Ergonomics
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/terg20

A neural network model for predicting postures during non-repetitive manual materials handling tasks

Miguel A. Perez [a] & Maury A. Nussbaum [b c]

[a] Center for Automotive Safety Research, Virginia Tech Transportation Institute , 3500 Transportation Research Plaza, Blacksburg, VA, 24061, USA

[b] Grado Department of Industrial & Systems Engineering , Virginia Tech , 250 Durham Hall, Blacksburg, VA, 24061, USA

[c] School of Biomedical Engineering & Sciences, Virginia Tech/Wake Forest University , 114 Randolph Hall, Blacksburg, VA, 24061, USA

Published online: 19 Sep 2008.

To cite this article: Miguel A. Perez & Maury A. Nussbaum (2008) A neural network model for predicting postures during non-repetitive manual materials handling tasks, Ergonomics, 51:10, 1549-1564, DOI: 10.1080/00140130802220570

To link to this article: http://dx.doi.org/10.1080/00140130802220570

Ergonomics
Vol. 51, No. 10, October 2008, 1549–1564

Taylor & Francis
Taylor & Francis Group

A neural network model for predicting postures during non-repetitive manual materials handling tasks

Miguel A. Perez[a]* and Maury A. Nussbaum[b,c]

[a]*Center for Automotive Safety Research, Virginia Tech Transportation Institute, 3500 Transportation Research Plaza, Blacksburg, VA 24061, USA;* [b]*Grado Department of Industrial & Systems Engineering, Virginia Tech, 250 Durham Hall, Blacksburg, VA 24061, USA;* [c]*School of Biomedical Engineering & Sciences, Virginia Tech/Wake Forest University, 114 Randolph Hall, Blacksburg, VA 24061, USA*

Posture prediction can be useful in facilitating the design and evaluation processes for manual materials handling tasks. This study evaluates the ability of artificial neural network models to predict initial and final lifting postures in 2-D and 3-D scenarios. Descriptors for the participant and condition of interest were input to the models; outputs consisted of posture-defining joint angles. Models were trained with subsets of an existing posture database before predictions were generated. Trained models predictions were then evaluated using the remaining data, which included conditions not presented during training. Prediction errors were consistent across these data subsets, suggesting the models generalised well to novel conditions. The models generally predicted whole-body postures with per-joint errors in the 5°–20° range, though some errors were larger, particularly for 3-D conditions. These models provided reasonably accurate predictions, even outperforming some computational approaches previously proposed for similar purposes. Suggestions for future refinement of such models are presented. The models in this investigation provide a means to predict initial and final postures in commonly occurring manual materials handling tasks. In addition, the model structures provide information about potential lifting strategies that may be used by individuals with particular anthropometry or strength characteristics.

Keywords: posture prediction; artificial neural networks; manual materials handling; simulation

1. Introduction

Diverse factors have been identified as important when assessing the risk of physical injury due to occupational tasks. Among these are static and/or non-neutral postures, which appear to be especially important for lifting exertions (Hsiang *et al*. 1997, National Institute for Occupational Safety and Health 1997, Keyserling 2000). This link between lifting posture and injury risk makes lifting posture a critical component of manual materials handling job analyses. This focus on posture is also reflected in the number and diversity of analytical techniques and software programs to assess lifting (Zhao and Badler 1994, Mirka *et al*. 2000, Russell *et al*. 2007). Job analyses are often based on existing tasks, in which case postures can be directly observed or measured. However, when tasks are being designed, or interventions for them being developed, relevant posture data are often

*Corresponding author. Email: mperez@vt.edu

ISSN 0014-0139 print/ISSN 1366-5847 online

DOI: 10.1080/00140130802220570
http://www.informaworld.com

not available and must be predicted. An ability to predict postures thus has potential benefits (Chaffin 2005), both in obviating the need for prototyping of changes to existing or simulated tasks and allowing for rapid and cost-effective examining of alternatives.

A variety of techniques have been employed for modelling of lifting or reaching motions, such as control theory, inverse kinematics, optimisation and fuzzy logic (Beck and Chaffin 1992, Dysart and Woldstad 1994, Jung and Park 1997, Faraway 2000, Zhang and Chaffin 2000). Reasonable levels of success have been achieved with these approaches, in terms of the errors between measured and predicted body segment configurations. Typically, these approaches model lifting or reaching motions as a series of instantaneous postures, often dependent on previous model steps, and success in prediction of initial and final postures may be assessed. However, these models are not explicitly designed to predict initial and final postures, but rather an overall motion trajectory, and their fidelity in simulating initial and final postures is usually not explicitly evaluated. Furthermore, some of these approaches need pre-defined boundary conditions, requiring that initial and/or final posture be provided as a model input instead of being an output. Therefore, this investigation proposes a model that is judged solely by its prediction of initial and final lifting posture, and independently of motion trajectory, which is not predicted.

Artificial neural networks (ANNs), which use highly interconnected nodes to 'learn' relationships and characteristics of data, and extrapolate these relationships to novel situations, have the potential to be useful tools for posture prediction by addressing some of the drawbacks in previous investigations. Evidence exists supporting manipulation of human motion control by collections of structures similar to ANNs (Windhorst 1996). While a number of investigators have used ANNs to predict motion (Jung and Park 1994, Lim *et al*. 1996, Bellan *et al*. 2000), posture prediction is a secondary consideration of these models; initial and/or final postures are either provided to the model as inputs, implicitly coded within the model parameters, or predicted only for some segments (e.g. arms).

This investigation involved relatively simple ANN-based models to predict full-body posture at the beginning and end of a lift in two and three dimensions. The models used hand locations and participant-specific coefficients (describing strength and anthropometry) as the only inputs. A description of the model is provided along with results of tests of predictive ability using an existing human posture database. Model results are also compared to those of past reports and predictions obtained from a commonly used commercial software package.

2. Methods

2.1. *Dataset*

The postures and motions used here were obtained from The University of Michigan's HUMOSIM project site (www.humosim.org). This site serves as a repository of posture and motion capture data. Methods for obtaining these data are briefly summarised here and the reader is referred to the HUMOSIM website for more detailed descriptions of the numerous studies that are included.

Data used here were taken from a study that included 10 male and 10 female volunteers. All were healthy and their ages ranged from 21 to 68 years. In total, 41 anthropometric and strength characteristics were measured. Participants performed 213 motion trials in two separate 2–3 h sessions. The trials required four separate motions: 1) beginning from a standardised 'home' position, deliver an object to a shelf; 2) return hands back to 'home' position; 3) go get the object; 4) bring the object back to the 'home' position. Only movements 1 ('deliver') and 4 ('bring back') were analysed here, as they required a load transfer.

Lift style and velocity were self-selected. To avoid fatigue effects, sufficient rest was provided between trials. The load centre was shifted at the beginning of each trial so that the distance to the load was maintained constant across participants; the distance was defined and measured with regard to the midpoint between the hips (H-point). The participant's left foot was not allowed to move. Trial targets were set at near and far distances, varied in height (five levels: lower; lower middle; middle; upper middle; upper) and were distributed around the participant (three levels: 0°, 45°, 90° rotation toward their right-hand side). Further details on the locations of these targets and trial constraints can be found in Perez and Nussbaum (2006) and the HUMOSIM project website.

During the trials, kinematic data (i.e. marker positions and body segment orientations, when applicable) were obtained from skin-mounted position tracking devices (optical and electromagnetic) and converted to joint locations (Park *et al.* 1999). Prior to subsequent analysis, motions with missing or inconsistent data were corrected or eliminated (as described in Perez and Nussbaum 2006). To obtain 2-D data from the 3-D joint coordinates, symmetrical joint positions were averaged on the sagittal plane (e.g. left knee and right knee 2-D coordinates to obtain a 2-D knee position vector). This approach was used in lieu of selecting only coordinates on one side of the body to reduce the effects of noise and non-sagittally symmetrical positioning on the derived 2-D dataset. Note that to ensure symmetry, this manipulation was only possible for the subset of lifting motions that nominally required 0° trunk rotation. These modified joint coordinates were then used to obtain segmental angles for the ankle, knee, hip, lower and upper torso, neck, clavicle, shoulder, elbow and wrist.

A commonly used and robust Euler angle sequence (Crawford *et al.* 1996) was used to convert Cartesian joint coordinates from the dataset into 3-D segment angles (i.e. the orientation of a segment relative to a global coordinate system) via the calculation of simple trigonometric relationships. This angle sequence is fully described by three rotations around three axes, which were defined based on the human-centric direction that they represented and denoted by α (around the inferior-superior axis), β (around the rotated anterior-posterior axis) and γ (around the rotated left–right axis). However, since the available joint and marker coordinates did not allow for estimation of axial rotation, only the first two Euler angles (α and β) were used to represent each of the segments (i.e. $\gamma = 0$ for all conditions). Detailed descriptions on how these calculations were performed can be found in Perez and Nussbaum (2006).

2.2. *Artificial neural network model*

Standard feed-forward neural networks were used to predict initial and final postures (Haykin 1999). The 2-D network had seven inputs, 10 outputs and a single hidden layer with 50 nodes. The corresponding 3-D network had eight inputs, 36 outputs and 100 hidden layer nodes (Figure 1). The additional input in the 3-D network indicated left–right location of the target, while the additional outputs represent the angles needed to characterise a 3-D posture compared to 2-D, and the additional hidden layer nodes were required to account for the increased complexity of the resulting 3-D network. Biases in the hidden and output layers were trained. A sigmoid transfer function was used in the hidden layer, while a linear transfer function was used for the output units. These networks were developed and trained (see below) using the Matlab (Mathworks, Natick, MA, USA) Neural Networks toolbox.

The ANNs were provided with information about the hand positions for each initial or final lifting posture being analysed. Five separate coefficients were also used as inputs to

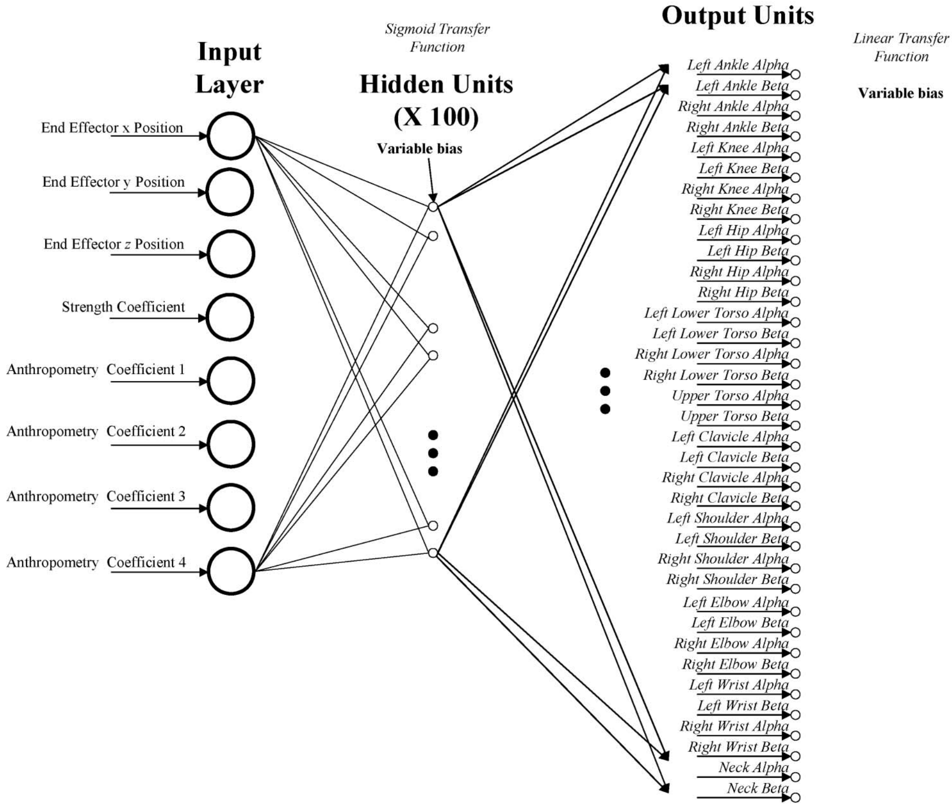

Figure 1. Network structure used for 3-D posture prediction.

make the ANNs adaptable to the characteristics of different individuals. These coefficients were derived using principal components (PC) regression (Johnson 1998) on the available anthropometric and strength characteristics of the participants. Separate strength and anthropometry PC regressions were run. A single component was chosen from the PC strength regression and four were chosen from the anthropometry regression (Tables 1 and 2, respectively). The number of components was selected by iteratively including the most predictive components (based on percent variance explained) until their contribution to model fidelity became marginal. Once the eigenvectors were determined, regressions could be run that converted each participant's strength and anthropometry characteristics into a unique and reduced set of strength and anthropometry coefficients, which could be input to the model.

The 2-D training dataset consisted of input vectors containing data for 80% of the participants (randomly selected), with the remaining 20% used to determine how well the trained network generalised (i.e. could make predictions) to a set of novel participants. These proportions were selected based on experimenter judgement, with the goal of providing enough training data to the network while saving sufficient data for testing generalisation to novel inputs. Standard error back-propagation was used to train the network (Werbos 1988). This process consists of an iterative presentation of input vectors, determination of predictions and prediction errors and modification of internal and

Table 1. Principal components analysis eigenvector used to generate the strength coefficient.

Variable	Eigenvector
Experimental object weight (N)	0.462
Mean shoulder flexion strength (N)	0.463
Mean shoulder abduction strength (N)	0.462
Mean shoulder 45° abduction strength (N)	0.453
Mean torso extension strength (N)	0.391

Table 2. Partial principal components analysis eigenvectors used to generate the anthropometry coefficients.

Variable	Eigenvector 1	Eigenvector 2	Eigenvector 3	Eigenvector 4
Weight (kg)	0.141	0.073	0.264	0.247
Stature without shoes (cm)	0.216	0.034	−0.035	0.038
Stature with shoes (cm)	0.216	0.009	−0.049	−0.047
C7 height (cm)	0.217	0.006	−0.052	−0.050
Suprasternale height (cm)	0.217	0.022	−0.045	0.016
L5 height (cm)	0.200	0.049	−0.078	−0.060
Greater trochanter height (cm)	0.182	0.033	−0.157	−0.091
Knee height (cm)	0.162	0.057	0.070	0.020
Malleolus height (cm)	0.042	0.088	−0.283	−0.195
Seated height (cm)	0.201	0.029	−0.049	−0.038
Seated eye height (cm)	0.192	0.053	−0.007	0.052
Seated C7 height (cm)	0.202	0.074	0.016	0.034
Seated ankle horizontal distance (cm)	0.109	−0.182	−0.096	0.129
Seated ankle vertical distance (cm)	−0.148	−0.178	−0.082	0.169
Head width (cm)	0.065	0.257	0.147	0.086
Head depth (cm)	0.188	0.016	0.043	0.183

Note: The complete eigenvectors included 45 variables.

structural network parameters with the goal of reducing prediction errors for the next iteration in the training process. For the networks in the current investigation, a maximum of 500,000 consecutive iterations (or epochs) of the training dataset were used. The training process was stopped if the mean square error (MSE) went below 0.003 rad^2 or stabilised (i.e. did not change by more than 10% within 20,000 epochs). The same training criteria were applied across the different networks. Once the training process was stopped, the internal and structural network parameters were held constant and network predictions calculated by presenting any desired input vectors and obtaining a corresponding set of predicted outputs. These predictions were compared to the empirically measured outputs to determine prediction errors.

A slightly different training dataset was used for the 3-D ANN. For the 2-D ANN, all available data from each of the participants included in the training dataset were used, resulting in one training dataset and one generalisation dataset. As such, generalisation was tested ‘across’ participants. The 2-D training dataset also included the complete range of initial and final hand locations. Based on the favourable results obtained using this approach (see Results), the 3-D ANN training dataset was further constrained by purposefully excluding, along with the pre-established 20% of participants, certain initial and final hand locations. This manipulation was intended to allow for exploration of the network's ability to generalise to completely novel conditions (i.e. participant and hand location). Specifically, the 3-D training dataset excluded the middle lower and middle upper targets, yielding one training dataset and three generalisation datasets. The training dataset contained all targets for 80% of the participants, excluding the middle lower and middle upper targets. The first generalisation dataset was used to determine the performance of the network for novel targets accessed by participants in the training dataset. It included the middle lower and middle upper targets for those participants present within the training dataset. The second generalisation dataset assessed conditions in which the participant was not in the training dataset but the target was, and contained all targets other than the middle lower and middle upper for those participants not included in the training dataset (i.e. 20% of participants). Finally, the third generalisation dataset examined network performance for conditions in which neither the participant nor the target were in the training dataset, including the middle lower and middle upper targets for those participants not included in the training dataset.

Repeated-measures ANOVA was used to detect significant differences between the ANN predictions and the empirical data. Root MSE (RMSE), determined as the angular differences at each joint between the model prediction and the corresponding empirical observation, were the dependent variables in these analyses. Independent variables were the type of dataset (training vs. generalisation), joint (10 joints for 2-D, 36 for 3-D), target locations (10 levels; near and far horizontal distances to five vertically distinct locations) and rotation (0°, 45° and 90°; only used for 3-D predictions). Post-hoc pairwise comparisons were done using Student–Newman–Keuls tests. A type I error rate of 5% was used for all statistical tests.

Given that a trained ANN represents a transformation function, plots were created that reflect the sensitivity of the 2-D and 3-D network outputs to changes in the network. Graphical depictions of these 'transformations' were complemented by plots comparing the variability observed from the 2-D and 3-D network outputs against the variability observed and quantified within the dataset. These 'variability envelope' plots were generated through systematic variation in the participant-dependent strength and anthropometry inputs.

As an additional evaluation of the 3-D ANN model performance, predictions were compared to those made by The University of Michigan 3DSSPP^{TM} model (Regents of The University of Michigan, 2005, version 5.0.0). Participant stature, body mass and gender were input to the model, along with the hand location, and the program output (i.e. joint locations) used to calculate joint angles, following the approach already described above for the HUMOSIM data. Data were used for all near and far targets, for all target locations and rotation conditions. Joint angle predictions for the ankle, wrist and neck could not be obtained from this software and were not used in the comparison. ANOVA was used to detect significant differences (5% type I error) in RMSE between the ANN and 3DSSPP^{TM} predictions, considering all the previously described independent variables except the type of dataset, which was not applicable.

3. Results

3.1. 2-D model

The error limit (0.003 rad^2) was not reached within the 500,000 epoch limit. After training, MSE was 0.05 rad^2, equivalent to a RMSE of 0.22 rad (12.8°). No significant differences in RMSE were found between the training dataset and the generalisation dataset. RMSE across the training and generalisation datasets differed significantly between joints ($F(9,190) = 39.27$, $p < 0.0001$) and target locations ($F(9,166) = 15.38$, $p < 0.0001$); the joint $\times$ location interaction was also significant ($F(81,1304) = 5.46$, $p < 0.0001$). Among the different joints, the clavicle exhibited the largest errors (mean RMSE: 33.7°), followed by the wrist (mean RMSE: 24.4°). The remaining angles were predicted with statistically similar errors (RMSE range: 6.6°–14.9°).

Across target locations, the near and far lower locations had significantly larger errors (mean RMSE: 15.8° and 14.1°) than other higher locations (RMSE range: 7.8°–11.0°). Other locations could not be distinguished statistically. Post hoc tests for the interaction effect showed differential magnitudes in the RMSE for different joints across some locations, but the overall patterns described for joint and target location effects were maintained.

3.2. 3-D model

The error limit (0.003 rad^2) was not reached within the 500,000 epoch limit. Following training, MSE was 0.52 rad^2, equivalent to a RMSE of 0.72 rad (41.5°). Significant differences in RMSE were observed as a function of the training/generalisation dataset used ($F(3,40) = 4.78$, $p = 0.0061$). However, a post hoc test showed that the training dataset did not differ significantly from any of the generalisation groups. Generalisation datasets with novel participants, however, had significantly higher RMSE than those with novel targets accessed by familiar participants. This last generalisation group had, in fact, the lowest overall levels of average error (mean RMSE: 10.3°), even lower than the errors for the training dataset (mean RMSE: 14.7°). Average RMSE for the remaining generalisation datasets was $\sim$20°.

As in the 2-D case, significant effects on RMSE were observed across joints and target locations and the interaction between these factors ($F(35,684) = 12.75$, $p < 0.0001$; $F(29,534 = 9.82$, $p < 0.0001$; $F(1015,18006) = 13.84$, $p < 0.0001$, respectively). Three different angles, the right knee α, left hip α and left shoulder α, exhibited significantly higher errors than angles for the remaining joints, which could not be differentiated statistically. The right knee α had the largest average error (mean RMSE: 58.1°), while the right ankle β had the lowest (mean RMSE: 4.6°). Overall, the observed errors were considerably larger than those obtained for the 2-D case.

With regard to target location, the far lower target location requiring no rotation exhibited the largest average error (mean RMSE: 50.1°), followed by the near lower target location requiring no rotation (mean RMSE: 30.9°). Remaining target locations could not be statistically distinguished. As for the 2-D ANN, post hoc tests for the interaction effect showed differential magnitudes in the RMSE for different joints across some locations, but the overall patterns highlighted by the joint and target locations were maintained. Altogether, these errors provided predictions that resulted in small to substantial deviations from actual postures (Figure 2).

3.3. Qualitative observation of network structure

Plots were created that represented the network transformation functions based on the input parameters. Recall that, in the 2-D model, seven inputs were provided to the network: two for the location of the end-effector, one for a participant-dependent strength coefficient and four that represented participant-dependent anthropometry coefficients. Before introducing these plots, note that the vertical axis on them is not consistently scaled

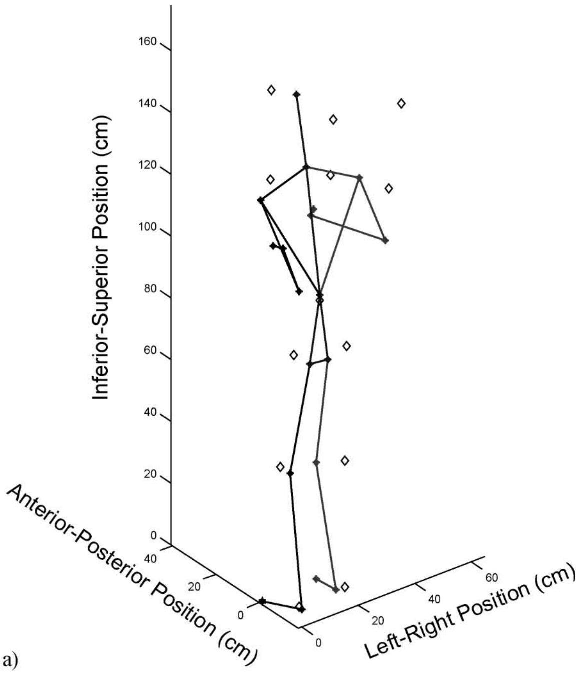

a)

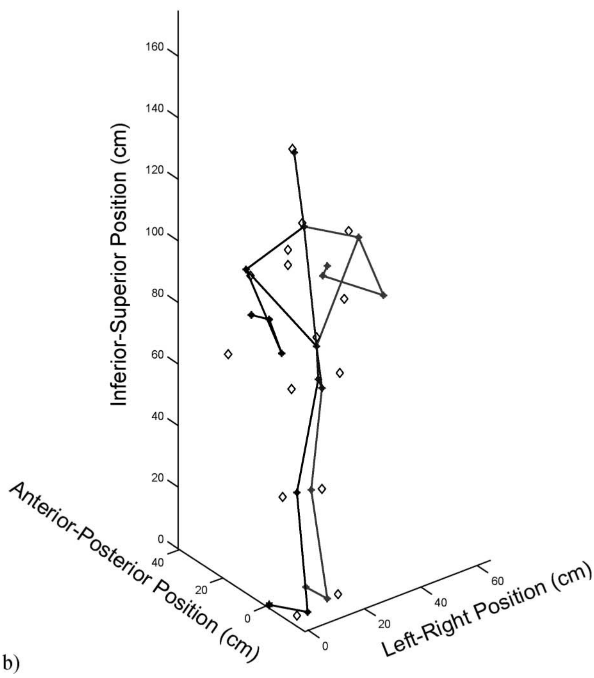

b)

Figure 2. Comparison of predicted posture (stick figure) and actual posture (joint positions are represented by hollow diamonds) for training cases where maximum (a) and minimum (b) root mean square error levels were observed.

across the different joint angles. This was done to purposefully amplify differences that, although present, would not be visible if all plots were presented on the same vertical scale. Thus, care is required in making relative comparisons between different joints. In addition, the plots are based on the complete dataset, rather than on training or generalisation subsets. This approach was selected to make the plots as representative of overall network performance as possible by including as much data as were available. Use of the entire dataset was also justified based on the similarity between prediction errors attained for both training and generalisation datasets.

For the first plot, only the box location (in 2-D) is manipulated; the remaining inputs are zeroed (Figure 3). Neck angle was not included in the plots since it had no effect on the kinematic chain. As is evident from the shape of the surface, the network provided a large importance to the inferior–superior (i.e. vertical) box position. Some joints show effects of the anterior–posterior (i.e. horizontal) box position (e.g. the lower torso), but these effects are smaller than those observed for the inferior–superior box position. This finding was not surprising, since the distance over which the anterior–posterior box position was varied in the dataset was relatively small (<0.6 m) and probably required only slight postural adjustments from participants.

Similar surfaces were obtained by modifying other model inputs. For example, the strength input varied from -4 to 4 (numbers are unitless). When the strength input is varied between these limits and the results are compared with those when the input was zeroed (Figure 4), there is a trend suggesting increased knee, hip and torso usage (i.e. higher angular deviations from neutral) with increased strength.

The 'variability envelope' plots for the 2-D network (Figure 5) indicate that the predicted variability tended to be smaller than the variability already present within the corresponding data. However, for many angles the shape of both the actual and predicted between-subject variability envelopes was quite similar. Prediction error can also be inferred, to some extent, in Figure 5; the more that bisecting lines for each of the shaded regions overlap, the smaller the error in the mean joint angle across participants, which is indicative of the overall degree of prediction error. Both results need to be considered together in determining overall predictive performance.

Similar plots for the 3-D ANN provided somewhat contrasting findings, since the end-effector position in all three axes influenced substantially the different predicted angles. Substantial strength and anthropometry effects could also be seen for many of the angles. The plots on which these observations are based are not included here, given that they are structurally similar to those already shown, albeit with a larger number of outputs. Substantial overlap in the empirical and predicted variability for a few 3-D angles was also observed and the predicted variability tended to be included within the envelope provided by the between-subjects variability.

3.4. Comparison to 3DSSPP^{TM}

Predictions obtained from 3DSSPP^{TM} had significantly larger RMSE than those made by the 3-D ANN ($F(1,38) = 70.01$, $p < 0.0001$). Across all target locations and joints, the RMSE obtained using the ANN was 25.2°, compared to 37.7° using 3DSSPP^{TM}. Significant differences were also found for the joint × model (i.e. ANN or 3DSSPP^{TM}) and target location × model interactions ($F(25,456) = 83.34$, $p < 0.0001$; $F(29,496) = 96.65$, $p < 0.0001$, respectively). While unbalanced data precluded the completion of post hoc tests for these interactions, qualitative observations indicate that certain 3DSSPP^{TM} joint predictions, especially for joints in the lower extremities, exhibited

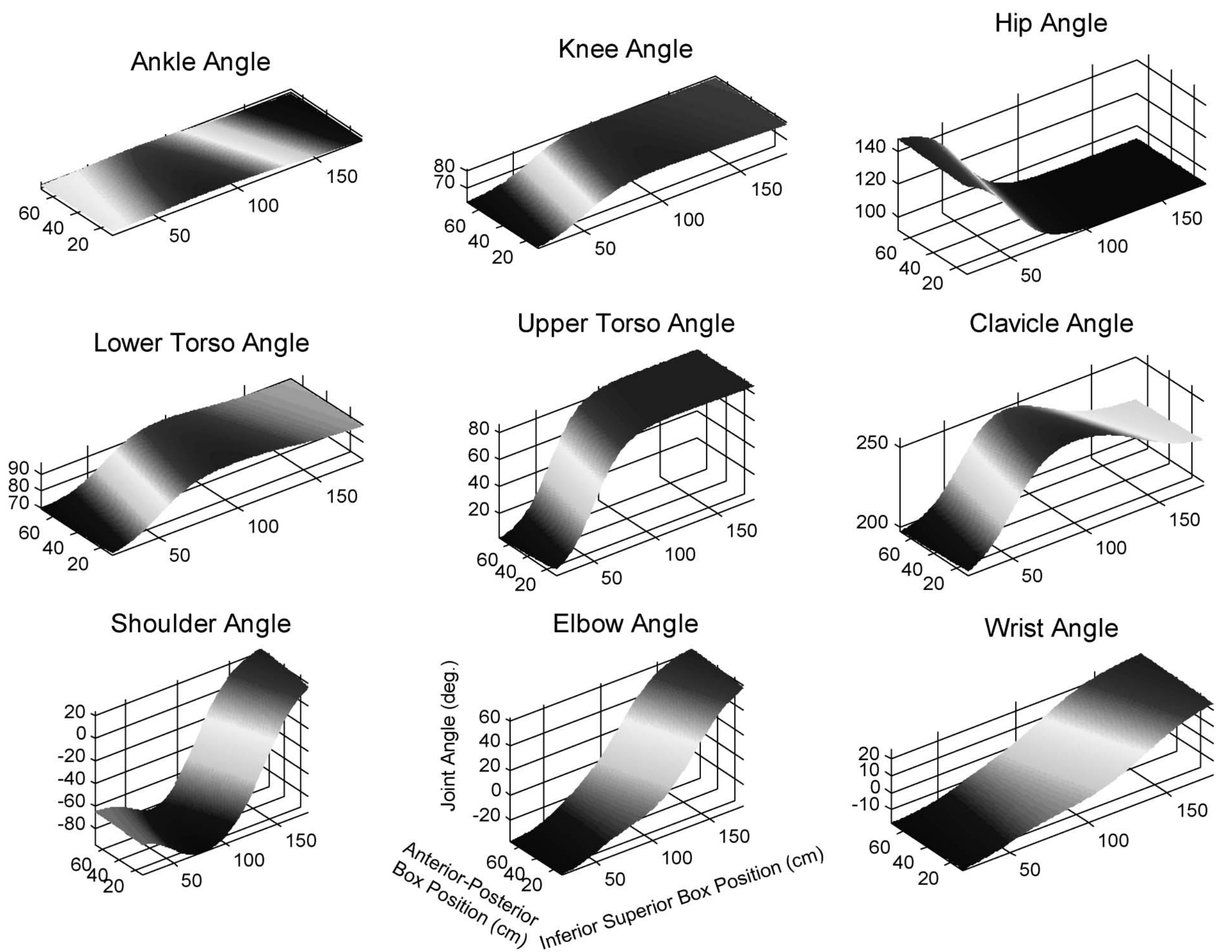

Figure 3. Predicted joint angles as a function of box position.

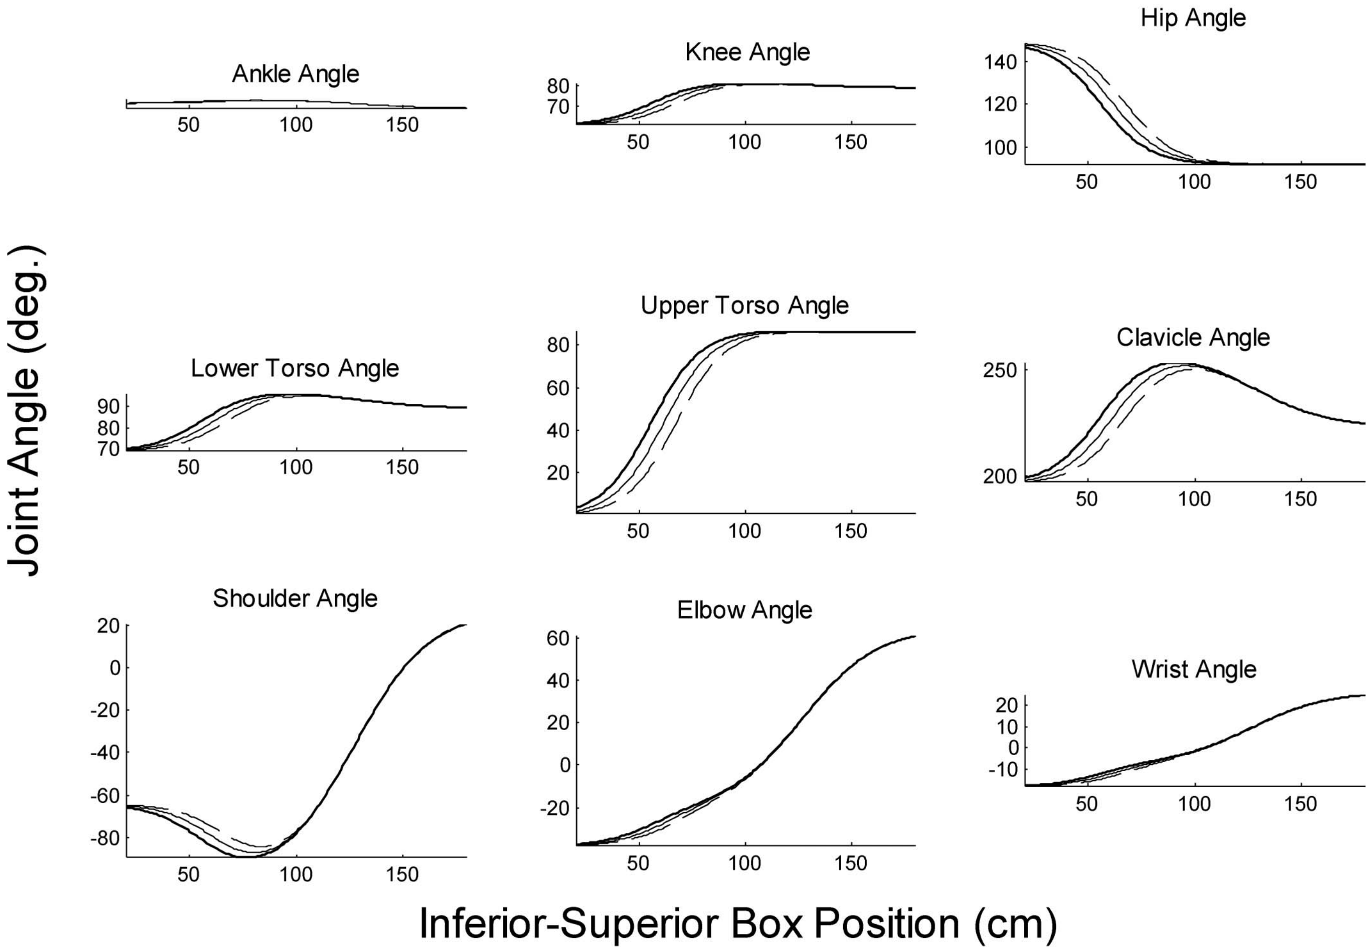

Figure 4. Predicted joint angles as a function of inferior–superior box position. The dashed line represents a low strength input (–4), the solid thin line a neutral strength input (0) and the solid wide line a high strength input (+4).

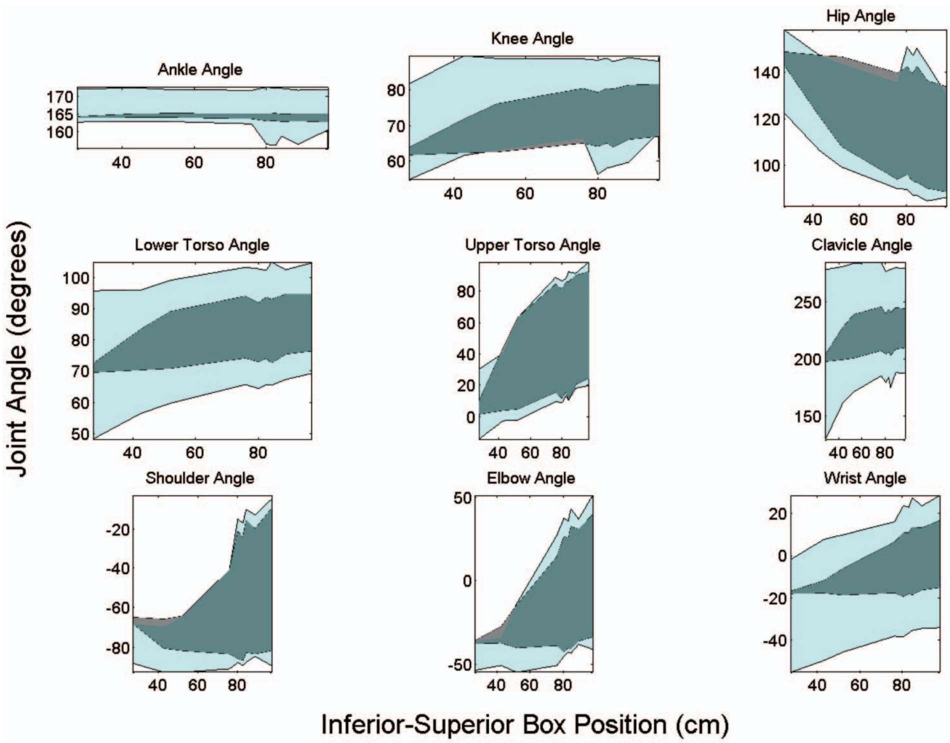

Figure 5. Comparison between observed (light grey) and predicted (dark grey) between-subject variability. The shaded regions correspond to a ± 1 SD region around the between-subjects mean for the observed and predicted angles.

disproportionately larger errors. The 3DSSPP^{TM} environment was unable to predict a posture for 1.3% of the input box positions (29 trials). It was not apparent whether these predictions were not calculated due to violation of balance constraints or due to anthropometry assumptions that made the location unreachable by the humanoid.

4. Discussion

In this investigation, ANNs were provided with input information to predict the initial and final postures for several lifting motions. The only motion-specific inputs provided to the ANN were the initial or final hand positions, which are often relatively easy to determine for existing or simulated lifting motions and are usually design constraints. The final ANN output, in the form of joint angles, could be easily input into a variety of software and/or analytical techniques that can then generate postural risk estimates or simulate the complete lifting motion.

The usefulness of 2-D and 3-D ANN-based predictions is also supported by comparisons with existing model predictions. Beck and Chaffin (1992) used inverse kinematics to predict postures during sagittally symmetric 2-D and 3-D movements. Prediction errors for different joints were typically between 5° and 15°. These errors are comparable, if slightly lower, to those obtained here. Note, however, that Beck and Chaffin report mean absolute errors, which in general tend to be slightly smaller than the

RMSE reported here. They also used a slightly smaller sample (five to 10 subjects, typically) and a more strenuous lifting task (75% of maximum load). In any case, the current model can be considered similar in performance to this inverse-kinematics-based approach.

The Beck and Chaffin (1992) model has been greatly modified and improved in its current use within 3DSSPP™, but comparisons of this program's predictions against the ANN predictions showed that the ANN predictions were significantly more accurate. Reasons for the higher error in the 3DSSPP™ predictions were not apparent. However, note that the 3DSSPP™ program does not advertise its posture prediction function as an accurate prediction, but rather as a feasible posture that should be modified by the user to represent the posture in which they are interested. Thus, the purpose of the function is likely not to explicitly maximise accuracy. The 3DSSPP™ had access to the same information that was available to the 3-D ANN-based model (only the 3-D model was used in this comparison). This information included the hand location and anthropometric information for each participant. In some cases, further performance-enhancing information could have been provided. For example, the 3DSSPP™ model can incorporate information as to whether the lower extremities were used in attaining a particular posture (e.g. for a squat lift). This information was not provided to the 3DSSPP™ model, however, to allow for a direct comparison with the current model.

Comparisons with other published models yielded similar favourable results. For example, the ANN model also outperformed a cost-based optimisation model proposed by Jung and Park (1997), who reported average mean absolute errors (not RMSE) on the order of 20° to 40°. An optimisation process was used in their approach, with relatively few constraints, which might explain the low performance levels, but which also makes it more comparable to the current model, which included no constraints.

Dysart and colleagues (Dysart 1994, Dysart and Woldstad 1994, 1996) reported errors in terms of a Euclidean distance parameter, which considers the squared prediction errors for the hip location and the forearm orientation. Qualitatively, however, their models (which employed optimisation with a number of different objective functions) yielded errors larger than those found here for two dimensions and in most lifting situations. Further developments of these models, reported in Woldstad (1997), still indicate relatively large errors employing a similar optimisation approach, on the order of 15° or more, although one of the objective functions showed mean absolute errors across joints between 5° and 20°.

Results for the 2-D and 3-D networks were encouraging, in that postures in both the training and generalisation datasets were predicted with comparable levels of error. It thus appears that the networks developed an ability to generalise, with no noticeable increase in error, to novel targets (3-D model) and participants (both 2-D and 3-D models). In addition, the results support the use of these networks as posture prediction tools rather than posture reproduction tools, the former reflecting that certain postural 'rules' were encoded within the networks. While these rules do not necessarily reflect the motor control mechanisms by which humans act, they may aid in understanding the relative importance of the different inputs provided to the networks. For example, the network may be using between-subjects (i.e. person-specific) anthropometric characteristics to define a 'likely' posture that is further modified based on additional inputs.

Sample visual depictions of these 'rules' are presented in Figures 3 and 4. These plots, along with similar plots generated for the remaining inputs, suggest that the 2-D network utilised the inferior–superior box position as its primary input, with the strength and first anthropometry inputs generating joint angle adjustments based on the particular individual being modelled. Anterior–posterior box position affected some angles, but to a lesser extent. The remaining anthropometry inputs seemed used for modulating inter-individual differences in particular segments (e.g. upper limbs) and their influences were smaller than those for the first anthropometry input. In turn, these plots can be used to infer the lifting strategies that the ANN predicts. For example, based on Figures 3 and 4, usage of the knee, hip and torso occurs mainly for inferior–superior locations lower than 100 cm. The knee and lower torso joints were not fully utilised (only a 20° utilisation range), in contrast with the hip and upper torso angles ($> 50^\circ$ utilisation range). Usage of the upper extremities was evident throughout the inferior–superior range of motion. In combination, these observations suggest the network tended to predict 'stoop' postures for lower object locations. Interestingly, joint utilisations suggesting the use of a 'squat' posture were predicted in combination with high individual strength, suggesting that high strength may make the use of 'squat' postures more likely. These strength and anthropometry effects could be formally tested in future work by obtaining similar posture data on individuals selected based on specific anthropometry and strength attributes, perhaps following approaches similar to Gill *et al.* (2007), who systematically varied several lift parameters.

The 'variability envelope' plots (Figure 5) suggest that the 2-D network performed marginally well in reproducing the variability within the available data. Variability in hip, upper torso, shoulder and elbow angles was modelled quite well and prediction errors for these joints were relatively low. However, the remaining angles, albeit not necessarily being subject to larger prediction error (in fact, these were smaller than those for the hip, upper torso, shoulder and elbow in many cases), were not modelled within the full bounds provided by the observed variability. In general, the 2-D ANN provided relatively good predictions of mean joint angles used by different individuals to perform lifting tasks similar to those described in this work and represented some of the variability demonstrated by those individuals.

Similar plots for the 3-D network suggested that more uniform use of inputs was made in comparison to the 2-D network, since most inputs resulted in observable differences in the 'transformation function' plots. The 'variability envelope' plots for the 3-D network produced similar patterns to those observed for the 2-D network. In general, and similar to the 2-D case, the 3-D results show that the network predictions modelled the data relatively well.

The 'variability envelope' plots demonstrate some important characteristics of posture data that have implications for modelling. One is the large inherent inter-individual variability that is present in these data, which has been more formally quantified in previous efforts (Perez and Nussbaum 2006). This large inter-individual variability may frustrate efforts to predict actual postures for specific individuals, rather than for a collection of individuals, since the mechanisms that cause this large variability remain unknown. Given the typically excellent data-assimilation characteristics of ANNs, and the positive results from comparisons against other modelling approaches (see below), the predictions reported here can be interpreted as approximations of the lower bounds of prediction errors. While this would be somewhat discouraging, predictions that approximate the empirical 'variability envelope', as many of the ANN-based model predictions did, could be very useful. These predictions could be used to describe populations of workers, which is often more important in industry than describing the behaviour of an individual worker. Such population predictions could be used as inputs to other models and processes (e.g. Monte-Carlo simulation) that in turn provide estimates of injury risks based on a variety of factors.

In general, the ANN-based models described here predicted 2-D and 3-D postures with varying degrees of error and some of these errors were considerable. However, prediction errors were at least comparable to, and in some instances lower than, those obtained using previously published approaches. Future work could be directed towards examining whether the observed results extend to circumstances that were not addressed in the dataset and/or models (e.g. lifts requiring substantial anterior–posterior translation). Potential improvements of the models could also be considered in future work, arising from changes to the network structure, introduction of additional inputs that help account for more of the variability within the dataset and the addition of physiological constraints to model outputs, among others. Overall, the results suggest that ANN-based models may be useful posture prediction tools. These models not only assimilated some key characteristics of the available data, but also used this 'knowledge' to predict postures for novel circumstances and individuals. Future models, however, should consider and attempt to account for the large inter-individual variance that characterises posture data.

References

Beck, D.J. and Chaffin, D.B., 1992. An evaluation of inverse kinematics models for posture prediction. *In*: *Proceedings of the international conference on computer-aided ergonomics and safety '92*, Tampere, Finland. Amsterdam: North Holland Publishing.

Bellan, Y., Gaia, E., and Rizzuto, F., 2000. ANNIE: standard bio-mechanics and cognitive models for computer aided ergonomics. *In*: *Proceedings of the 44th annual meeting of the Human Factors and Ergonomics Society – 'ergonomics for the new millennium'*. San Diego, CA: Human Factors and Ergonomics Society.

Chaffin, D.B., 2005. Improving digital human modelling for proactive ergonomics in design. *Ergonomics*, 48, 478–491.

Crawford, N.R., Yamaguchi, G.T., and Dickman, C.A., 1996. Methods for determining spinal flexion/extension, lateral bending, and axial rotation from marker coordinate data: Analysis and refinement. *Human Movement Science*, 15, 55–78.

Dysart, M.J., 1994. Development and validation of a posture prediction algorithm. *Industrial and Systems Engineering*. Blacksburg, VA: Virginia Polytechnic Institute and State University.

Dysart, M.J. and Woldstad, J.C., 1994. Development and validation of a posture prediction algorithm for a static lifting task. *In*: *Proceedings of the triennial congress of the International Ergonomics Association – IEA '94*. International Ergonomics Association.

Dysart, M.J. and Woldstad, J.C., 1996. Posture prediction for static sagittal-plane lifting. *Journal of Biomechanics*, 29, 1393–1397.

Faraway, J.J., 2000. *Modeling reach motions using functional regression analysis*. Detroit, MI: Society of Automotive Engineers.

Gill, K.P., *et al*., 2007. Regional changes in spine posture at lift onset with changes in lift distance and lift style. *Spine*, 32, 1599–1604.

Haykin, S., 1999. *Neural networks: A comprehensive foundation*. Upper Saddle River, NJ: Prentice Hall.

Hsiang, S.H., Brogmus, G.E., and Courtney, T.K., 1997. Low back pain (LBP) and lifting technique – a review. *International Journal of Industrial Ergonomics*, 19, 59–74.

Johnson, D.E., 1998. *Applied multivariate methods for data analysts*. Pacific Grove, CA: Duxbury Press.

Jung, E.S. and Park, S., 1994. Prediction of human reach posture using a neural network for ergonomic man models. *Computers and Industrial Engineering*, 27, 369–372.

Jung, E.S. and Park, W., 1997. Comparison of biomechanical cost functions with weighting factors for posture prediction of human load lifting. *In*: *Proceedings of the triennial congress – International Ergonomics Association*. International Ergonomics Association.

Keyserling, W.M., 2000. Workplace risk factors and occupational musculoskeletal disorders, Part 1: A review of biomechanical and psychophysical research on risk factors associated with low-back pain. *American Industrial Hygiene Association Journal*, 61, 39–50.

Lim, S.Y.E., Fok, S.C., and Tan, I.T.Y., 1996. Neural network investigation of posture and motion. *In*: S.A. Robertson, ed. *Proceedings of the annual conference of the Ergonomics Society*. London: Taylor and Francis.

Mirka, G.A., *et al*., 2000. Continuous assessment of back stress (CABS): A new method to quantify low-back stress in jobs with variable biomechanical demands. *Human Factors*, 42, 209–225.
National Institute for Occupational Safety and Health. ed. 1997. *Musculoskeletal disorders and workplace factors*. Cincinnati, OH: National Institute for Occupational Safety and Health.
Park, W., *et al*., 1999. Integration of magnetic and optical motion tracking devices for capturing human motion data. *In*: *Proceedings of SAE digital human modeling for design and engineering conference*. The Hague, Netherlands: Society of Automotive Engineers.
Perez, M.A. and Nussbaum, M.A., 2006. Posture and motion variability in non-repetitive manual materials handling tasks. *Human Movement Science*, 25, 409–421.
Russell, S.J., *et al*., 2007. Comparing the results of five lifting analysis tools. *Applied Ergonomics*, 38, 91–97.
Werbos, P.J., 1988. Generalization of backpropagation with application to a recurrent gas market model. *Neural Networks*, 1, 339–356.
Windhorst, U., 1996. On the role of recurrent inhibitory feedback in motor control. *Progress in Neurobiology*, 49, 517–587.
Woldstad, J.C., 1997. Further evaluations of a revised posture prediction algorithm for static lifting. *In*: B. Das and W. Karwowski, eds. *Advances in occupational ergonomics and safety II*. International Ergonomics Association. Amsterdam: IOS Press.
Zhang, X. and Chaffin, D.B., 2000. A three-dimensional dynamic posture prediction model for simulating in-vehicle seated reaching movements: development and validation. *Ergonomics*, 43, 1314–1330.
Zhao, J. and Badler, N.I., 1994. Inverse kinematics positioning using nonlinear programming for highly articulated figures. *ACM Transactions on Graphics*, 13, 315–336.