Emergent Effector-Independent Internal Spaces: Adaptation and Intermanual Learning Transfer in Humans and Neural Networks

Martin V. Butz, Alexandra Lenhard, Oliver Herbort

Abstract— **Psychological studies have shown immense behavioral flexibility in arm reaching tasks. Intermanual learning transfer (ILT) tasks have shown that both reaching movements adapt to distorted spaces rather rapidly and the adaptation generalizes to the behavior of other limbs. In this paper, we present an ILT experiment and replicate it with feedforward neural network (NN) architectures. We show that the NN architecture is the key to successfully replicating the experiments. Moreover, we show that dependent on the architecture and the initial training schedule applied, an internal space representation emerges that enables ILT. The results confirm that internal body spaces, identified in neuroscience and cognitive psychological research, can emerge solely due to an interdependence between different limb movements and the right neural architecture. We hypothesize that, in order to develop internal spatial representations observed in animals and humans, it might be sufficient to enforce the integration of multiple correlated sensory and motor information into one compact internal representation.**

I. INTRODUCTION

In psychological research, intermanual learning transfer (ILT) has been recorded for at least 150 years, e.g. [1], [2]. Training one limb does not only improve performance of the trained but also of the untrained, contralateral limb. Although many studies have investigated this topic and have shown various types of ILT, it remains unclear how learning transfer is accomplished in the brain.

In this paper, we present a recent study on ILT. Participants showed learning transfer in an aiming task with a target-specific prism-like visual displacement. After training with the right arm, the left arm exhibited the same adapted aiming behavior without any actual training. Thus, ILT of a distorted workspace was observed.

We model this behavior with two multi-layer neural network (NN) architectures, trained with standard backpropagation techniques [3]. We show that it is impossible to qualitatively replicate the observed learning transfer with a naïve approach. Based on the assumption that eye-centered coordinates are transformed into an effector-independent spatial representation, which controls both arms, we show that ILT can be modeled using a Y-shaped NN architecture. We show that when the two arms are trained with common goals, an internal representation emerges that allows for effective learning transfer. That is, also the untrained arm adapts to the distorted work-space environment, as observed in the experiments with humans.

We hypothesize that an internal effector-independent body-space representation emerges, because the internal NN layers

The authors are with the Department of Psychology, University of Würzburg, Germany, (phone: +49 931 312808; fax: +49 931 312815; email: butz@psychologie.uni-wuerzburg.de).

suitably transfer knowledge gained from one effector to the other. The representations form due to the enforced compact integration of multiple sensory and motor information and their interdependence. Thus, internal bodyspaces may form simply because they are the most effective representation to ensure an adaptive and flexible manipulation of the outside environment by multiple means—choosing the limb or extremity currently most suitable and available for the task at hand.

The paper is structured as follows. We first provide background on ILT. We then detail the experiment conducted in our laboratory that shows learning transfer from the trained to the untrained hand in an aiming task. Next, we detail the NN architectures used and show when learning transfer is possible within the architectures. The paper ends with a final discussion, summary, and conclusions.

II. INTERMANUAL LEARNING TRANSFER

Investigators have used a diversity of tasks to demonstrate ILT. One of the most typical ones is pointing under transformed visual feedback [4]. Despite high experimental effort, the mechanisms underlying ILT are not yet fully understood.

A. ILT Mechanisms

To perform goal-directed movements, a number of problems have to be solved A spatial target has to be selected, usually based on visual information. Additionally, a final arm posture needs to be selected that is sufficiently close to the target [5]. To do so, the target location has to be transformed from spatial, vision-based to effector-based coordinates. The problem of learning such a transformation is often termed an *inverse kinematics problem*. Other factors including trajectory formation, motor command generation, obstacle avoidance, hand posture selection, and the consideration of other environmental task constraints will not be further addressed in this study.

It is yet unclear which of these processes contributes to which degree to ILT. Some researchers suggest that the learned pattern of motor commands is transferred [6]. For example, it was found that right-handers were more accurate in drawing meaningless figures with the right hand, if they had previously practiced the drawing of mirror-reversed versions of the same figures with the left hand [7].

Other researchers have claimed that ILT occurs because the actor learns a new mapping of vision-based target coordinates onto effector-based target coordinates, which can be generalized to the contralateral limb [8], [9]. Recent neurological and behavioral data suggest that signals from a variety of

modalities, including vision, proprioception, audition, and vestibular sensation are combined to form modular spatial representations in the posterior parietal cortex [10], [11]. These effector-independent spatial representations may be used to affect different effector systems and may consequently constitute the basis of ILT.

Support for this hypothesis can be found in a study in which participants had to learn an artificial transformation of elbow- and shoulder-joint angles during an aiming task. Thereby, also the untrained arm adapted to the transformation [9]. According to the authors, the adjustments had been made to a central representation not specific to either arm but common to both arms. Other spatial mapping alterations can by found in the literature [12], [13], which all confirm high behavioral plasticity in humans. Humans are able to quickly adapt to spatial alterations, whereby the adaptivity is not confined to a single extremity but appears to be transferable to any available limb and even to tool usage [11].

To summarize, profit from prior contralateral training has been demonstrated for mirror-reversed as well as for non-reversed versions of the originally practiced movements. The results suggest that ILT occurrence and type are mainly dependent on the task-dependent goal representations activated during training. The high relevance of activated goal representations was also shown in bimanual behavioral studies in which object manipulations [14] showed congruency effects depend on task-dependent goal representations but not on the symmetry of the actually executed movements. Also the efficient execution of bimanual circular motion patterns was shown to depend on non-conflicting goal representations [15]. However, it remains unclear to what degree transfer depends on particular body-, task-, training-, and representational constraints. In the remainder of this work, we focus on ILT based on coordinate re-mappings. The following study evaluates the capability of humans to remap coordinates and transfer the remapping in an aiming task.

B. ILT in an Aiming Task

In an experiment by Lenhard *et al.* [16], [17] participants exhibited ILT in an aiming task. The goal of the study was to further investigate whether ILT is based on a re-interpretation of the visually perceived target locations or on a transfer of muscle-specific motor commands.

1) Experimental Setup: An experimental setup was used in which visual feedback could be displaced for individual targets (cf. Figure 1). Participants had no visual feedback of their own hand but rather saw a virtual spot representing the position of the hand. An array of nine squares distributed in a 3 grid depicted potential targets. Each critical movement started from a square in the middle column and aimed at one of the neighboring squares. For one single target the feedback was manipulated. The virtual spot indicated that movements towards this target were shorter than the actual movements. Accordingly, the visual feedback signaled a hit when the real position of the hand was beyond the target square.

To determine the degree of ILT, performance of both hands was tested before and after adaptation. Hereby, visual
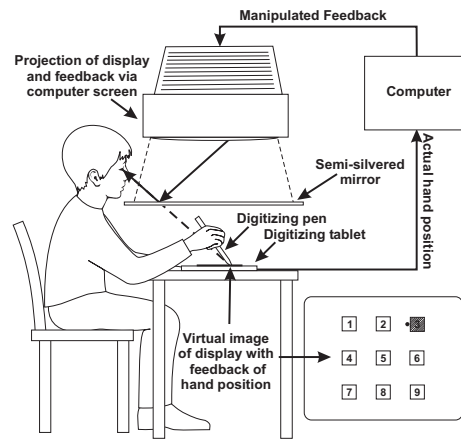

Fig. 1. Participants could not see their hand. Instead, a virtual spot indicated the current hand location. Neighboring targets had to be reached starting from one of the three locations in the center column. The visual feedback for one of the outside targets was manipulated so that the displayed hand location did not correspond with its actual one.

feedback was given which always indicated a hit for any aiming movement that ended outside of the starting square. Consequently, error-based learning was impossible during testing.

If ILT in this task was based on a re-interpretation of the visually perceived location of the manipulated target, movements of the untrained hand to the target should shifted in the same direction as movements of the trained hand. Thus, if the trained hand has learned to make farther movements in the direction of the manipulated target, the untrained hand should also show farther movements towards that target. If, however, the ILT was based more on a directional representations, a mirror-reversed version of the adapted movements should be observed. Besides these two possible directions of adaptation, also the generalization of the adaptation to targets in the near vicinity was of interest.

2) Results: After adaptation training, the results of participants that were not aware of the spatial manipulation (based on a subsequent questionnaire) clearly indicated that the trained hand adapted to the altered location. Also generalization to the neighboring targets in the same movement direction was observable, aiming at them with slightly farther movements. The adaption decreased with target distance to the manipulated target and did not generalize to opposite movement directions. Figure 3 (bottom bars for each target) shows the adaptation in arm movements.

Test trials with the untrained hand revealed ILT. In fact, no significant difference between the right and the left hand adaptation was found. ILT took place in various conditions with respect to location of the manipulated target, but not in the opposite movement direction. Therefore, a transfer of motor commands to homologous muscles of the other arm could be ruled out as an explanation for the observed learning transfer. Rather, it appeared that a target location re-interpretation was used to control the aiming movements of both hands.

III. SIMULATION OF ILT

Despite an extensive literature search, we were not able to find a neural network (NN) model that simulated ILT. The aim of the following NN architectures is to model the transfer observed in the human participants with a simple multi-layered NN, which is trained using backpropagation.

A. Training and Testing Data

The NN architectures tested were trained to learn the coordinate transformations from eye-orientation-based target coordinates to joint angle coordinates of the arm, which yield a hand position that coincides with the target coordinates. To generate suitable training data, the two-dimensional body model shown in Figure 2 was used, which included two arms and the eyes fixating a particular target. Hereby, the upper arm and forearm were assumed to be 26cm and 30cm long, respectively. Eye distance was set to 6.4cm and shoulder distance to 34cm. These distances correspond on average to the lengths of an adult human, where the forearm length includes part of the hand.

In an initial long training phase, ten NNs were trained with different random seeds for each architecture with a learning rate $\delta = .1$. The NNs learned the inverse kinematics of the two arms, that is, they learned to transform eye angles, which encode target locations, to arm angles, which determine the corresponding hand location. To do so, 1,000 data pairs of two eye angle pairs $((\alpha_l, \beta_l), (\alpha_r, \beta_r))$ and corresponding joint angle pairs $((\gamma_l, \delta_l), (\gamma_r, \delta_r)$, see Figure 2) were created. Both hands always had to reach the same point in space. The locations were uniformly randomly generated within a $20cm \times 20cm$ area distant $10cm$ in front of the eyes. All networks were trained for 10^7 iterations, sampling the data uniformly randomly with replacement.

In a second much shorter adaptation phase with only 10^4 iterations, the NNs were trained on the distorted workspace task used in the experiments with humans. In this phase, only nine locations were considered distributed in a 3×3 grid. Rows and targets within a row were distant 3.2cm. Corresponding to the experiments with humans, the location of the upper left target was distorted by associating to the eye
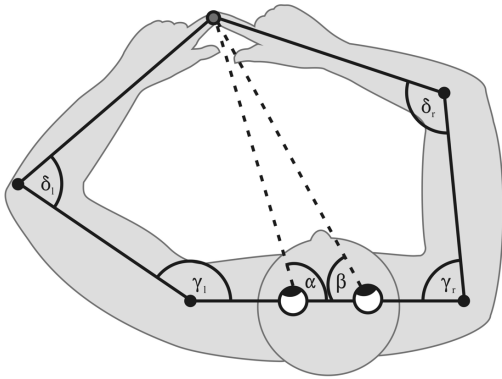

Fig. 2. To mimic the experimental setup with humans, a spatial model was used from which the training and testing data was generated.

angles joint angles that reached 1.96cm further to the left. Furthermore, the middle column was sampled twice as often to mimic the experiments with humans even closer. During this adaptation phase, only one arm received feedback (γ_r, δ_r), whereas the other received no feedback.

Finally, network adaption was tested feeding in the nine eye angle locations that corresponded to the nine target locations, recording the joint angles and corresponding hand locations generated by the NNs.

B. A Multilayer Feed-forward Neural Network

The neural network architectures were generated using the T-learn software, publicly available online [18]. Figure 4 shows the initial multilayer feed-forward network used in the initial experiments (NN 1). The network was constructed with nine hidden layers to be able to mimic multiple coordinate transformations and investigate effects of bottleneck location. All hidden layers consisted of six neurons except for the sixth layer, which consisted of four units, creating a processing bottleneck. The bottleneck further enforces compact encodings, which is in accordance to the dimensionality reduction technique recently shown to efficiently and compactly encode multidimensional data in neural networks [19]. All neurons had standard sigmoid activations [18], except for the last hidden layer and the output layer, which had linear activation patterns to be able to learn unnormalized radian measures. The input layer consisted of the four input values (α_l, β_l, α_r, β_r), to encode two parallel but separate targets, and a bias neuron, whose activity was set to one throughout the simulations.

After the 10^7 initial training iterations, performance tests of the ten independently trained NN 1 networks on the training data yielded a roots mean square error of RMS=.0088± .0029 in joint angle coordinates. This confirms that the NN 1 networks accurately learned to transform two pairs of identical eye angle coordinates into two separate pairs of joint angles.

After the initial training, the adaption phase was applied to each of the ten networks. Figure 3A (top bars for each target) shows the adaptation in horizontal direction after 10^4 adaption steps. The trained arm exhibits a very effective adaptation moving $1.01 \pm .348$ further to the left than necessary (target 1). Also joint angles for the neighboring targets (target 2 and target 4) exhibit significant adaptation. Thus, human data was mimicked which exhibited a generalized adaptation to target 4 but not to target 2 (Figure 3A bottom bars), because the center column was the start column [16].

However, when testing the untrained arm, no systematic adaptation was observable (Figure 3B top bars). The large standard deviations indicate that the untrained arm was affected by network crosstalk, but not in any directed way. This shows that NN 1 itself is not able to account for ILT.

In order to evaluate if an effector-independent representation evolves in NN 1 that can account for the transfer observed in humans, we fixed all the weights of the connections from the output layer until the bottleneck layer. Doing so, errors encountered during the adaption phase were back-propagated

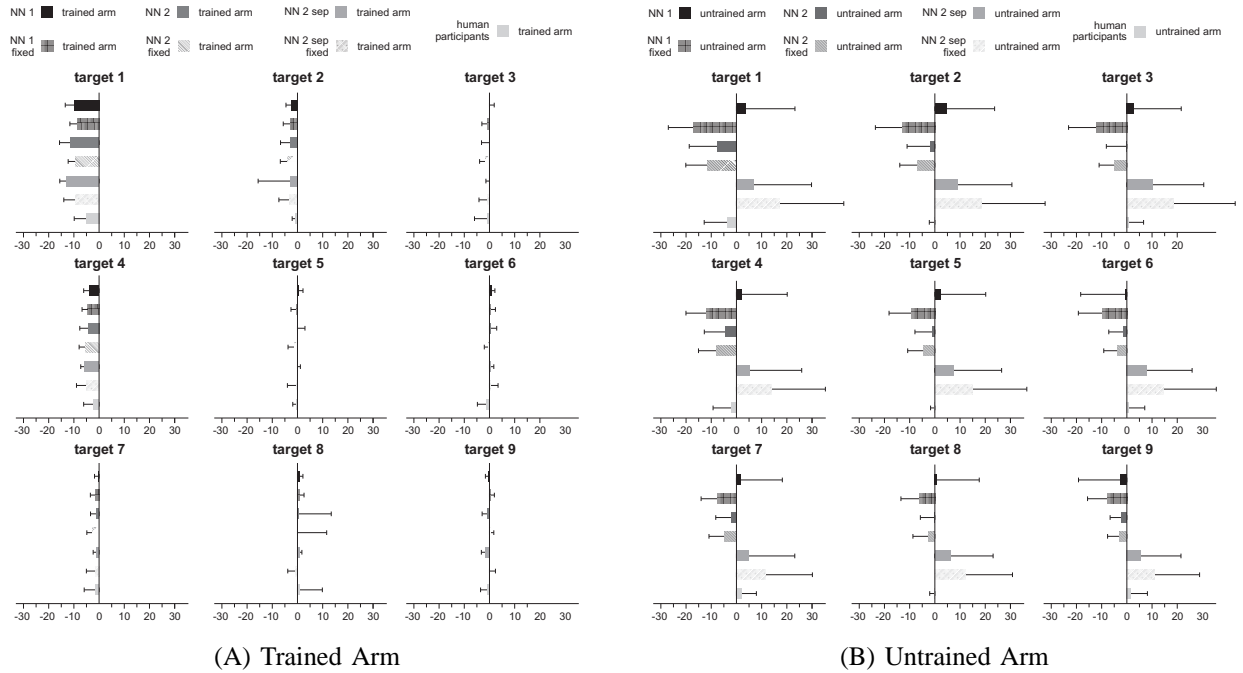(A) Trained Arm   (B) Untrained Arm

Fig. 3. The figure shows horizontal deviations of the hand from the nine visually perceived target locations (in millimeters). During the adaptation phase, target 1 was re-located 19.6mm to the left. In all NN settings, the trained arm adapted to the spatial distortion, exhibiting different amounts of generalization to the neighboring targets (A). The adaptation of the untrained arm (B) strongly depends on the NN architecture (NN 1 vs. NN 2), if weight adaptation was unconstrained or restricted to the layers before the bottleneck layer (free vs. fixed), and if the NN was trained on common goals or separate goals only (NN 2 vs. NN 2 sep). Human data (lowest bar in all target plots) shows typical adaption of the trained arm (A) and ILT in the untrained arm (B).

until the bottleneck layer without weight adaptation. From the bottleneck layer on then, the errors were back-propagated further and weights were adjusted until the input layer.

Figure 3A shows slightly smaller adaptations of the trained arm compared to the adaptations with unconstrained weight adaptation. However, now also the untrained arm exhibits strong adaptation tendency towards the left (Figure 3B). Thus, transfer is taking place. The transfer is not locally confined to target 1, though, but it extends to all target locations. This indicates that earlier hidden layers in NN 1 commonly represent the surrounding space but transformation effects have a global rather than a local impact. Thus, ILT is not well replicated.

C. Emergent Learning of an Effector-Independent Space

The results so far indicate that in NN 1 the other arm is affected by the adaptation, but rather in an unspecific fashion. Thus, the spatial representation of the two arms somewhat interacts within the hidden layers, but the interaction is not locally confined.

In neuroscience and psychology, though, it is well know that internal body spaces exhibit local encodings. Moreover, effector specific cortical areas can be distinguished from effector unspecific, pre-motor areas [11]. Thus, we generated a Y-shaped neural network, NN 2, that processed the eye angles in several common hidden layers and splits into two parts after the bottleneck layer (see Figure 4). We expected that the coordinate transformation from effector-independent representations to the two pairs of joint angles is accomplis-
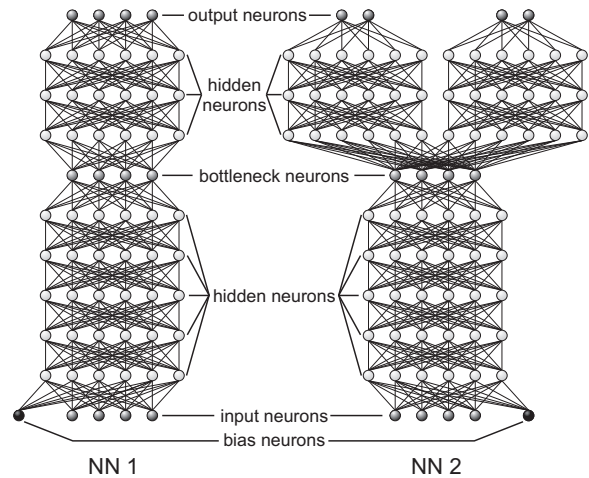

Fig. 4. A multilayer neural network architecture (NN 1) and a Y-shaped neural architecture was used in the NN experiments. A bottleneck layer enforced compact encodings.

hed after the bottleneck layer. Before the bottleneck layer, the eye angles may be suitably transformed into a maximally compact internal representation, which is enforced by the bottleneck layer.

To test the emergent plasticity of the resulting networks, we again trained 10 sets of NN 2 weights, which yielded a performance of $RMS = .0067 \pm .0015$ on the training data (in joint angle coordinates). After the adaptation phase and as in NN 1, the trained hand shows suitable adaptation

to the transformed target 1 as well as to the surrounding targets (Figure 3A). Even more interestingly, though, also the untrained arm adapts locally to the transformation. The adaptation of targets 1, 2, and 4 show adaptations to the left comparable to the trained arm and in line with the human data. Moreover, adaptations to the other targets differ from the actual correct position only slightly. Thus, local ILT occurs, similar to the one observed in humans.

Despite this very promising result, the adaptation in the y-direction shows that the transfer is not confined to the horizontal direction but also affects the vertical axis. The transformation in the vertical direction shows an oblique, negative distortion away from target 1. Target 9 was undershot with $-17.6 \pm 9.8mm$ (full data not shown). In NN 1, these distortions were not as extreme—a maximum displacement of $4.73 \pm 3.96mm$ was observe for target 3. Thus, the horizontal direction exhibits ILT but additional crosstalk disrupts adaptation in the vertical direction.

To investigate whether the adaptation in the trained arm before the bottleneck layer caused this effect, we again fixed the weights between the output and bottleneck layers during adaptation. Figure 3 shows that the displacements are comparable to the ones without fixed weights, albeit the adaptation is slightly more global. In comparison to the adaption in NN 1, though, adaption remains much more locally confined. Moreover, the adaptation in the vertical direction is significantly decreased (maximum distortion of $-4.10 \pm 8.41mm$ for target 9). Thus, this setup most closely replicates the local ILT observed in humans.

D. Training with Different Goal Location Pairs

Recent insights from psychology furthermore suggest that generalization in motor learning strongly depends on the history of motor actions. For example, it was shown that knowledge transfer occurs from shoulder to wrist manipulations but not vice versa, presumably since the wrist also moves through space when the shoulder is moved, but not the other way round. The authors suggest that an interdependence between effectors needs to be present in order to enable knowledge transfer [20].

To see if knowledge transfer depends on the interdependence between arm movements in our model, we generated another training set in which each arm had to reach a different target. Thus, the α_l, β_l pair for the left arm now differed from the α_r, β_r pair. Again, 1,000 data entries were generated and ten NN 2 networks were trained on the data for 10^7 steps with different random seeds. As expected, the task was harder for the networks so that one out of the ten networks did not succeed in learning the transformation accurately. We consequently replaced that network with an additional network trained with an additional random seed. Performance evaluations on the training data yielded an RMS of $.0186 \pm .0030$ in joint angle coordinates in the ten successfully trained networks.

Figure 3A shows that also this network (NN 2 sep) is able to adapt to the transformation locally with the trained arm. However, the untrained arm does not show any systematic transformation, as indicated by the large standard deviations and the displacement in the opposite direction (Figure 3B). Thus, common, interdependent training of the two arms seems necessary to form emergent internal representations that support ILT.

IV. DISCUSSION

Our neural network simulations have shown that internal NN representations that are suitable for ILT only emerge under the following conditions. (1) The two simulated effectors need to be trained interdependently. In our simulations, this was accomplished by the training of common goals during training. (2) The network architecture needs to enforce independent processing of the inverse kinematics of the two arms, while relying on an internally hidden common code. This was accomplished by a Y-shaped NN architecture (NN 2 in Figure 4). (3) If the weight adaptation is back-propagated to the common layers and does not affect the separate layers, the ILT is even more similar to the one observed in humans. This suggests that the internal adaptation of an effector-independent internal common code yields the most accurate effects.

Clearly there is lots of room to improve the applied neural architecture and the utilized data. Research suggests that internal body spaces are encoded with population codes with various, body-related topologies [11]. Moreover, bidirectional, Bayesian-like processing occurs in the brain [21]. An interesting challenge lies in the exploitation of this knowledge to accomplish even more accurate and reliable ILT in artificial learning systems.

A. Population-Encoded Body Spaces

Recent neuro- and cognitive psychological investigations have shown that the brain encodes the surrounding space in various, multimodal, body-related spatial representations [11], [22], [23]. For example, cortical areas were identified in monkeys that encode a *peripersonal* arm space. A neuron in this area may respond if the arm is touched at a certain location, if an object is close to that location, and even if the position is simply looked at [22].

These spaces are mostly encoded by population codes, that is, neural representations in which each neuron has a local receptive field in the encoded (body-) space. Integrating neural activity by means of single cell recordings in monkeys, it is now even possible to derive current arm posture and current arm movements from this activity [23]. Also psychological experiments indicate that partially effector-independent internal neural representations exist [11]. As suggested, these effector-independent spatial representations might be used to accomplish re-mappings and some types of ILT observed in humans [9], [8], [12], [13].

Our simulations focused on re-mappings from eye to joint angles. The currently applied networks, however, did not apply population encodings. Due to the sparsity of all layers, and the bottleneck layer in particular, it can also not be expected that the neurons in the layers exhibit local receptive field properties. Thus, a future research challenge will be

to account for ILT by means of NNs based on population encodings.

Nonetheless, our simulations indicated that fundamental learning biases are necessary to evolve internal codes suitable to accomplish ILT. Thus, it can be expected that also population encoded internal spatial representations will depend on these factors.

B. Multimodular Bayesian Processing

Besides the population encoding, it was shown that body spaces are highly modular and are present in various modalities. Moreover, the spaces are highly interactive integrating available sensory, motor, and sensorimotor information [11], [23].

To be able to integrate multiple sources of information, Bayesian-based processing mechanisms seem necessary [21]. Information should be integrated dependent on its current reliability. For example, it was shown that arm representations in the motor cortex mainly rely on proprioceptive feedback, whereas the ventral premotor cortex primarily represents visual information in a distorted movement task [23]. Thus, neuron activity depends on information reliability, availability, but also on intention.

Clearly, such higher level processing mechanisms are not possible in the investigated NN architectures. For example, bidirectional processing is not possible with a feedforward neural network. Nonetheless, the Bayesian brain hypothesis and the reliance on different sensory and motor feedback in different brain areas indicates that adaptations in the brain can occur in different locations dependent on information reliability. Thus, fixing the weights in the layers close to motor output and propagating encountered spatial distortions to deeper neural layers is well in line with neurological findings [23].

C. Conclusions

Although the NNs used to simulate ILT are certainly far too simple to account for actual neural processing mechanisms in the brain, the findings identified learning biases necessary to shape the internal neural network structure. These biases might need to be integrated in other ways into other neural structures, such as population-encoded structures. Nonetheless, the results suggest that interdependence between multiple effectors, interdependencies inside the neural architectures, and effective neural adaptation within will be the key to shape emergent internal body spaces suitable for ILT. Future research will show which other constraints are necessary to generate such spaces reliably in more neurologically plausible learning architectures.

Acknowledgments

References

[1] A. W. Volkmann, "Über den Einfluss der Übung auf das Erkennen räumlicher Distanzen," *Berichte der sächsischen Gesellschaft der Wissenschaften. Mathematisch-physische Klasse*, vol. 10, p. 38, 1858.

[2] H. von Helmholtz, *Handbuch der physiologischen Optik*. Leipzig : Leopold Voss, 1867, vol. III.

[3] S. Haykin, *Neural Networks: A Comprehensive Foundation*. Upper Saddle River, NJ: Prentice Hall, 1999, 2nd edition.

[4] J. Wang and R. L. Sainburg, "Interlimb transfer of visuomotor rotations depends on handedness," *Experimental Brain Research*, vol. 175, pp. 223–230, 2006.

[5] D. A. Rosenbaum, L. D. Loukopoulos, R. G. J. Meulenbroek, J. Vaughan, and S. E. Engelbrecht, "Planning reaches by evaluating stored postures," *Psychological Review*, vol. 102, no. 1, pp. 28–67, 1995.

[6] R. E. Hicks, C. T. Gualtieri, and S. R. Schroeder, "Cognitive and motor components of bilateral transfer," *American Journal of Psychology*, vol. 96, pp. 223–228, 1983.

[7] G. Thut, N. D. Cook, M. Regard, K. L. Leenders, U. Halsband, and T. Landis, "Intermanual transfer of proximal and distal motor engrams in humans," *Experimental Brain Research*, vol. 108, pp. 321–327, 1996.

[8] M. L. Latash, "Mirror writing: Learning, transfer, and implications for internal inverse models," *Journal of Motor Behavior*, vol. 31, pp. 107–111, 1999.

[9] H. Imamizu, Y. Uno, and M. Kawato, "Adaptive internal model of intrinsic kinematics involved in learning an aiming task," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 24, pp. 812–829, 1998.

[10] R. A. Andersen, L. H. Snyder, D. C. Bradley, and J. Xing, "Multimodal representation of space in the posterior parietal cortex and its use in planning movements," *Annual Review of Neuroscience*, vol. 20, pp. 303–330, 1997.

[11] A. Maravita, C. Spence, and J. Driver, "Multisensory integration and the body schema: Close to hand and within reach," *Current Biology*, vol. 13, pp. 531–539, 2003.

[12] K. M. Mosier, R. A. Scheidt, S. Acosta, and F. A. Mussa-Ivaldi, "Remapping hand movements in a novel geometrical environment," *Journal of Neurophysiology*, vol. 94, pp. 4362–4372, 2005.

[13] R. Shadmehr and S. P. Wise, *The Computational Neurobiology of Reaching and Pointing: A foundation for motor learning*. Cambridge, MA: MIT Press, 2005.

[14] W. Kunde and M. Weigelt, "Goal congruency in bimanual object manipulation," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 31, pp. 145–156, 2005.

[15] F. Mechsner, D. Kerzel, G. Knoblich, and W. Prinz, "Perceptual basis of bimanual coordination," *Nature*, vol. 414, no. 6859, pp. 69–73, 2001.

[16] A. Lenhard, J. Hoffmann, and A. Sebald, "Intra- and intermanual transfer of adaptation to unnoticed virtual displacement under terminal and continuous visual feedback," 2002, diploma Thesis, University of Würzburg.

[17] A. Lenhard and J. Hoffmann, "Adaptation von Handbewegungen an targetspezifische virtuelle Zielverlagerungen," in *Beiträge zur 46. Tagung experimentell arbeitender Psychologen*, D. Kerzel, V. Franz, and K. Gegenfurther, Eds. Lengerich, Germany: Pabst Science Publishers, 2004, p. 156.

[18] K. Plunkett and J. L. Elman, *Exercises in Rethinking Innateness: A Handbook for Connectionist Simulations*. The MIT Press, 1997. [Online]. Available: http://crl.ucsd.edu/innate/index.shtml

[19] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, pp. 504 – 507, 2006.

[20] J. W. Krakauer, P. Mazzoni, A. Ghazizadeh, R. Ravindran, and R. Shadmehr, "Generalization of motor learning depends on the history of prior action," *PLoS Biology*, vol. 4, no. 10, 2006.

[21] S. Deneve and A. Pouget, "Bayesian multisensory integration and cross-modal spatial links," *Journal of Physiology - Paris*, vol. 98, pp. 249–258, 2004.

[22] G. Rizzolatti, L. Fadiga, L. Fogassi, and V. Gallese, "Enhanced: The space around us," *Science*, vol. 277, pp. 190–191, 1997.

[23] A. B. Schwartz, D. W. Moran, and G. A. Reina, "Differential representation of perception and action in the frontal cortex," *Science*, vol. 303, pp. 380–383, 2004.