Neural Networks 29–30 (2012) 60–69

Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

A computational model of use-dependent motor recovery following a stroke: Optimizing corticospinal activations via reinforcement learning can explain residual capacity and other strength recovery dynamics

David J. Reinkensmeyer [a,b,c,*], Emmanuel Guigon [d,e], Marc A. Maier [f,g]

[a] Department of Anatomy and Neurobiology, University of California at Irvine, USA
[b] Department of Mechanical and Aerospace Engineering, University of California at Irvine, USA
[c] Department of Biomedical Engineering, University of California at Irvine, USA
[d] UPMC Univ Paris 06, UMR 7222, ISIR, F-75005, Paris, France
[e] CNRS, UMR 7222, ISIR, F-75005, Paris, France
[f] CNRS, UMR 8194, Université Paris Descartes, Sorbonne Paris Cité, F-75006 Paris, France
[g] Univ Paris Diderot, Sorbonne Paris Cité, F-75205 Paris, France

ARTICLE INFO

Article history:
Received 1 October 2010
Received in revised form 31 January 2012
Accepted 7 February 2012

Keywords:
Stroke
Motor control
Reinforcement learning
Use-dependent plasticity

ABSTRACT

This paper describes a computational model of use-dependent recovery of movement strength following a stroke. The model frames the problem of strength recovery as that of learning appropriate activations of residual corticospinal neurons to their target motoneuronal pools. For example, for an agonist/antagonist muscle pair, we assume the motor system must learn to activate preserved agonist-exciting corticospinal neurons and deactivate preserved antagonist-exciting corticospinal neurons. The model incorporates a biologically plausible reinforcement learning algorithm for adjusting cell activation patterns – stochastic search – using generated limb force as the teaching signal to adjust the synaptic weights that determine cell activations. The model makes predictions consistent with clinical and brain imaging data, such as that patients can achieve an increase in strength after appearing to reach a recovery plateau (i.e., "residual capacity"), that the differential effect of a dose of movement practice will be greater earlier in recovery, and that force-related brain activation will increase in secondary motor areas following a stroke. An interesting prediction that could be explored clinically is that temporarily inhibiting subpopulations of more powerfully connected corticospinal neurons during late movement training will allow the motor system to optimize corticospinal neurons with a weaker influence, whose optimization was blocked by the rapid optimization of more strongly connected neurons early in training.

1. Introduction

Over 700,000 people experience a stroke in the US each year (Broderick et al., 1998). About 80% of acute stroke survivors experience hemiparesis, with this percentage decreasing to about 50% in chronic stroke (Gresham, Duncan, & Stason, 1995). Stroke patients typically undergo several months of rehabilitative movement training aimed at improving strength and coordination, but the neural mechanisms that promote motor recovery in response to movement practice are not well understood. There are currently intensive efforts to develop new stroke rehabilitation techniques (Langhorne, Coupar, & Pollock, 2009), including robotic and virtual reality approaches (Brochard, Robertson, Médée, & Rémy-Néris, 2010), but there is a lack of rigorous, theoretical frameworks to guide these efforts.

Developing mathematical models of stroke recovery could improve the understanding of stroke motor recovery and help guide the design of new therapies. Currently, however, only a few previous studies have used a computational approach to model motor function following stroke. The focus of these studies has primarily been on explaining changes in receptive fields or cortical maps following stroke (Goodall, Reggia, Chen, Ruppin, & Whitney, 1997; Lytton, Stark, Yamasaki, & Sober, 1999; Sober, Stark, Yamasaki, & Lytton, 1997) or on explaining kinematic features of movement impairment such as decreased smoothness and increased variability (Reinkensmeyer, Iobbi, Kahn, Kamper, & Takahashi, 2003; Rohrer et al., 2002) rather than on the dynamics of motor performance recovery. One recent exception is the model by Han, Arbib, and Schweighofer (2008), which modeled the

* Correspondence to: Department of Mechanical and Aerospace Engineering, 4200 Engineering Gateway, Irvine, CA 92697-3975, USA. Tel.: +1 949 824 5218; fax: +1 949 824 3535.
E-mail address: dreinken@uci.edu (D.J. Reinkensmeyer).

0893-6080/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2012.02.002

phenomenon of "learned non-use", in which a stroke patient will chose to not use their weakened limb because of difficulty in using it effectively; this non-use is hypothesized to result in further deterioration in the ability to control the limb. The Han et al. model used a population vector coding paradigm with simulated lesions to model directional errors in a center-out reaching task, and an action-choice module that learned by reinforcement learning the value of using each arm for reaching in different directions. They found that if motor training brought directional reaching errors below a threshold, then spontaneous use of the arm continued to train the arm, and they were able to model cortical reorganization as a redistribution of population vector directions due to the training. The model was supported by self-reports of functional activity from a large multi-site randomized controlled trial of constraint-induced therapy (Schweighofer, Han, Wolf, Arbib, & Winstein, 2009).

The Han et al. model focused on a cognitive issue—the choice whether or not to use the impaired arm, relating directional errors to functional use. The goal of the work described in this paper was to develop a computational model of motor recovery following stroke, and specifically, of the recovery of the ability to generate force using disrupted corticospinal pathways. We focused on the recovery of distal upper limb strength because strength strongly predicts upper extremity functional activity (Bohannon, Warren, & Cogman, 1991; Canning, Ada, Adams, & O'Dwyer, 2004; Harris & Eng, 2007; Lang, Wagner, Edwards, & Dromerick, 2007), and thus understanding its recovery might be expected to generalize to a wide range of functions. In addition, it was possible to base a model of force production on single cell neurophysiological studies of wrist force production in primates (Fetz & Cheney, 1980; Kasser & Cheney, 1985; Maier, Perlmutter, & Fetz, 1998; Mewes & Cheney, 1991; Perlmutter, Maier, & Fetz, 1998).

An initial goal in developing this model was to gain insight into the phenomenon of "residual capacity" or "functional potential", which refers to the finding that additional movement training can improve motor function, including strength, even years following a stroke (Page, Gater, & Bach-Y-Rita, 2004; Rijntjes, 2006; Stinear, 2007). Upper extremity motor recovery following a stroke reaches an apparent plateau in the first year after the initial incident by clinical (Duncan et al., 1994; Heller et al., 1987; Sunderland, Tinson, Bradley, & Hewer, 1989) and biomechanical measures (Mirbagheri, Tsao, & Rymer, 2008) (Fig. 1). However, there is now extensive evidence that the time course of recovery is not fixed, but rather that additional movement practice can enhance movement and strength in both the sub-acute and chronic phases following a stroke (Ada, Dorsch, & Canning, 2006; Barreca, Wolf, Fasoli, & Bohannon, 2003; French et al., 2007; Kwakkel, Kollen, & Krebs, 2008; Kwakkel et al., 2004; van der Lee et al., 2001) (Fig. 1). The effect size of additional movement practice is typically small (French et al., 2007; Kwakkel et al., 2004) leaving patients short of a full recovery, but additional recovery is statistically significant. Understanding the possible neural bases of residual capacity should provide an insight into how to improve recovery. We sought to gain an insight into these neural bases by modeling the recovery of strength as a reinforcement learning problem in which the limb force experienced on attempts to move the limb guides the refinement of activation in preserved corticospinal pathways. Portions of this work have been published in abstract form (Reinkensmeyer et al., 2009).

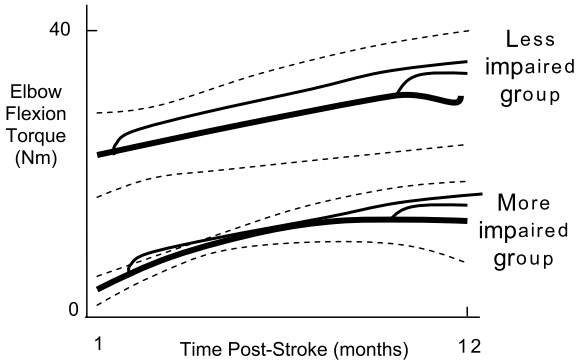

Fig. 1. Strength recovery and residual capacity following a stroke. The strength recovery curves (thick solid lines), which were identified by measuring maximum isometric elbow flexion torque from 20 stroke patients 5 times over a 12 month period, and grouping more and less impaired patients into two groups using a growth-mixture model technique (thick solid line = mean, dashed lines ±1 standard deviation (SD)). For reference, unimpaired age-matched maximum elbow flexion strength is about 80 N m (Dewald & Beer, 2001). The superimposed thin, solid lines show the hand-drawn, predicted effects of strength training in the acute and chronic phase following a stroke based on a systematic review of 14 strength training studies (Ada et al., 2006). In the Ada review, the average effect size for strengthening acute weak patients was 0.33 SD, and for chronic patients 0.18 SD. *Source:* The strength recovery curves (thick solid lines) are copied from Mirbagheri et al. (2008).

2. Methods

2.1. Model description

The model presented here is intended to model strokes that cause weakness by damaging motor areas that give rise to descending white matter tracts. The model ignores muscle plasticity since it has been shown that electrical stimulation of muscle can produce near normal limb forces after chronic stroke, indicating that muscle atrophy is not the main cause of weakness (Landau & Sahrmann, 2002). We apply the model to the task of activating motor networks to generate a flexion force with the wrist, a task commonly used in primate neurophysiological studies. We propose that the motor system must learn appropriate activation of residual corticospinal (CS) cells from movement practice. Other potential mechanisms for recovery after a stroke include structural changes in dendrites and dendritic trees, activation of neural stem cells, and changes in the extracellular matrix (Cramer, 2008), but we focus here on the following question: "To what extent can the dynamics of stroke motor recovery be explained by the process of optimizing activity in residual, fixed pathways to motoneurons, based on experience of movement practice?" The model is characterized by two key features.

2.1.1. Model feature 1: summed activity from corticospinal cells determines muscle force

The first key feature is that the force the wrist muscles generate is determined by the weighted, summed activity of CS cells that activate the wrist motoneurons (MNs) (Fig. 2). The influence of both mono-synaptically connected systems (i.e. corticomotoneuronal CM cells) and multi-synaptically connected systems are captured by fixed, functional connectivity weights c_{fi} and c_{ei} that represent the net excitatory or inhibitory effect of the cell on the MN pool. We modeled the distribution of these weights based on the primary motor cortex CM system because it is likely the most functionally important of the premotor systems for humans, and has been well characterized using spike-triggered averaging techniques in primates (Fetz & Cheney, 1980; Kasser & Cheney, 1985). Most CM cells (99%) originating from primary motor cortex facilitate either the flexor or extensor muscle groups for wrist flexion/extension movement; about a third or less of these cells at the same time inhibit the antagonist (see Fig. 17 in Perlmutter et al. (1998)). Changing the percentage of reciprocal or inhibitory connections did not alter the basic findings from the model, except that when no inhibitory cells were included, the levels of simulated co-contraction tended to be higher. Note that a basic feature of the model is