. Tremendous efforts have been undertaken to (i) develop controllable and appealing training conditions [for examples, see overviews in the field of surgery by Bartlett et al. (1) and in the field of sport by Neumann et al. (2)], (ii) design augmented feedback to foster active participation and guide the trainee [for an overview, see Sigrist et al. (3)], and (iii) comprehensively monitor task performance by sensors of the robot or additional wearables [e.g., summarized by Kos et al. (4)].

Commonly, a human trainer defines a training goal; chooses a strategy for achieving this goal; selects training conditions, including type of feedback; and interprets available data to switch between conditions or change strategies and goals. In particular, if many training conditions and a huge amount of performance data are available, a robotic trainer may assist the human trainer to select the most appropriate training condition in a reasonable and reliable manner.

An automated trainer is a closed-loop controller that incorporates the trainee in the control loop and selects a training condition based on the trainee's performance. An automated trainer is termed a robotic trainer if haptic interaction with the trainee is present. A training condition might incorporate a virtual trainer, i.e., an avatar that demonstrates the target movement, sometimes displayed simultaneously with an avatar of the trainee [e.g. (5–7)]. In robot-assisted training, parameters of haptic guidance have been automatically adjusted in relation to the current task performance [e.g. (8–10)]. Thus, adaptations within a training condition have been proposed; however, another training condition featuring another feedback design might be more effective. Switching to another feedback design to further improve seems appropriate because the effectiveness of a feedback design depends on the task characteristics to be learned, individual preferences, and current skills

(3, 11, 12). To reasonably switch to another feedback design, an automated trainer has to monitor the deviation of the current performance to a defined training goal, e.g., movement error. In addition, an algorithm that implements the selection of the goal will be required.

Recently, we have presented a robotic trainer for a typical rowing movement: During trunk-arm sweep rowing, participants had to learn a predefined oar blade trajectory. The spatial deviation from the trajectory and the velocity error were monitored. To select the next training condition, we predicted error reductions due to different types of feedback by a linear mixed model. This model was established on previous evaluations of the available types of feedback. Observed error reductions were used to adjust the model to the individual participant. Our predictive concept demonstrated a reasonable selection of training conditions featuring different types of feedback (13). However, a considerable amount of data was required to develop and justify the underlying model. To avoid such modeling, an automated trainer could be based, to a further extent, on the expertise of a human trainer. For instance, an automated yoga trainer that prioritized errors based on the location of the issue was developed. This automated trainer provided verbal instructions on how to improve, following the suggestions of yoga experts (14). Human expertise has also been incorporated in other concepts of an automated trainer, e.g., for squats (15). However, none of these concepts were evaluated in terms of motor learning.

We adopted the approach of incorporating human expertise in a robotic trainer by addressing the largest error out of different performance metrics by a dedicated type of augmented feedback. The movement task was to learn an oar blade trajectory in trunk-arm sweep rowing, as in our previous study (13). The oar had to be moved according to a predefined velocity profile, i.e., fast in the water and slow in the air phase. Deviations to the velocity profile were considered as one error metric. In addition, four spatial errors that represent deviations from the reference path in the four distinct phases of an oar blade movement in rowing were considered: immersion into the water, movement through the water, movement out of the water, and movement through the air. Each of the five errors was linked to one type of unimodal concurrent feedback (either auditory or visual cues) and one type of multimodal

[1]Sensory-Motor Systems Lab, Department of Health Sciences and Technology, ETH Zurich, Zurich, Switzerland. [2]BIROMED-Lab, Department of Biomedical Engineering, University of Basel, Basel, Switzerland. [3]Medical Faculty, University of Zurich, Zurich, Switzerland.
*Corresponding author. Email: georg.rauter@unibas.ch

concurrent feedback (haptic guidance was added to the unimodal type) that had previously been shown to facilitate error reduction (*16*, *17*). The errors were calculated for non-feedback test conditions and normalized by the mean errors observed in the retention tests of former studies on the same task. The robotic trainer selected the feedback that addressed the largest normalized (i.e., dominant) error or switched to another feedback design if the dominant error stayed the same. Thus, the robotic trainer provided an individualized feedback sequence. We hypothesized that the participants receiving this automatically individualized feedback sequence would learn faster than those receiving the feedback sequence of matched participants. If confirmed, this would demonstrate that, given a training goal (such as accurately tracking an oar movement) whose underlying metrics are continuously evaluable (such as spatial and velocity error), a closed-loop controller in combination with a simple training strategy (such as addressing the largest error and a cascade of feedback designs that address different errors and provide cues in additional modalities when errors are repeated) can significantly contribute to learning of complex motor tasks (such as trunk-arm sweep rowing).

Fig. 1. Normalized errors over the course of the study for one participant of the experimental group and his matched control. Training blocks (T) are indicated by dark gray background. Baseline (BL) and retention tests at the beginning of day 2 (RE2) and day 3 (RE3) are indicated in white background. No-feedback trials (NF) are indicated in light gray background. BL, RE2, and NF were used as test conditions to assess the spatial error during catch, drive, release, recovery (see Fig. 4), and the velocity error. Triangle face color represents the dominant error based on which feedback was provided.

RESULTS

Baseline differences

There were no significant differences between the groups at baseline for any variable.

Dominant error occurrences

Participants of the experimental group received an individualized sequence of feedback designs based on their dominant error as assessed immediately before each training session, whereas the control participants received the feedback sequence of their matched participant in the experimental group. In consequence, participants in the control group could receive feedback for their dominant error in single trainings if, by chance, they had the same dominant error as the participant they were matched to (see Fig. 1). To simplify the understanding and interpretation of our results, Fig. 1 shows learning curves for one participant of the experimental group and his matched control participant.

The experimental group always trained with a feedback design according to their dominant error (8 participants × 10 training conditions = 80 times the desired feedback design over 2 days). In contrast, participants of the control group received feedback designs according to their dominant error in 36 of 80 training conditions (for details, see Table 1).

Learning results from baseline to retention tests

No learning could be observed for the outcome spatial error at the catch phase. Both groups showed significant learning for the outcome spatial error at r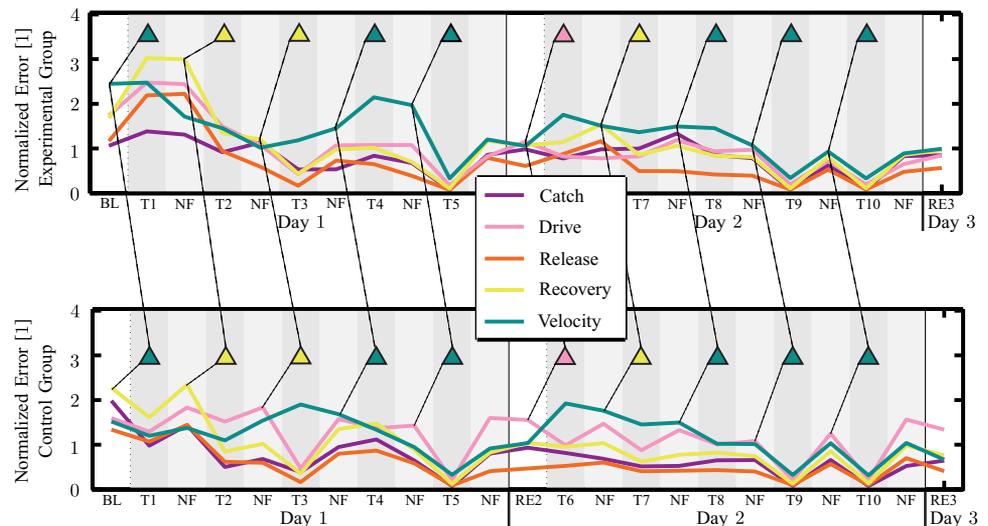elease and velocity error over the entire cycle. For the remaining outcomes—spatial error drive, spatial error recovery, and spatial error entire cycle—learning was only observed within the experimental group (Fig. 2). The absence of learning from the control group in these outcomes should not be interpreted as an indication of group differences, because the repeated-measures analysis of variance (ANOVA) did not reveal significant group effects or significant group test interactions for these outcomes (Table 2). However, velocity error entire cycle showed a trend ($P < 0.1$) toward a significant group test interaction.

A significant group difference ($F_{1,14} = 5.275$, $P = 0.038$, $\eta_p^2 = 0.274$) was observed for the learning rate of the velocity profile, i.e., reduction of the velocity error, from baseline to retention test on day 2 (Fig. 2, bottom middle). No further significant group differences were found.

DISCUSSION

Currently, huge efforts are invested in machine learning algorithms for automating decision-making in a variety of fields, e.g., disease recognition in radiology (*18–21*) and financial transactions on the stock market (*22–25*). Decision-making algorithms can be used in open-loop settings to condense data, making information comprehensible for humans, such as in radiology (*18–21*). The next level in autonomy after open-loop settings is closed-loop settings that directly influence humans or digital systems. Hereby, the closed-loop system observes and records the resulting actions of the influenced human or system. On the basis of the newly observed data, decision-making can be improved. For example, automated closed-loop systems that already affect us in our daily life are personalized advertisements (*26–28*). Because of the increasing digitalization and volume of available data, the importance of data mining and closed-loop decision making also increases.

However, big data are often missing in fields where (i) instrumentation that needs expert knowledge for installation is required, (ii) setting up data collection systems is time consuming, (iii) only a few data points can be recorded at once, (iv) data acquisition or required technical

Table 1. Dominant error for each participant for each training condition. Augmented feedback was only provided during the training conditions T1 to T5 on day 1 and T6 to T10 on day 2. Cat, spatial error at catch; Drv, spatial error during drive phase; Rel, spatial error at release; Rec, spatial error during recovery phase; Vel, velocity error over the entire cycle. Bold font indicates a match, meaning that both participants had the same dominant error by chance, and therefore, the control participant received feedback focusing on his/her dominant error as well.

| Participant	Day 1					Day 2				
	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10
Exp-1/Ctrl-1	Drv/Vel	Drv/Vel	Drv/Vel	Drv/Vel	Drv/Vel	Drv/Vel	**Vel/Vel**	Drv/Vel	Drv/Vel	**Vel/Vel**
Exp-2/Ctrl-2	Vel/Rec	**Rec/Rec**	Rec/Vel	Rec/Vel	**Vel/Vel**	Vel/Drv	**Rec/Rec**	**Vel/Vel**	Rec/Vel	Vel/Drv
Exp-3/Ctrl-3	Vel/Rec	**Rec/Rec**	Rec/Drv	**Vel/Vel**	Vel/Drv	**Drv/Drv**	Rec/Vel	**Vel/Vel**	Vel/Drv	Vel/Drv
Exp-4/Ctrl-4	**Vel/Vel**	Rec/Vel	Cat/Vel	Drv/Vel	Vel/Drv	Cat/Vel	Vel/Cat	Rec/Cat	**Cat/Cat**	Cat/Drv
Exp-5/Ctrl-5	**Vel/Vel**	Drv/Rec	**Vel/Vel**	Drv/Rec	**Vel/Vel**	**Vel/Vel**	Drv/Rec	Drv/Rec	Vel/Rec	Vel/Rec
Exp-6/Ctrl-6	**Vel/Vel**	Rec/Vel	**Vel/Vel**	**Vel/Vel**	Rec/Drv	**Vel/Vel**	Vel/Drv	**Vel/Vel**	**Vel/Vel**	**Vel/Vel**
Exp-7/Ctrl-7	**Vel/Vel**	**Vel/Vel**	**Vel/Vel**	**Vel/Vel**	**Vel/Vel**	**Vel/Vel**	Rec/Vel	**Vel/Vel**	**Vel/Vel**	**Vel/Vel**
Exp-8/Ctrl-8	**Rec/Rec**	Rec/Drv	**Rec/Rec**	Rec/Drv	**Rec/Rec**	Rec/Drv	Rec/Drv	Drv/Rec	**Vel/Vel**	Rec/Vel
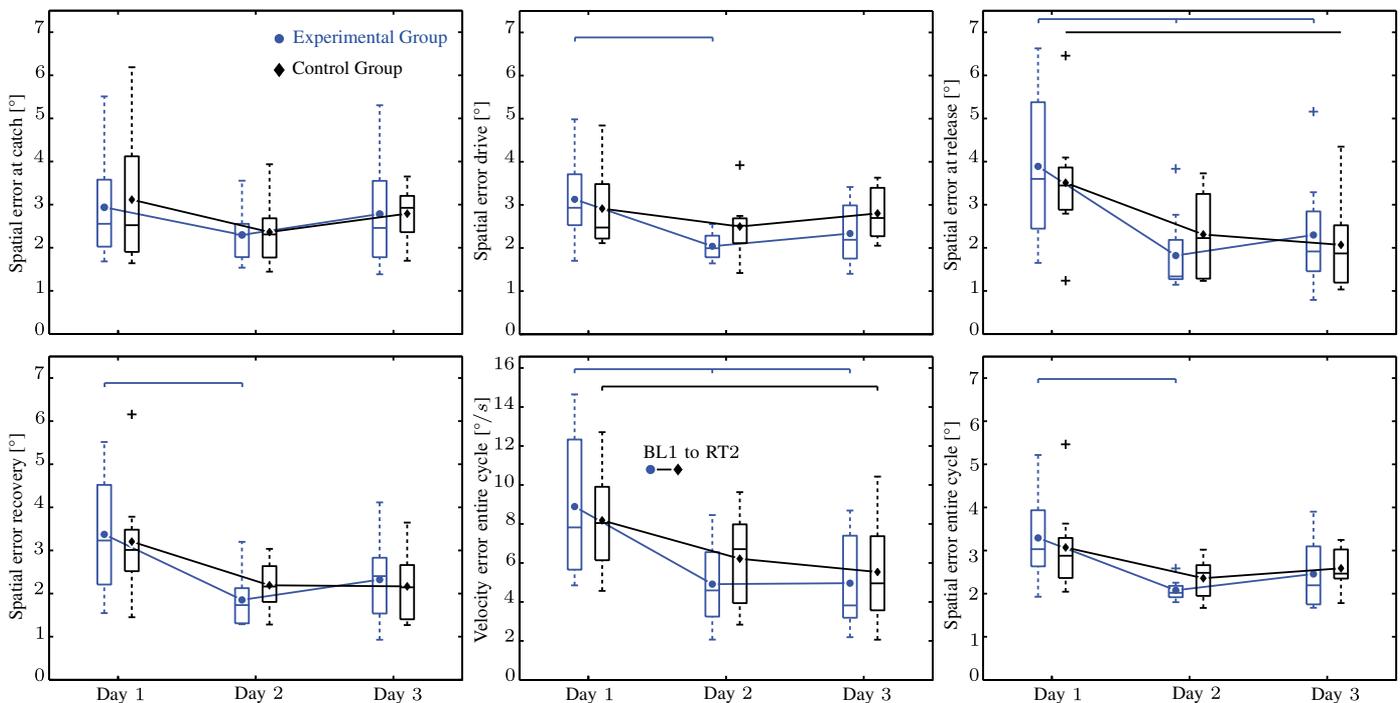
Fig. 2. Development of the six primary outcomes over the main test conditions. Group means are connected between the baseline test on day 1 (BL) and the retention tests on day 2 (RE2) and day 3 (RE3). For each box, the central horizontal line is the median, the edges of the box are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points and cover 99.3% of the data. Outliers are plotted individually (marked by a cross in the color of the corresponding group). Horizontal lines in group colors above the boxes indicate significant main effects of test within the groups. Small vertical bars attached to those horizontal lines indicate Bonferroni-corrected significant differences in the test conditions. Small black lines ending with group symbols and a star above lie between the test conditions and indicate significant group differences in learning rate between the two test conditions.

systems are costly, (v) systems that require particular safety prerequisites are used, or (vi) systems that need particular legal authorization are used. Most of these issues apply to current developments in the field of human-robot interaction. Therefore, it is not unexpected that machine learning has not yet advanced deeply into this field (29). Applications of human-robot interaction that may benefit from automated decision-

making are rehabilitation robots, surgical robotics, and robot-assisted motor learning devices (30). Because of the missing big data, our paper presents an alternative to machine learning techniques: automated decision-making by modeling expert knowledge directly into decision rules.

In the field of motor learning, a few concepts for automated decision-making based on human expertise have been presented (14, 15). However,

Table 2. Learning effects from the baseline test on day 1 (BL) to the retention tests on day 2 (RE2) and day 3 (RE3) between and within groups.

Variable	Interaction and between group effects			Within experimental group			Within control group		
	Group(2)*Test(3) group	P	η_p^2	Main effect post hoc	P	η_p^2	Main effect post hoc	P	η_p^2
Spatial error at catch	$F_{1.40,19.57} = 0.02$ $F_{1,14} = 0.08$	0.943 0.782	0.002 0.006	$F_{2,14} = 0.61$	0.559	0.080	$F_{2,14} = 0.86$	0.444	0.110
Spatial error recovery phase	$F_{1.33,18.61} = 0.45$ $F_{1,14} = 0.00$	0.566 0.992	0.031 0.000	$F_{2,14} = 10.79$ BL–RE2 BL-RE3	0.001 0.019 0.069	0.606	$F_{1.23,8.57} = 2.70$	0.134	0.278
Spatial error drive phase	$F_{2,28} = 1.56$ $F_{1,14} = 0.69$	0.228 0.421	0.100 0.047	$F_{2,14} = 7.07$ BL-RE2 BL-RE3	0.008 0.045 0.116	0.502	$F_{2,14} = 0.88$	0.438	0.111
Spatial error at release	$F_{1.41,19.76} = 1.06$ $F_{1,14} = 0.01$	0.341 0.943	0.071 0.000	$F_{2,14} = 21.57$ BL-RE2 BL-RE3	<0.001 0.006 0.004	0.755	$F_{2,14} = 4.01$ BL-RE2 BL-RE3	0.042 0.223 0.215	0.364
Velocity error entire cycle	$F_{2,28} = 2.52$ $F_{1,14} = 0.09$	0.099 0.768	0.152 0.006	$F_{2,14} = 23.18$ BL-RE2 BL-RE3	<0.001 0.003 0.004	0.768	$F_{2,14} = 10.16$ BL-RE2 BL-RE3	0.002 0.122 0.005	0.592
Spatial error entire cycle	$F_{1.405,19.67} = 0.54$ $F_{1,14} = 0.07$	0.532 0.801	0.037 0.005	$F_{2,14} = 6.78$ BL-RE2 BL-RE3	0.009 0.031 0.196	0.492	$F_{1.12,7.86} = 1.99$	0.199	0.221

Fig. 3. Coauthor demonstrating the rowing simulator. Participants were seated in a real rowing boat placed in a CAVE. The participant's task was to perform a predefined rowing oar movement only with arms and upper body (roll seat was strapped back to prevent the rower from bending the legs). When the virtually prolonged oar was moved through the virtual water, the rower saw, heard, and felt realistic oar-water interactions. In addition to rendering haptic oar-water interactions, the tendon-based parallel robot was used to haptically guide the participant if the related feedback design was selected. The current photo shows visual feedback during the drive phase (green and red dots drawn right above and below the blade).

this work differs by evaluating such concepts in a motor learning study and thereby also involving human-in-the-loop control in connection with a robotic device. Our motor task was complex and requested the minimization of different types of movement errors. We realized real-time data recording and evaluation and applied a simple training strategy, i.e., giving feedback on the dominant error. Therefore, the feedback was selected from a cascade of augmented concurrent feedback designs.

Dominant error occurrences

Commonly, concepts of automated learning address the reduction of one type of error (14, 15), where participants had to minimize spatial and velocity errors. To identify the dominant error, all error values were normalized by the best results obtained from previous studies that were reevaluated for this paper. Normalization of the error metrics allowed a direct comparison of errors that would otherwise have different units (° for the spatial errors and °/s for the velocity error). In addition, normalization by the smallest errors observed in previous studies allowed us to reasonably scale the magnitude of different error metrics to each other based on effectively reachable performance. Therefore, no manual tuning of error metrics had to be performed, and only minimal prior knowledge or data were required compared with other approaches relying on large datasets (13).

Although the velocity error was dominant for 50% of all trainings and the spatial error at release was not dominant for a single training, the sequence of dominant errors was rather individual (Table 1). Because the velocity error was predominant (40 times in the experimental group and 47 times in the control group, matching 28 times), the dominant error in a specific assessment of matched participants was often the same. In 45% of the trainings, the participants of the control group received feedback on their dominant error. Thus, differences between the two training groups could only arise from 55% of the training conditions. Because the participants of our control group had already experienced our chosen training strategy (focusing on the dominant error) so often, the sensitivity of our study with respect to finding group differences was reduced. However, other control groups are less suitable for answering our research question (i.e., whether the participants would benefit from an automatically individualized feedback sequence). For example, a control group relying only on intrinsic feedback without receiving any augmented feedback would be reasonable to answer other research questions, such as whether our task could be learned intrinsically, which is not the case [see (13)]. Other control groups receiving randomly selected feedback designs or feedback designs

focusing on nondominant errors could provide higher sensitivity to detect potential differences between control and experimental groups. However, such control groups would result in different random occurrences of feedback designs (and switching between those). Unfortunately, for such more sensitive control groups, the resulting group differences might be based purely on the differences in feedback design occurrence. Such control groups would lack specificity to answer our question of whether the personalized dominant training strategy was truly superior. Instead, all potential group differences might result from the fact that just one specific feedback design (or combination of a few) is more effective than other designs. Therefore, we are convinced that enforcing feedback occurrences (and switching between those) equally in both groups was the most suitable choice to answer our research question considering both sensitivity and specificity.

Missing significant group effects in the repeated-measures ANOVA

Because of within-group variability, it is usually difficult to find significant interactions between groups based on small group samples. This is also true for our study. In addition, the feedback for both groups was different only in 55% of the training sessions. We observed learning effects within each group particularly for the spatial error at release and the velocity error over the entire cycle. Because both groups learned, a tremendous improvement by one group might be required from one day to another to show interactions between groups. However, potential flooring effects, e.g., owing to the feedback designs, might have limited further learning from day 2 to day 3 (Fig. 2). A trend ($P < 0.1$) for a group interaction could be found for the velocity error (Table 2). This trend indicates that differences in the development between the experimental and the control group are likely and might become significant with a larger group sample. Therefore, we assume that the missing significant group effects in the repeated-measures ANOVA were caused by the small group sizes and high baseline variability.

Significant group difference in learning rate

In contrast to repeated-measures ANOVA on the six error metrics (Fig. 2), learning rates for the one-way ANOVA were calculated on the basis of their baseline values. This normalization of learning rates by the baseline values reduced their sensitivity to the high intersubject variability at baseline (Fig. 2) and allowed a more detailed insight into the difference of error reduction between groups.
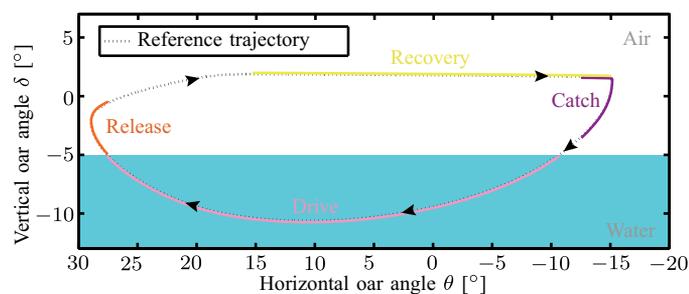
A significant group difference between baseline on day 1 and retention on day 2 in the velocity error was found. Therefore, individual differences in the provided versus optimal feedback have a considerable impact on the learning rate. In our opinion, the significant group difference in the learning rate could be expected in the early phase of the study, because most learning was observed between day 1 and day 2. In particular, the significant group difference for the velocity error could be explained by the fact that this error was trained most frequently. Furthermore, this significant difference in learning rate supports our interpretation of the tendency observed in repeated-measures ANOVA. To conclude, this significant difference in learning rate outlined that our closed-loop controlled, automated feedback selection worked as intended.

Interpretation of learning effects

The fact that the learning effects could be observed mainly on the velocity error is not unexpected because this error was trained the most. Learning was also observed in both groups for the spatial error at release, even though feedback specific to the release was never selected. However, all secondary feedbacks included haptic information on the complete movement. Eventually, the haptic information on release contained in these secondary feedbacks was enough for the participants to learn. It is possible that the spatial error at release is easier to improve compared with the other spatial errors, because most of the arm movement is close to the participant's trunk and the trunk might have acted as a physical end stop preventing large errors.

With respect to the velocity error and the spatial error entire cycle, both groups showed improvements and reached absolute performance levels similar to those in previous studies of the same task (13, 16, 31). For this reason, the subdivision of visual feedback into the four rowing phases and the novel multimodal feedback combinations proved to be suitable for learning the rowing movement.

CONCLUSION

Our automated feedback selection strategy showed a significantly higher learning rate (Fig. 2) despite relying on a very simple concept: provision of feedback addressing the dominant error, which could be obtained by normalizing observed errors based on previous knowledge. Common knowledge of human trainers was implemented in two decision rules for the robotic trainer. First, focusing on more severe errors is more likely to result in stronger improvements. Second, choosing another feedback design that provides additional or alternative information, e.g., in another modality, might be beneficial if the same error stays dominant for a trainee.

We consider the herein presented methods as generalizable to any type of automated feedback selection based on arbitrary quantitative metrics. In contrast, we consider the specific results on which one performance metric was learned faster by our rowing novices because of personalized automated feedback selection limited to the herein selected trunk-arm rowing task.

We believe that the herein proposed concept contributes particularly to fields where automated decision-making is desired, but little prior data are available—such as when devices or concepts are evaluated for the first time, when study protocols are changed, when new feedback designs or training strategies are under evaluation, or when effects of a device are difficult to predict. The only required knowledge is a minimum amount of data to allow comparison of different types of error metrics, e.g., by normalization. These few values for normalization can be obtained in a pilot study or by estimation through an expert.
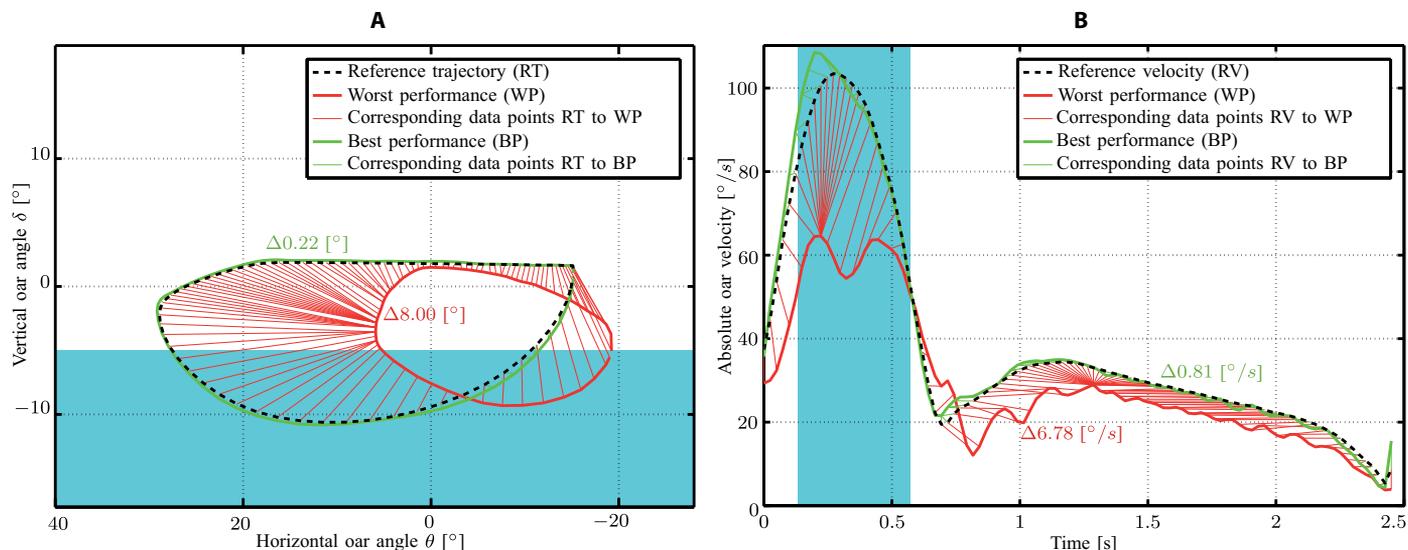

Fig. 4. Desired oar movement. The arrowheads indicate the desired direction of the cyclic oar movement. The desired oar movement was divided into four phases: the region around the horizontal turning point before the oar enters the water (catch), the drive phase, the region around the horizontal turning point after the oar exits the water (release), and the region where the oar moves horizontally through air (recovery phase).

Fig. 5. Visualization of performance metrics mapping to the reference movement. DTW is used to correlate the measured spatial (**A**) and velocity (**B**) data to the reference trajectory. The two measured datasets show the worst and best performance measured during this study. The best performance was measured during position control. The spatial error (A) is indicated by the mean value over the entire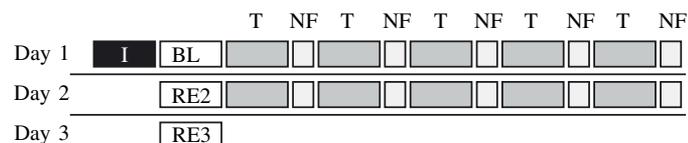 rowing cycle in the color of the corresponding performance. The velocity error (B) is indicated by the mean value over the entire rowing cycle in the color of the corresponding performance.

Fig. 6. Study protocol for the three consecutive training/measurement days. The no-feedback trials lasted for 1 min (indicated by small block). All other measurement blocks lasted for 3 min. I, instruction; BL, baseline measurement; T, training; NF, no feedback trial; RE2, retention test on day 2; RE3, retention test on day 3.

Outlook

This paper highlights a possible implementation of a working closed-loop feedback selection. In our opinion, extensions and modifications of our concept have the potential to result in increased learning rates, an increased level of automation (including the adaptation of the learning goal), or both. Such extensions/modifications could be (a) the inclusion of novel performance metrics; (b) different normalization and weighting; (c) selection rules, including additional knowledge of human trainers; and (d) implementation of a long-term training plan. As a possible example of (a), including variability as a performance metric seems promising, because low variability is an indicator of expert performance (32). To increase real-life relevance of closed-loop feedback selection, normalization and weighting (b) could be performed, reflecting the individual performance metrics' effects on a higher-level goal, such as a maximization of the mean boat velocity. Additional trainer knowledge (c), such as more variations within training, can lead to better training effects (33–35). This could be realized by a selection rule that switches feedback more often and has a larger repertoire of feedback designs. Also, the implementation of a long-term training plan (d) could be beneficial for the learning of complex tasks over a longer period of time. Usually, training plans focus first on the learning of simple parts of a movement before these simple parts are combined to form a more complex complete movement (36, 37). As more data are available, we expect that expansion of our concept by machine learning algorithms

for the automated feedback selection will lead to even larger effects in group differences. Even extensions to automated strategy adaptation or automated training goal change itself are conceivable.

MATERIALS AND METHODS
Apparatus
Over the past few years, a rowing simulator that can be used as a controlled setting for motor learning studies has been developed at ETH Zurich (38, 39). This rowing simulator enables rowers inside a Cave Automatic Virtual Environment (CAVE) to see, hear, and feel their interactions with a virtual rowing scenario. For the current study, sweep rowing was set up, where the rower inside the CAVE held one bow side oar at the oar handle (Fig. 3). A customized tendon-based parallel robot was attached to the blade side end of the oar and could render water forces as well as apply haptic feedback (sample frequency, 1 kHz) (40, 41). The visual scenario was displayed on three large screens surrounding the rower in a U-shape (4.4 m by 3.3 m, projectors: Projection Design F3+, Norway; update frequency, >30 fps). Auditory rendering and feedback were provided through standard stereo headphones (update frequency, ~30 Hz).

Task
Participants were asked to learn a desired trunk-arm rowing movement. Trunk-arm rowing incorporates a complex oar blade movement that is commonly used during warm-up or technical training. Reproducing a desired trunk-arm movement was considered a suitable motor learning task for rowing novices. In addition, movement reproduction also provides the possibility for quantitative performance evaluation. The desired oar blade movement incorporated several challenges that can be separated spatially into four distinct phases of the rowing cycle: (i) the immersion of the oar blade into the water, i.e., "catch"; (ii) the movement of the oar through the water, i.e., "drive phase"; (iii) the movement of the oar out of water, i.e., "release"; and (iv) the horizontal movement of the oar through air back to catch, i.e., "recovery phase." The catch is determined by a distinctive turn of the oar blade and a fast vertical

movement of the oar into the water. During the drive phase, the rower applies high forces to the oar, pushing the boat forward. The release is characterized by a fast turn of the oar blade and a deceleration. During the recovery phase, the rower slowly moves the oar back to the start for the next rowing cycle on a horizontal trajectory while avoiding undesired deceleration of the boat (Fig. 4). The desired oar blade movement from an expert rower was recorded.

In postprocessing, the trajectory was smoothed in space and velocity to form a cyclic C^2-continuous trajectory, i.e., a trajectory that can be continuously differentiated twice and has neither start nor end point. Furthermore, the desired trajectory was scaled to an angular horizontal range of 44° and a vertical angular range of 12°. This corresponded to a horizontal range of the oar handle's tip of 0.67 m and a vertical range of 0.19 m. The desired movement was designed to be feasible only by trunk and arm movements for participants taller than 1.65 m. The desired movement should be executed at a desired stroke rate of 24 strokes/min.

Kinematic evaluation

The participants' performance was analyzed with respect to five different errors. Four of these separately represented the spatial deviation of each rowing stroke from the reference movement during the four rowing phases: catch, drive, release, and recovery (Fig. 4). A velocity error assessed the deviation of the measured velocity profile of the entire movement for each rowing stroke from the reference velocity profile. After the first three strokes in a test condition (which were needed to accelerate the boat; these strokes were not considered in the analysis), participants had to perform at least eight strokes that were within the desired stroke rate of 24 strokes/min with a tolerance of ±2 strokes/min. For each of the five error types, the error values of these eight or more strokes were averaged over one test condition.

Kinematic evaluation was performed both online during the test conditions used for automated feedback selection and in postprocessing of recorded raw data to analyze and report study results. Online evaluation was performed on an xPC Target (MATLAB/Simulink) that controlled the tendon-based parallel robot and rendered the haptic interactions of the virtual rowing scenario. Although the rowing scenario and the robot control were controlled at 1 kHz, the online data analysis was performed at the reduced sample frequency of 100 Hz. Raw movement data were also saved for postprocessing at a sample frequency of 100 Hz. Postprocessing evaluation was performed with custom scripts (MATLAB).

To divide the measured data according to the defined phases onto the desired trajectory, we mapped the measured data to the desired trajectory stroke by stroke. The measured data were first cut into single strokes at catch (except for analysis of errors at catch, for which the rowing strokes were cut at release). In a subsequent step, the cut data were resampled to 250 data points for comparison with the desired trajectory that also consisted of 250 data points. Online evaluation was conducted on a lower temporal resolution to reduce calculation time by resampling to only 101 data points instead of 250. On each of the cut rowing strokes, a dynamic programming algorithm, namely, dynamic time warping (DTW), was applied for performance evaluation. This DTW found the optimally correspondent data points on the reference trajectory for the measured data points (42):

$$E_c[\xi, \tau] = \int \left[|\xi(t)|^2 + \lambda \tau(t)^2 \right] dt \qquad (1a)$$

$$\frac{d\tau}{dt} > -1 \qquad (1b)$$

$$\tau(t_{\min}) = \tau(t_{\max}) = 0 \qquad (1c)$$

In the optimization cost function (Eq. 1a), E_c denoted the error cost of the spatial shift vector ξ and the temporal shift τ. The positive constant λ denoted the relative weight between spatial and temporal errors (42). The first optimization constraint (Eq. 1b) ensures that assignment of corresponding data points cannot cross, and the causal order of measured data point sequences is preserved. The second optimization constraint (Eq. 1c) sets temporal shifts of the first and last data points to zero to enforce their assignment in both trajectories. For all performance measures, the relative weight constant λ was set to zero to match spatially the measured data as closely as possible. To evaluate one rowing stroke of interest without the reduced sensitivity caused by the constraint (Eq. 1c), we applied DTW to align a trajectory of three reference strokes against a triplet of recorded rowing strokes including the stroke before, the stroke of interest, and the stroke after. Error outcomes were then calculated on the basis of the alignments of only the rowing stroke of interest. For all spatial performance measures, the spatial shift vector ξ incorporated deviations in θ and δ directions (Fig. 5, left). For the velocity errors, the spatial shift vector ξ from Eq. 1a was adapted to represent the shift in velocity magnitude (Fig. 5, right).

To allow a comparison of the five errors, we normalized the resulting average errors by the results of two previous motor learning studies (16, 31). In these two studies, the same desired oar movement had to be learned, the same study protocol was applied, and similar or the same feedback strategies were used. We reevaluated the recorded raw data of these studies to assess the performance in the five error types of interest. For normalization, the best group mean from the retention test on day 3 in each error type was taken and rounded to one digit. Numeric group means used for normalization were $\bar{\epsilon}_{\text{spatial, catch}} = 2.7°$, $\bar{\epsilon}_{\text{spatial, drive}} = 1.7°$, $\bar{\epsilon}_{\text{spatial, release}} = 2.7°$, $\bar{\epsilon}_{\text{spatial, recovery}} = 1.7°$, and $\bar{\epsilon}_{\text{velocity, entireCycle}} = 3.4°/\text{s}$.

In addition to the five primary outcomes, the spatial error of the entire rowing cycle was evaluated in postprocessing using DTW. All distances between each point of the desired rowing trajectory and the corresponding point on the measured trajectory obtained through DTW were averaged for the complete rowing stroke. In contrast to the four spatial rowing errors, the spatial error of the entire rowing cycle should provide an impression of the participants' overall spatial performance and simplify our comparison with previous studies (16, 31). Besides the evaluation of the test conditions, the performance during training with augmented feedback was also evaluated in the same way and is reported in the appendix.

Automated feedback selection focusing on dominant error

The dominant error was obtained as the highest normalized mean error from the online kinematic evaluation of the last condition. For each type of dominant error, two augmented feedback designs were available, corresponding to primary and secondary designs. The primary feedback design was provided if an error type emerged as dominant for the first time in a row. If the same error type was dominant a second time or even more often in a row, then the secondary feedback design was provided. In total, 10 different feedback designs (primary and secondary for each of the five different errors) were available for selection from the planned 10 training conditions.

Augmented feedback designs

The primary augmented feedback design was always unimodal, whereas the secondary augmented feedback design was always multimodal. All

primary feedback designs on spatial errors were based on visual cues. The primary feedback design on the velocity error was based on auditory cues. For the secondary feedback, haptic guidance was added to the primary feedback. For spatial errors, a path controller in addition to visual feedback should help to further minimize spatial deviations from the requested oar blade trajectory, whereas a position controller in addition to auditory feedback should help to further minimize velocity errors.

Rowers could always see their virtually prolonged oar on the right screen. For the primary spatial feedback, this virtually prolonged oar drew a trace when it deviated by more than 3.6° vertically or 1.9° horizontally from the target trajectory (the values differ because the blade is more tall than wide). The trace was only drawn in the area of the spatial error type that should be corrected. The trace was green close to the trajectory, became an increasing shade of red the larger the deviation, and saturated to fully red at deviations larger than 11.5°. The trace faded out after 8 s, i.e., the trace was visible for about three rowing strokes. This visual feedback design was chosen because it was successfully applied in a previous motor learning study where the same task was taught (31).

As a secondary feedback strategy for spatial errors, the primary spatial feedback was combined with haptic guidance in terms of a path controller. The path controller supported the rower to stay in the vicinity of the reference along the entire desired trajectory without imposing timing, i.e., forces pointing normal toward the reference were applied if the rower was outside of a tolerance or dead-band around the reference. A more detailed description of the applied path controller can be found elsewhere (16). A path controller was chosen because, in comparison with three other types of haptic guidance, it was most effective in teaching spatial aspects of the same desired rowing trajectory in a previous study (16).

The primary feedback for velocity comprised a sonification of the rower's own oar movement (right speaker) and the desired oar movement (left speaker). The same sonification function, i.e., mapping of movement data to sound dimensions, was used for the rower's and the desired oar movement. In consequence, the rower could synchronize his or her own oar movement with the desired oar movement. Horizontal oar angles were sonified by violin sounds that changed their pitch from 54.5 Hz (approximately A1) to 91.58 Hz (approximately F#2) when the oar was outside the water. When the oar was inside the water, purling sounds were faded in, whereas the violin sounds were faded out, depending on the deepness of the oar inside the water. In a previous study, the same auditory feedback in combination with visual feedback for motor learning of the same task represented an advantage compared with visual feedback alone. This advantage is supposed to originate from the auditory feedback's potential to teach dynamic aspects of the oar movement (31).

As a secondary feedback on the velocity error, the auditory, primary feedback was combined with a position controller. The position controller enforced the desired oar trajectory with spatial errors lower than 0.5° and velocity errors lower than 1°/s. The position controller was realized with a proportional component (P_{pos} = 6000 N/rad) and a derivative component (D_{pos} = 170 Ns/rad). In a previous study of the same rowing task, this position controller was the most effective feedback design to learn the velocity profile in comparison with three other types of haptic guidance (16).

Participants and groups

For the study, 16 participants were recruited according to the following inclusion criteria: (i) no prior rowing experience; (ii) aged between 18 and 35 years; (iii) at least 0.5 hours of sport per week; (iv) healthy (no physical impairments or discomforts); (v) normal hearing, normal (or corrected to normal) vision; and (vi) taller than 1.65 m.

Overall, seven women and nine men [mean age, 27.7 years; SD, 1.9 years (from 25 to 32 years)] were included. The study was approved by the ethics commission of ETH Zurich (EK_2010-N-53). All participants signed a written consent form before the study. In addition, all participants were informed verbally about the study procedure, the risks, and their option to quit at any time without a reason and without consequences.

The participants were randomly assigned to one of two training groups, i.e., eight participants to the experimental group and eight participants to the control group. Each participant of the control group was matched to one participant of the experimental group. Participants of the experimental group received an automatically compiled, individual sequence of feedback designs addressing their dominant error, and participants of the control group obtained the feedback sequence of their matched participant of the experimental group.

Protocol

The measurement protocol was designed to assess motor learning over three consecutive days. On days 1 and 2, all participants performed an initial baseline (day 1) or retention test (day 2) before the training started. The training itself consisted of five training blocks composed of training with feedback followed by a no-feedback trial. On day 3, participants performed only a retention test. Measurements on day 1 took 70 min, on average, including administrative work, an introduction of the rowing simulator, explanation of the task, and mechanical adaption of the simulator to the participant. Measurements on day 2 took about 45 min, and measurements on day 3 lasted for about 15 min (Fig. 6).

Before the study, the roll seat in the rowing simulator was fixed at the same position for all participants. Therefore, the foot stretcher in the rowing boat had to be set up individually so that all participants could be seated with their legs stretched. The foot stretcher position was noted and reset for each participant before each measurement. On the first day, after the rowing simulator was introduced to each participant, one of the principal investigators took place in the simulator. The principal investigator demonstrated and explained the handling of the rowing simulator and its safety features to the participant. After watching the investigator for 1 min from outside the boat, the participant took his or her place in the rowing simulator and could test the foot stretcher settings and the reachability of all points of the desired trajectory by the same path controller, which was used for augmented feedback, for about half a minute, i.e., the oar could be moved freely inside a tunnel around the desired path. The investigator told the participant to focus well on the velocity profile, the shape, and the stroke rate of the desired rowing trajectory (without mentioning the value of the desired stroke rate). In addition, all participants were told that the robot would select feedback designs for training individually adapted to their largest errors. Further instructions were not provided. Training of day 1 started by the instruction through position control (also used for augmented feedback) for 3 min (72 rowing strokes). After the instruction, the participant was asked to reproduce the instructed movement in the best way possible (baseline measurement). Then, the first training block started and consisted of 3 min of training with feedback followed by a no-feedback trial for 1 min. Training blocks were repeated five times on day 1. Day 2 started with a retention test of 3 min, followed by five training blocks like on day 1. On day 3, participants only performed a retention test (Fig. 6).

During the baseline measurements, the non-feedback trials, and the retention tests, the participants experienced only the rendering of the

rowing scenario. To keep a constant nominal task complexity level for all participants during all the test conditions, we instructed the participants to row faster or slower as soon as they deviated more than 2 strokes/min from the desired stroke rate of 24 strokes/min.

Statistical evaluation

Statistical analysis was performed on all performance measures in IBM SPSS Statistics for Windows, version 20.0 (IBM Corp., NY). Two-sided P values lower than 0.05 indicated significant differences. The mean values of all performance measures for each 3-min test condition (baseline and retention test on days 2 and 3) and for each participant were used as an input for the statistical analysis. Consequently, the input data for the statistical analysis of each test condition contained 16 mean values (two groups of eight participants each). Statistical analysis was performed independently for each of the six outcomes of the kinematic evaluation.

A one-way ANOVA at baseline between the groups was performed to ensure that no significant group differences at baseline would affect our results. Levene's test was performed to ensure that the assumption of equal variance was not violated.

To assess whether automated feedback selection accelerates motor learning, we performed two different tests to find group differences. First, to assess learning, a 2-by-3 repeated-measures ANOVA was performed for the two feedback groups and the three test conditions: baseline, retention test on day 2, and retention test on day 3. To assess learning within the individual groups, we performed a within-group follow-up repeated-measures ANOVA. Post hoc Bonferroni tests were applied for multiple comparisons between the tests. Violations of sphericity were tested with Mauchly's test and corrected using Greenhouse-Geisser correction.

Second, to compare learning between groups in early and late stages, we compared the development of a variable from baseline to the retention test on days 2 and 3 by running one-way ANOVA on the learning rates. The learning rate denotes the difference from one test condition to the other, normalized by the value of the first test condition. Levene's test was used to check for violations in the assumption of equal variance, whereas Tukey honestly significantly different post hoc test was used for multiple comparisons between groups.

SUPPLEMENTARY MATERIALS

robotics.sciencemag.org/cgi/content/full/4/27/eaav1560/DC1
Table S1. Between- and within-group differences during training with and without feedback.
Movie S1. M³-rowing simulator

REFERENCES AND NOTES

1. J. Bartlett, J. E. Lawrence, M. E. Stewart, N. Nakano, V. Khanduja, Does virtual reality simulation have a role in training trauma and orthopaedic surgeons? *Bone Joint J.* **100**, 559–565 (2018).
2. D. L. Neumann, R. L. Moffitt, P. R. Thomas, K. Loveday, D. P. Watling, C. L. Lombard, S. Antonova, M. A. Tremeer, A systematic review of the application of interactive virtual reality to sport. *Virtual Reality* **2017**, 1–16 (2017).
3. R. Sigrist, G. Rauter, R. Riener, P. Wolf, Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review. *Psychon. Bull. Rev.* **20**, 21–53 (2013).
4. A. Kos, V. Milutinovic, A. Umek. Challenges in wireless communication for connected sensors and wearable devices used in sport biofeedback applications. *Future Gener. Comput. Syst.* **92**, 582–592 (2018).
5. J. C. P. Chan, H. Leung, J. K. T. Tang, T. Komura, A virtual reality dance training system using motion capture technology. *IEEE Trans. Learn. Technol.* **4**, 187–195 (2011).
6. I. de Kok, F. Hülsmann, T. Waltemate, C. Frank, J. Hough, T. Pfeiffer, D. Schlangen, T. Schack, M. Botsch, S. Kopp, The intelligent coaching space: A demonstration, in *International Conference on Intelligent Virtual Agents* (Springer, 2017), pp. 105–108.
7. P. Kelly, A. Healy, K. Moran, N. E. O'Connor, A virtual coaching environment for improving golf swing technique, in *Proceedings of the 2010 ACM Workshop on Surreal Media and Virtual Cloning* (ACM, 2010), pp. 51–56.
8. H. I. Krebs, J. J. Palazzolo, L. Dipietro, M. Ferraro, J. Krol, K. Rannekleiv, B. T. Volpe, N. Hogan, Rehabilitation robotics: Performance-based progressive robot-assisted therapy. *Auton. Robots* **15**, 7–20 (2003).
9. J. Huegel, M. K. O'Malley, *Haptics Symposium, 2010 IEEE* (IEEE, 2010), pp. 343–350.
10. L. Kahn, W. Rymer, D. Reinkensmeyer, Adaptive assistance for guided force training in chronic stroke, in *26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (IEEE, 2004), vol. 1, pp. 2722–2725.
11. L. Marchal-Crespo, M. van Raai, G. Rauter, P. Wolf, R. Riener, The effect of haptic guidance and visual feedback on learning a complex tennis task. *Exp. Brain Res.* **231**, 277–291 (2013).
12. L. Marchal-Crespo, P. Wolf, N. Gerig, G. Rauter, L. Jaeger, H. Vallery, R. Reiner, The role of skill level and motor task characteristics on the effectiveness of robotic training: first results, in *2015 IEEE International Conference on Rehabilitation Robotics (ICORR)* (IEEE, 2015), pp. 151–156.
13. N. Gerig, P. Wolf, R. Sigrist, R. Riener, G. Rauter, Automated feedback selection for robot-assisted training. *Int. J. Comput. Sci. Sport* **16**, 149–174 (2017).
14. K. Rector, C. L. Bennett, J. A. Kientz, Eyes-free yoga: An exergame using depth cameras for blind & low vision exercise, in *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ACM, 2013), p. 12.
15. F. Hülsmann, C. Frank, T. Schack, S. Kopp, M. Botsch, Multi-level analysis of motor actions as a basis for effective coaching in virtual reality, in *Proceedings of the 10th International Symposium on Computer Science in Sports (ISCSS)* (Springer, 2016), pp. 211–214.
16. G. Rauter, R. Sigrist, R. Riener, P. Wolf, Learning of temporal and spatial movement aspects: A comparison of four types of haptic control and concurrent visual feedback. *IEEE Trans. Haptics* **8**, 421–433 (2015).
17. R. Sigrist, G. Rauter, R. Riener, P. Wolf, Terminal feedback outperforms concurrent visual, auditory, and haptic feedback in learning a complex rowing-type task. *J. Mot. Behav.* **45**, 455–472 (2013).
18. S. Wang, R. M. Summers, Machine learning and radiology. *Med. Image Anal.* **16**, 933–951 (2012).
19. Z. Obermeyer, E. J. Emanuel, Predicting the future—Big data, machine learning, and clinical medicine. *N. Engl. J. Med.* **375**, 1216–1219 (2016).
20. E. I. Zacharaki, S. Wang, S. Chawla, D. Soo Yoo, R. Wolf, E. R. Melhem, C. Davatzikos, Classification of brain tumor type and grade using MRI texture and shape in a machine learning scheme. *Magn. Reson. Med.* **62**, 1609–1618 (2009).
21. P. Lakhani, B. Sundaram, Deep learning at chest radiography: Automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology* **284**, 574–582 (2017).
22. X. Ding, Y. Zhang, T. Liu, J. Duan, *Ijcai* (2015), pp. 2327–2333.
23. R. Akita, A. Yoshihara, T. Matsubara, K. Uehara, Deep learning for stock prediction using numerical and textual information, in *2016 IEEE/ACIS 15th International Conference on Computer and Information Science (ICIS)* (IEEE, 2016), pp. 1–6.
24. J. Patel, S. Shah, P. Thakkar, K. Kotecha, Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. *Expert Syst. Appl.* **42**, 259–268 (2015).
25. S. Mullainathan, J. Spiess, Machine learning: An applied econometric approach. *J. Econ. Perspect.* **31**, 87–106 (2017).
26. J. Brouhard, System and method of a media delivery services platform for targeting consumers in real time. U.S. Patent 9,811,836 (2017).
27. D. L. Selinger, T. Kohn, M. DeCourcey, S. Ahuja, J. Osial, A. Sunwoo, System and process for improving product recommendations for use in providing personalized advertisements to retail customers. U.S. Patent 8,924,265 (2014).
28. V. Bhagwan, J. Savla, D. Sharp, R. Teotia, J. Andersen, D. Agarwal, J. Wang, Advertisement generator. U.S. Patent 14/586,886 (2016).
29. A. Holzinger, Interactive machine learning for health informatics: When do we need the human-in-the-loop? *Brain Inform.* **3**, 119–131 (2016).
30. O. C. Santos, Training the body: The potential of AIED to support personalized motor skills learning. *Int. J. Artif. Intell. Educ.* **26**, 730 (2016).
31. R. Sigrist, G. Rauter, L. Marchal-Crespo, R. Riener, P. Wolf, Sonification and haptic feedback in addition to visual feedback enhances complex motor task learning. *Exp. Brain Res.* **233**, 909–925 (2015).
32. K. A. Ericsson, The path to expert golf performance: Insights from the masters on how to improve performance by deliberate practice, in *Optimising Performance in Golf*, P. R. Thomas, Ed. (Australian Academic Press, 2001), pp. 1–57.
33. D. C. Shapiro, *The Development of Movement Control and Coordination* (Springer, Dordrecht, 1982), pp. 113–150.
34. C. H. Shea, R. M. Kohl, Specificity and variability of practice. *Res. Q. Exerc. Sport* **61**, 169–177 (1990).

35. W. X. Yao, A. Cordova, W. De Sola, C. Hart, A. F. Yan, The effect of variable practice on wheelchair propulsive efficiency and propulsive timing. *Eur. J. Phys. Rehabil. Med.* **48**, 209–216 (2012).

36. J.-H. Park, H. Wilde, C. H. Shea, Part-whole practice of movement sequences. *J. Mot. Behav.* **36**, 51–61 (2004).

37. J. Klein, S. J. Spencer, D. J. Reinkensmeyer, Breaking it down is better: Haptic decomposition of complex movements aids in robot-assisted motor learning. *IEEE Trans. Neural Syst. Rehabil. Eng.* **20**, 268–275 (2012).

38. J. von Zitzewitz, P. Wolf, V. Novakovic, M. Wellner, G. Rauter, A. Brunschweiler, R. Riener, Real-time rowing simulator with multimodal feedback. *Sports Technol.* **1**, 257–266 (2008).

39. G. Rauter, J. von Zitzewitz, A. Duschau-Wicke, H. Vallery, R. Riener, A tendon-based parallel robot applied to motor learning in sports, *2010 3rd IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)* (IEEE, 2010), pp. 82–87.

40. J. von Zitzewitz, A. Morger, G. Rauter, L. Marchal-Crespo, D. Wyss, T. Bruckmann, F. Crivelli, R. Riener, A reconfigurable, tendonbased haptic interface for research into human-environment interactions. *Robotica* **18**, 687–689 (2012).

41. G. Rauter, R. Sigrist, C. Koch, F. Crivelli, M. van Raai, R. Riener, P. Wolf, Transfer of complex skill learning from virtual to real rowing. *PLOS ONE* **8**, e82145 (2013).

42. M. A. Giese, T. Poggio, Morphable models for the analysis and synthesis of complex motion patterns. *Int. J. Comput. Vision* **38**, 59–73 (2000).

Citation: G. Rauter, N. Gerig, R. Sigrist, R. Riener, P. Wolf, When a robot teaches humans: Automated feedback selection accelerates motor learning. *Sci. Robot.* **4**, eaav1560 (2019).

Science Robotics

When a robot teaches humans: Automated feedback selection accelerates motor learning

Georg Rauter, Nicolas Gerig, Roland Sigrist, Robert Riener and Peter Wolf

ARTICLE TOOLS	http://robotics.sciencemag.org/content/4/27/eaav1560
SUPPLEMENTARY MATERIALS	http://robotics.sciencemag.org/content/suppl/2019/02/15/4.27.eaav1560.DC1
REFERENCES	This article cites 27 articles, 0 of which you can access for free http://robotics.sciencemag.org/content/4/27/eaav1560#BIBL
PERMISSIONS	http://www.sciencemag.org/help/reprints-and-permissions

Use of this article is subject to the Terms of Service