Eliminating Motion Artifacts from Fabric-mounted Wearable Sensors

Brendan Michael and Matthew Howard

***Abstract*— Sensors embedded into clothing for measuring human movement are becoming more widespread in research, with applications in clinical diagnostics or rehabilitation studies. A major issue with their use is the undesired effect of fabric motion artifacts corrupting movement signals. This paper presents a method for learning body movements, viewing the undesired motion as stochastic perturbations to the sensed motion, and utilising errors-in-variables models. Experiments, both in simulation and with a physical fabric-mounted sensor, indicate improved prediction accuracy as compared to standard learning methods.**

I. Introduction

Measuring and analysing human movement is important in a wide range fields such as analysing human gait for rehabilitation [1], medical diagnostics [2], or even animation control [3]. In the laboratory or studio environment, a variety of devices are available for taking such measurements, including sophisticated inertial motion capture systems such as the XSens Motion Capture System [4] or other strap-on sensors [5].

However, outside the laboratory environment, these systems can present a number of practical problems. These include (i) size and weight of sensors effecting user movements (undesirable for analysis/diagnostic studies), (ii) sensors being visually obtrusive (detrimental to patient experience [6]), and (iii) inability to record over a continuous time-period lasting days or even weeks.

A natural solution to this problem is to incorporate such systems into items already in use by patients. The development of *fabric-embedded sensors*, also known as *e-textiles*, is an emerging technology that uses small sensors mounted onto items of fabric, such as clothing [7]. These sensors are very small in size (order of millimetres in diameter) and are built in mind of the user's comfort by minimising physical and visual invasiveness. In comparison to laboratory devices, they are also very inexpensive.

While these new, soft sensing technologies offer significant potential for inexpensive and unobtrusive capture of human movement data, there remain a number of problems in their use. An important issue, is that of how to deal with motion artifacts corrupting data recordings, as caused by the unpredictable motion of fabric sensors with respect to the body (see Fig. 1). Previous fabric systems have attempted to solve this problem by fitting sensors tightly to the body (e.g., by use of straps or other tight-fitting garments [7]). However, this is unsatisfactory if sensors are to be incorporated in an unobtrusive way, into everyday items of clothing.

As an alternative, this paper proposes the use of statistical methods from the errors-in-variables field to learn models of the wearer's movement that eliminate the effects of fabric motion artifacts. The proposed approach is computationally efficient, and can be easily implemented in an embedded system for on-board (i.e., on-wearer) prediction of movements. Experimental results in learning and predicting movement from acceleration data from a physical e-textile device indicate superior performance as compared to standard learning approaches.

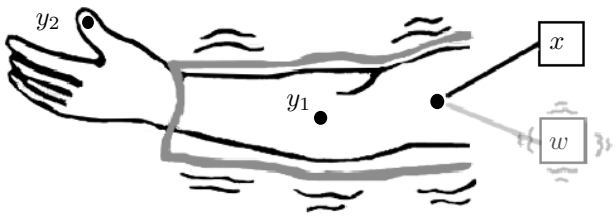

Fig. 1: Prediction of wearer movement at different points on the body (e.g., forearm y_1, fingertip y_2) based on sensor readings. Fabric motion with respect to the body causes unpredictable artifacts to be introduced into e-textile sensor readings w, as compared to those derived from a rigidly attached sensor x.

II. Problem Definition

To motivate the proposed approach, consider the effect that non-rigid attachment of body-mounted sensors may have on estimation of a wearer's movement (see Fig. 1). In a typical problem, one might have access to readings w from a sensor (e.g., accelerometer) mounted on an item of clothing (e.g., shirt sleeve). From these, it is desired to estimate the corresponding motion (accelerations) of different points on the wearer's body. For example, these might include estimating the acceleration of the forearm y_1, or a fingertip y_2.

Given readings from a sensor located rigidly on the arm x, this can be achieved in a straight forward way, either through forward kinematics (given an analytic model of the wearer's arm), or by learning the mapping from x to y_1 or y_2 from a set of calibration data.

However, in the case of sensing from e-textiles, the loose coupling between the wearer motion and sensor readings, means that artifacts are introduced from a number of sources, for example, external air currents, air resistance on the fabric, vibrations and the dynamics of the fabric itself. These unpredictable disturbances may cause a significant discrepancy between the sensor readings w and the underlying motion of the wearer.

A. Explicit Models of Fabric Dynamics

One approach to deal with these issues would be to explicitly model the wearer/fabric interaction dynamics. For

example, in the field of cloth animation, fabric dynamics have been simulated using mass-spring models or particle systems based on motion capture data [8]. This, however, is a complex procedure, due to the need to first analyse the mechanics of interactions, then build a suitable model to represent these interactions [9]. The computational demands of such approaches make them unsuitable for a light-weight, embedded wearable system. Furthermore, different materials may interact in different ways, depending on their internal fabric, or fibre structure [10]. This means that any given model formed in this way would be restricted to a particular class of fabric, and would become more complex with the addition of differing materials.

B. Learning Fabric Motions

In the absence of a detailed model of the wearer/fabric interaction dynamics, an appealing approach is to use statistical learning techniques to form a predictive model of the wearer's motion. Addressing the problem in this way provides many benefits over analytical fabric modelling, since it allows unpredictable motion artifacts to be treated as stochastic perturbations to the underlying motion. It removes the need to estimate physical quantities such as mass and fibre structure of the fabric, and through use of simple parametric models, can be computationally very inexpensive.

To apply such an approach, a calibration stage is required in which data from the the target quantity $y \in \mathbb{R}$ and the fabric sensor readings $\mathbf{w} \in \mathbb{R}^{\mathcal{P}}$ are gathered for training the model. In the setting considered here (ref. Fig. 1), such data may be gathered by subjecting the system to various movements while data is recorded both from the fabric sensors and from a sensor measuring the target quantity.

Note that, since the latter is only needed temporarily (i.e., during the calibration), a larger, rigidly-attached sensor can be used, that may otherwise not be suitable for long-term use. For example, one might chose to use a more intrusive, but higher fidelity motion capture sensor to obtain high quality readings, knowing that once the calibration is complete, the rigid sensor may be discarded, in favour of the predictions obtained from the fabric sensor readings.

C. Standard Least Squares Estimators

While the approach described in Sec. II-B is appealing for dealing with fabric-mounted sensor data, close examination of the usual assumptions made in standard learning approaches suggest its direct application may be problematic. This is due to an important mismatch in the sources of error assumed by these approaches, and those actually encountered in the data.

Specifically, the standard assumption made by such techniques is that data are generated according to a model of the form

$$y = f(\mathbf{x}) + \epsilon \tag{1}$$

where ϵ denotes additive noise on y and f denotes the functional relationship between the sensed inputs $\mathbf{x}$ and the target outputs y. Given samples $\{\mathbf{x}_n, y_n\}_{n=1}^{\mathcal{N}}$ the goal of the learning approach is to form an estimate of the function f.

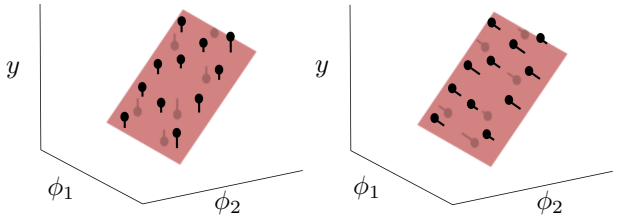

Fig. 2: Fitting the model (1) with standard least squares (left) minimises the residuals due to error in the target variable y but ignores errors in the inputs. By minimising the residuals orthogonal to the fit, total least squares (right) reduces their effect.

A common approach for this is to minimise the sum of the squared residuals

$$S_o = \sum_{n=1}^{\mathcal{N}} (y_n - \tilde{y}_n)^2 \tag{2}$$

where $\tilde{y}_n$ denotes the prediction of the model on the nth data point. A convenient class of function approximators are the generalised linear models

$$\tilde{y} = \boldsymbol{\phi}(\mathbf{x})^\top \tilde{\boldsymbol{\theta}} \tag{3}$$

where $\boldsymbol{\phi}(\cdot) \in \mathbb{R}^{\mathcal{J}}$ is a suitable feature vector or set of basis functions, such as Gaussian radial basis functions or polynomials and $\tilde{\boldsymbol{\theta}} \in \mathbb{R}^{\mathcal{J}}$ is the parameter. Note that, for convenience, in this paper it is assumed that $\boldsymbol{\phi}(\cdot)$ contains the term $\phi_{\mathcal{J}}(\mathbf{x}) := 1$ to encode any constant offset in the target function (1).

The optimal choice for the parameter (with respect to (2)) is

$$\tilde{\boldsymbol{\theta}} = (\boldsymbol{\Phi}^\top \boldsymbol{\Phi})^{-1} \boldsymbol{\Phi}^\top \mathbf{y} \tag{4}$$

where $\boldsymbol{\Phi} := (\boldsymbol{\phi}_1^\top, \cdots, \boldsymbol{\phi}_{\mathcal{N}}^\top)^\top \in \mathbb{R}^{\mathcal{N} \times \mathcal{J}}$ is the data matrix, containing independent sample features $\boldsymbol{\phi}_n := \boldsymbol{\phi}(\mathbf{x}_n)$ on each column.

In the context of fabric-based sensing, however, difficulties occur due (1) being a poor model of the noise encountered in the data. As described in Sec. II-B, the major source of noise in fabric sensor data is that of the motion of fabric with respect to the wearer, i.e., *noise on the inputs* $\mathbf{x}$. This has a number of implications with respect to the reliability of movement prediction models computed according to this standard approach. For example, ignoring these perturbations, and using the ordinary least squares estimate (4) may result in (i) bias in the parameter estimation (ii) loss of power in detecting relationships, and (iii) the masking of features in non-linear models [11].

III. Method

To deal with these problems, in this paper it is proposed to explicitly account for non-negligible noise in the *independent variables* $\mathbf{x}$ during learning and prediction. Specifically, the latter are assumed to be subject to zero-mean, additive noise $\boldsymbol{\varepsilon}$, that corrupt the sensor readings

$$\mathbf{w} = \mathbf{x} + \boldsymbol{\varepsilon} \tag{5}$$

Given samples $\{\mathbf{w}_n, y_n\}_{n=1}^{\mathcal{N}}$, the task is to form a prediction model (3), that takes account of stochasticity both in $\mathbf{x}$ and y in the data.

A. Model Estimation through Total Least Squares

An intuitive approach to achieve this, is to modify the objective function (2) such that the parameter estimate minimises the squared residuals orthogonal to the predicted curve, an approach known as Total Least Squares (TLS) fitting [12], [13]. The following describes how TLS can be applied to fit a generalised linear model (3) given data $\{\mathbf{w}_n, y_n\}_{n=1}^{\mathcal{N}}$.

In particular, augmenting the feature vector[1] $\boldsymbol{\phi}(\cdot)$ with the targets y, (3) can be re-written

$$\boldsymbol{\zeta}^\top \tilde{\boldsymbol{\omega}} = 0 \tag{6}$$

where $\boldsymbol{\zeta}(\mathbf{x}, y) := (\phi_1(\mathbf{x}), \ldots, \phi_{\mathcal{J}-1}(\mathbf{x}), y)^\top \in \mathbb{R}^{\mathcal{J}}$ and $\tilde{\boldsymbol{\omega}} \in \mathbb{R}^{\mathcal{J}}$ is the vector of parameters.

In this augmented space, instead of minimising the residuals in y (as in (2)), the proposed approach minimises the sum of squared *orthogonal residuals*

$$S = \sum_{n=1}^{\mathcal{N}} d_n^2 \tag{7}$$

where d_n is the orthogonal distance from the nth data point to the plane defined by (6)

$$d_n = \boldsymbol{\zeta}_n^\top \hat{\boldsymbol{\omega}} \tag{8}$$

with $\boldsymbol{\zeta}_n := \boldsymbol{\zeta}(\mathbf{w}_n, y_n)$ and $\hat{\boldsymbol{\omega}} := \tilde{\boldsymbol{\omega}}/||\tilde{\boldsymbol{\omega}}||$ (see Fig. 2). It can be shown [12], [13] that the plane minimising (7) must pass through the centroid of the data

$$\bar{\boldsymbol{\zeta}} := \frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} \boldsymbol{\zeta}_n. \tag{9}$$

Hence, minimisation of (7) can be achieved by minimising its upper bound

$$\overline{S} = \sum_{n=1}^{\mathcal{N}} ((\boldsymbol{\zeta}_n - \bar{\boldsymbol{\zeta}})^\top \hat{\boldsymbol{\omega}})^2 \tag{10}$$

or in matrix notation

$$\overline{S} = ||\mathbf{Z}^\top \hat{\boldsymbol{\omega}}||^2 \tag{11}$$

where $\mathbf{Z} := ((\boldsymbol{\zeta}_1 - \bar{\boldsymbol{\zeta}})^\top, \cdots, (\boldsymbol{\zeta}_{\mathcal{N}} - \bar{\boldsymbol{\zeta}})^\top)^\top$.

The optimal solution is retrieved by forming the singular value decomposition of $\mathbf{Z}$, and setting $\hat{\boldsymbol{\omega}} = \underline{\mathbf{u}}$, the eigenvector corresponding to the smallest singular value. The parameter estimate is then retrieved as

$$\tilde{\boldsymbol{\theta}} = (\hat{\boldsymbol{\omega}}^\top, \hat{\omega}_0)^\top \tag{12}$$

where the last element (intercept term) is

$$\hat{\omega}_0 = -\bar{\boldsymbol{\zeta}}^\top \hat{\boldsymbol{\omega}}. \tag{13}$$

[1]Note that, to avoid biasing effects due to the mapping of $\boldsymbol{\varepsilon}$ into the feature space, the feature vector as far as possible should be chosen to satisfy $\mathbb{E}_{\boldsymbol{\varepsilon}}[\phi_j(\mathbf{w})] = \phi_j(\mathbf{x})$ for $j \in 1, \cdots, \mathcal{J}$, where $\mathbb{E}_{\boldsymbol{\varepsilon}}[\cdot]$ denotes expectation over $\boldsymbol{\varepsilon}$. In practice, this condition is not found to be crucial in obtaining a superior fit over approaches that ignore the input errors $\boldsymbol{\varepsilon}$.

B. Motion Prediction from Noisy Sensor Readings

Having learnt the model parameters $\tilde{\boldsymbol{\theta}}$, the next step is to form predictions based on incoming sensor readings.

In standard function approximation (see Sec. II-C), the movement estimate $\tilde{y}^*$ for a given query point $\mathbf{w}^*$ is simply

$$\tilde{y}^* = \boldsymbol{\phi}(\mathbf{w}^*)^\top \tilde{\boldsymbol{\theta}}.$$

However, because this fails to take account of the noise in $\mathbf{w}^*$, it can result in poor accuracy.

The ideal prediction would be obtained by directly feeding $\mathbf{x}^*$ to the model (1), but in the present context this reading is not directly accessible. It is therefore necessary to build an estimate $\tilde{\mathbf{x}}^*$ based on the data available.

In this paper, this is achieved through use of *replicate data* [11], whereby for any query point $\mathbf{x}^*$, the availability of $\mathcal{K}$ noisy replicates

$$\mathbf{w}_k = \mathbf{x}^* + \boldsymbol{\varepsilon}_k \tag{14}$$

is assumed within the training data. Under the assumption of zero-mean distribution of errors $\boldsymbol{\varepsilon}_k$ (ref. Sec. III), this means that a simple estimate $\mathbf{x}^*$ can be obtained by taking the sample mean of the replicates over $\mathcal{K}$. This allows the final prediction to be made as

$$\tilde{y}^* = \boldsymbol{\phi}(\tilde{\mathbf{x}}^*)^\top \tilde{\boldsymbol{\theta}}. \tag{15}$$

Note that, in general, the accuracy of the prediction $\tilde{\mathbf{x}}^*$ depends on the number of replicate data available at that point. For sensors measuring continuous variables (as considered here), exact replicates $\mathbf{w}_k$ of sensor readings at a given $\mathbf{x}^*$ are unlikely to be available. However, in practice, a good estimate can still be found from approximate replicates (i.e., using samples $\mathbf{w}_n = \mathbf{x}_n + \boldsymbol{\varepsilon}_n$ where $||\mathbf{x}_n - \mathbf{x}^*||$ is small). In the experiments reported here, approximate replicates are obtained through a simple, heuristic binning procedure whereby the $\mathbf{w}_n$ are grouped according to their value into a set of fixed-width bins, regularly spaced across the range of the input space.

IV. Evaluation

In this section, the proposed approach is evaluated through a simulation study, and through an experiment on acceleration data from a fabric-embedded device.

A. Simulation

The goal of the first evaluation is to characterise the performance of the proposed approach for learning and predicting movements from noisy sensory inputs. For this, learning is tested on artificial data from models with both linear and non-linear relationships between the input x, the sensed w, and the target quantity y. For example, x may represent the acceleration of a body segment (e.g., forearm), w the fabric sensor readings (e.g., from a shirt sleeve), and y the corresponding acceleration of another segment (e.g., hand), see Fig. 1. The experimental procedure is as follows.

As training data, a set of $\mathcal{M} = 50$ independent sample inputs are drawn from the uniform random distribution $x_m \sim U[-1, 1]$. For each of these, to simulate noise and motion

		Linear	Quadratic	Sinusoidal
LS	$\|\|\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}\|\|$	0.467 ± 0.053	3.188 ± 0.164	-
	NMSE	0.344 ± 0.043	0.748 ± 0.009	0.493 ± 0.041
TLS	$\|\|\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}\|\|$	0.189 ± 0.068	1.234 ± 0.368	-
	NMSE	0.067 ± 0.013	0.404 ± 0.173	0.265 ± 0.052

TABLE I: Mean norm difference between estimated and ground truth parameters and normalised mean squared error (NMSE) in predictions $\tilde{y}^*$. Results are mean±s.d. over 20 trials.

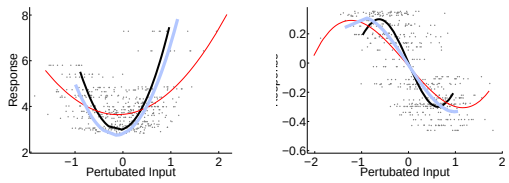

Fig. 3: Least squares (thin red) and total least squares (thick light purple) predictions overlaid on the ground truth (black) for f_2 (left), and f_3 (right) when learning on noisy data (grey dots).

artifacts, a set of $\mathcal{K} = 10$ readings are generated, corrupted with additive Gaussian noise

$$w_n = w_{m,k} = x_m + \varepsilon_k \tag{16}$$

where $\varepsilon_k \sim N(0, \sigma_\varepsilon^2)$ and $\sigma_\varepsilon^2 = 0.15$. At the same time, the corresponding target quantities y_n are computed for each of the readings w_n

$$y_n = y_{m,k} = f(x_m) + \epsilon_n \tag{17}$$

where $\epsilon_n \sim N(0, \sigma_\epsilon^2)$ and $\sigma_\epsilon^2 = 0.01$. In the following, results are reported for generative functions f that are (i) linear $f_1(x) = 1.5x + 3$, (ii) quadratic $f_2(x) = 4x^2 + 0.75x + 3$, and (iii) sinusoidal $f_3(x) = -0.3\sin(2.5x)$ in the inputs.

The resultant $\{w_n, y_n\}_{n=1}^{\mathcal{N}}$ are used to train the approximator (3) through the total least squares (TLS) method outlined in Sec. III. For this, $\boldsymbol{\phi}$ is chosen according to the model (f_1, f_2 and f_3) used to generate the data. In particular, for f_1 and f_2 basis functions exactly capturing the parametric form of the model (e.g., for f_2, $\boldsymbol{\phi}(x) := (x^2, x, 1)^\top$) are used, while for f_3, a 3rd order polynomial basis is used. For comparison, identical models are trained on the same data through standard least squares (LS), using the approach outlined in Sec. II-C. The procedure is repeated for 20 trials on different data sets.

Table I summarises the results. There, it can be seen that the parameters estimated by TLS for f_1 and f_2 are are much closer to the ground truth, as compared to those learnt through standard least squares. This is reflected in the normalised mean squared error values, that indicate good predictive accuracy of the models. Likewise, for f_3, TLS obtains lower NMSE than LS, despite the exact parametric form of the function being unavailable in this case. Interestingly, the worst performance for both methods is found when learning the quadratic function f_2. This appears to be due to the specific form of this function, where high input

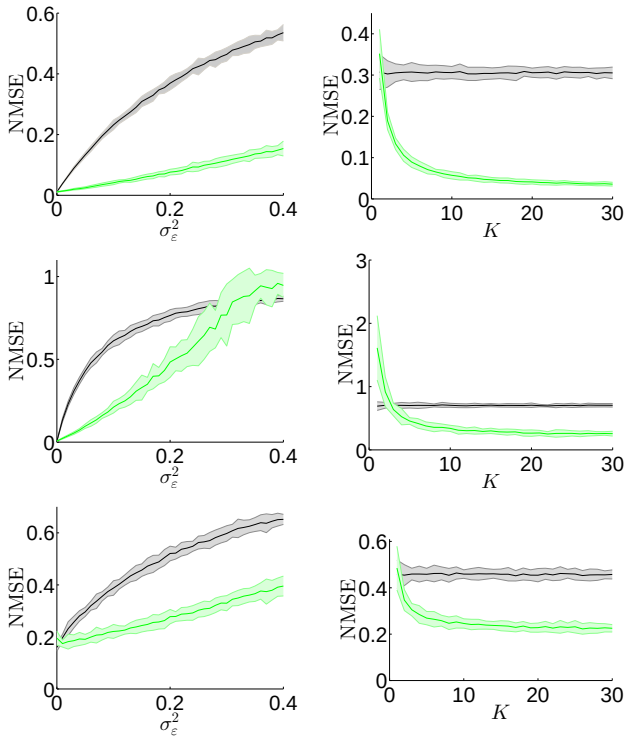

Fig. 4: Effects of varying σ_ε^2 (left) and $\mathcal{K}$ (right) on NMSE for linear (top), quadratic (middle) and sinusoidal (bottom) functions when learning with least squares (black) and total least squares (light green). Shown mean± two s.d. over 20 trials.

noise tends to cause overlap of data between the two 'arms' of the parabola, resulting in interference in learning.

These results are verified by examining the predictions of the learnt models over the range of training data. In Fig. 3, the predictions of the learnt TLS models for f_2 and f_3 are plotted, overlaid upon the ground truth predictions, and those of standard LS[2]. It can be seen that the models learnt with TLS are in good agreement with the underlying ground truth functions. In contrast, those learnt through standard LS suffer a bias towards zero, causing attenuation of the predictions and thereby higher errors.

To further assess the performance of learning, the experiment was repeated, varying (i) the noise in the data, $0 \leq \sigma_\varepsilon^2 \leq 0.4$, and (ii) the number of replicates available, $1 \leq \mathcal{K} \leq 30$. Note that, the former corresponds to increasing the 'slack' of the fabric-mounted sensor (since looser coupling between sensor and wearer is likely to result in larger motion artifacts), while the latter corresponds to differences in the size and density of the data set recorded during the calibration stage (ref. Sec. II-B).

The results for functions f_1, f_2 and f_3 are plotted in Fig. 4. There, it can be seen that, as expected, there is a decrease in accuracy for both TLS and LS as the noise level increases. However, the divergence of the TLS and LS lines indicates a

[2]Note that, for standard LS, the predictions extend over a wider range of inputs since the $w_n = x_n + \varepsilon_n$ usually extends beyond the maximum and minimum x_n due to the symmetrically distributed additive noise.

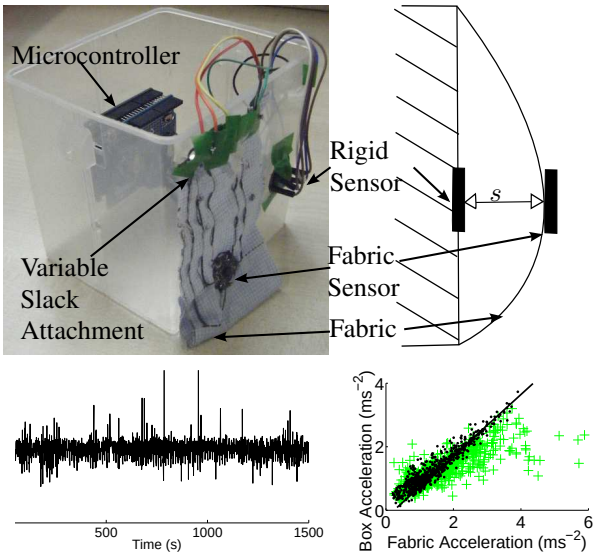

Fig. 5: Experimental set up (top). Difference of fabric and box acceleration magnitude against time (bottom left). Plot of box acceleration magnitude against that of the box (bottom right) showing ground truth mapping (black line), data with no slack (black circles) and data with slack (crosses).

much quicker degradation of performance for the latter. For the concave function f_2, this is particularly pronounced, an effect that may also be attributed to the non-monotonicity of the function: increasing noise causes greater overlap of data from the two arms of the parabola resulting in greater interference during learning.

Looking at the learning curves for varying $\mathcal{K}$ (Fig. 4, right), it can be seen that the error in the prediction NMSE for TLS drops rapidly as the number of replicates found in the data increases, levelling off at around $\mathcal{K} = 10$ for all functions. This suggests that the proposed approach is able to use the data efficiently to obtain a good fit. The LS line, in contrast, does not change significantly, despite the increase in the amount of data available.

B. Experiment

In the second evaluation, the proposed approach is tested for learning from physical data with the goal of predicting the movement of an object through space from a fabric-mounted sensor. The experimental platform used is shown in Fig. 5.

The platform consists of a pair of LilyPad ADXL335 tri-axial accelerometers mounted onto a plastic box. Of these, the first is attached rigidly, to provide a ground truth measurement of the box acceleration α, while the second is sewn onto a light-weight strip of cloth, and measures the fabric acceleration β. The cloth attachment is designed such that the slackness of the fabric s (defined as the maximum displacement from the box admitted by the fabric, see Fig. 5) can be adjusted between $s = 0\,cm$ (taut against the box) and $s = 6\,cm$.

During motion, readings from the two sensors are sampled synchronously at a rate of $23\,Hz$ using an Arduino Uno (Atmega-328P microcontroller, 16-bit ADC), and sent wirelessly to a PC base-station for analysis. The fabric-mounted sensor is connected to the Arduino using conductive thread to ensure minimal interference with the fabric motion (as might occur, for example, with use of wiring). Note that, while this reduces invasiveness of the sensor, it also adds further noise to the sensor readings [14], making the learning task in this experiment especially challenging.

As data, signals from the two sensors recorded during random shaking of the box for sessions of $60\,s$ each are used. The raw signals are preprocessed by converting the ADC values to acceleration in g, and then computing the acceleration magnitude for each time step. The acceleration magnitude is commonly used in clinical movement recording studies [15] and is used in this experiment to verify the usefulness of the proposed approach in such settings. The resultant data $\{\beta_n, \alpha_n\}_{n=1}^{\mathcal{N}}$ consists of $\mathcal{N} = 1350$ samples of box- and fabric-mounted sensor readings, respectively. These are randomly split into training and test data sets of equal size.

An sample data set is illustrated in Fig. 5 (bottom left and right), for displacement $s = 6\,cm$. On the left, the difference between fabric and box acceleration magnitude is shown against time, showing a significant amount of noise, much greater than would be expected from ordinary electrical noise. On the right, the fabric acceleration magnitude is plotted against that of the box, for two different slacks. In this case, the sensors are calibrated against one another so that there exists an identity relationship between the two (black line). When the slack is zero ($s = 0\,cm$, black dots) the sensor readings lie closely along the line, however for greater slack ($s = 6\,cm$, green crosses) a much broader distribution of readings is observed.

For learning, the proposed TLS approach is then used to train a linear model $\boldsymbol{\phi} := (x, 1)$ mapping the measured fabric accelerations β to box acceleration α on the training data. Note that, at the prediction stage, exact replicates of the form (14) are not available for forming the estimate ε_n (ref. Sec. III-B). Instead, approximate replicates are obtained by grouping similar values of β_n together into 50 discrete bins of equal size and spacing, and treating the data in these bins as the replicates.

In the following, results are reported for 5 trials of this experiment, in which data were recorded at slackness levels $s = 0$, 2, 4, 5, and $6\,cm$. As a baseline for comparison, the experiment is also repeated using the same model but learnt through the standard LS approach (ref. Sec. II-C).

In Fig. 6 (left) the mean prediction NMSE of the LS and TLS models are plotted against the slackness s. There it can be seen that, as the slackness increases, there is a gradual decrease in accuracy for both methods. This is in agreement with the simulation study (see Sec. IV-A), where it was seen that increasing noise (motion artifacts) resulted in a similar trend. However, comparing TLS with the standard LS approach, it is seen that the proposed approach consistently outperforms the latter across the range of slacknesses, with a more gradual decrease in accuracy in the face of greater

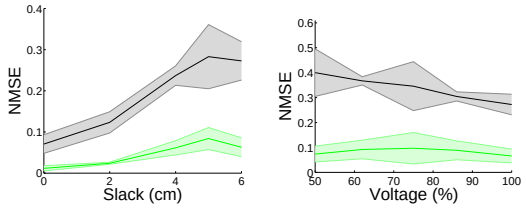

Fig. 6: NMSE of least squares (black) and total least squares (light green) of box motion when varying fabric slacknesses (left) and fabric sensor voltage (right). Results are mean±s.d. over 5 trials.

noise.

To further test performance, the experiment was also repeated using data from a decalibrated pair of sensors. This corresponds to the situation described in Sec. II-B, where a temporary (potentially heterogeneous) sensor is used for gathering data in a calibration stage.

In this experiment, the same test platform in Fig. 5 is used and data is collected following the same experimental procedure (with the slack of the fabric fixed at $s = 6\,cm$). However, to induce differences between the two accelerometers, the fabric sensor's input voltage is altered with a potentiometer, reducing the strength of the signal[3]. In the following, results are reported for data collected when the fabric sensor was supplied with (i) 4.3V (100% normal operating voltage), (ii) 3.7V (86%), (iii) 3.2V (75%), (iv) 2.7V (62%), and (v) 2.2V (50%). Note that, decalibrating the sensors in this way induces a non-identity mapping between the sensors, so that, for example, when the fabric sensor operates at 2.15V, the readings of the box sensor readings should be approximately twice the magnitude of those of the fabric sensor.

In Fig. 6 (right), the prediction NMSE of the test data against the fabric sensor input voltage are shown for 5 trials of this experiment. As can be seen, the proposed approach outperforms standard LS across the range of voltages (sensor calibration factors).

V. Discussion

In this paper, the application of statistical learning techniques in estimating motion through minimally invasive, e-textile sensing has been investigated. A crucial difficulty in this area, is the ability to deal with unpredictable motion artifacts introduced into sensing by the complex, unpredictable dynamics of the sensor fabric. The unsuitability of analytic modelling of the latter, both from the view point of model identification and computational tractability, suggests the use of learning approaches. It is also seen, however, that even learning techniques need to be adapted to this problem, due to the mismatches in noise assumptions found in standard learning models.

With a view to this, it has been proposed to exploit statistical methods from the errors-in-variables field to deal with the effects of stochastic perturbations to the sensory inputs in clothing-based sensing. An approach to model estimation and movement prediction has been presented, based on the use of total least squares regression for fitting, and estimation of query point inputs during prediction. Evaluation of the proposed approach in simulation has shown its ability to outperform standard regression in fitting non-linear models in face of significant input noise, and to efficiently make best use of replicate readings found in data. Experiments in learning from a physical cloth-mounted accelerometer to estimate the motion of an object in space, and compensate for calibration mismatches, show that it is always an improvement to apply orthogonal regression techniques over ordinary least squares. The computational efficiency of the proposed approach also makes it feasible for implementation onto an embedded device, making it an appealing option for the long-term gathering of data, e.g., in rehabilitation studies.

In future work, the performance of learning from higher dimensional data will be investigated, in order that more complex movements, such as walking, can be captured using sensors integrated into everyday items of clothing. Furthermore, the relationship between the structure of the e-textile fabric and the distribution of errors encountered will be investigated with a view to further improving learning and prediction accuracy.

[3]Note that, at each voltage, the sensors still have each axis calibrated at zero acceleration.

Acknowledgements

The authors would like to thank Rahul Pathak for building the experimental platform and collecting data.

References

[1] E. Jovanov, A. Milenkovic, C. Otto, and P. C. De Groen, "A wireless body area network of intelligent motion sensors for computer assisted physical rehabilitation," *J. Neuroeng. & Rehab.*, vol. 2, no. 1, p. 6, 2005.

[2] J. Barth, J. Klucken, P. Kugler, T. Kammerer, R. Steidl, J. Winkler, J. Hornegger, and B. Eskofier, "Biometric and mobile gait analysis for early diagnosis and therapy monitoring in Parkinson's disease," in *IEEE Trans. Eng. Medicine & Biology Soc.*, 2011.

[3] T. Shiratori and J. K. Hodgins, "Accelerometer-based user interfaces for the control of a physically simulated character," 2008.

[4] D. Roetenberg, H. Luinge, and P. Slycke, "Xsens MVN: full 6dof human motion tracking using miniature inertial sensors," Xsens Motion Technologies, Tech. Rep., 2009.

[5] J. Chen, K. Kwong, D. Chang, J. Luk, and R. Bajcsy, "Wearable sensors for reliable fall detection," in *IEEE Trans. Eng. Medicine & Biology Soc.*, 2005.

[6] J. Cancela, M. Pastorino, M. T. Arredondo, and O. Hurtado, "A telehealth system for Parkinson's disease remote monitoring. the PERFORM approach," in *IEEE Int. Conf. Eng. in Med. & Bio. Soc.*, 2013.

[7] R. Slyper and J. K. Hodgins, "Action capture with accelerometers," in *ACM SIGGRAPH/Eurographics Symp. Computer Animation*, 2008.

[8] S. Min and M. Tianlu, "Cloth animation based on particle model with constraint," in *W.S. Digital Media & Content Man.*, 2011.

[9] H. Li and Y. Wan, "An object-oriented system for dynamics-based 3d cloth simulation," in *IEEE Int. Conf. Inf. Sci. & Tech.*, 2012.

[10] B. Chen and M. Govindaraj, "A physically based model of fabric drape using flexible shell theory," *Textile Res. J.*, vol. 65, no. 6, pp. 324–330, 1995.

[11] R. J. Carroll, D. Ruppert, L. A. Stefanski, and C. M. Crainiceanu, *Measurement Error in Nonlinear Models: A Modern Perspective*. Chapman and Hall, 2006.

[12] Y. Nievergelt, "Total least squares: State of the art in numerical analysis," *SIAM Review*, vol. 36, no. 2, pp. 258–264, 1994.

[13] G. H. Golub and C. F. Van Loan, "An analysis of the total least squares problem," *SIAM J. Num. Analysis*, vol. 17, no. 6, pp. 883–893, 1980.

[14] T. Linz, C. Kallmayer, R. Aschenbrenner, and H. Reichl, "Embroidering electrical interconnects with conductive yarn for the integration of flexible electronic modules into fabric," in *IEEE Int. Symp. Wearable Computers*, 2005.

[15] M. V. Albert, K. Kording, M. Herrmann, and A. Jayaraman, "Fall classification by machine learning using mobile phones," *PLoS ONE*, vol. 7, no. 5, p. e36556, 2012.