

FabricTouch: A Multimodal Fabric Assessment Touch Gesture Dataset to Slow Down Fast Fashion

Temitayo Olugbade^{*†}, Lili Lin^{*}, Alice Sansoni^{*§}, Nihara Warawita^{*},
Yuanze Gan^{*}, Xijia Wei^{*}, Bruna Petreca[‡], Giuseppe Boccignone[§],
Douglas Atkinson[¶], Youngjun Cho^{*}, Sharon Baurley[‡], Nadia Bianchi-Berthouze^{*||}
^{*}University College London, [†]University of Sussex, [‡]Royal College of Art,
[§]Università degli Studi di Milano, [¶]Manchester Metropolitan University
^{||}nadia.berthouze@ucl.ac.uk

Abstract—Touch exploration of fabric is used to evaluate its properties, and it could further be leveraged to understand a consumer’s sensory experience and preference so as to support them in real time to make careful clothing purchase decisions. In this paper, we open up opportunities to explore the use of technology to provide such support with our *FabricTouch* dataset, i.e., a multimodal dataset of fabric assessment touch gestures. The dataset consists of bilateral forearm movement and muscle activity data captured while 15 people explored 114 different garments in total to evaluate them according to 5 properties (warmth, thickness, smoothness, softness, and flexibility). The dataset further includes subjective ratings of the garments with respect to each property and ratings of pleasure experienced in exploring the garment through touch. We further report baseline work on automatic detection. Our results suggest that it is possible to recognise the type of fabric property that a consumer is exploring based on their touch behaviour. We obtained mean F1 score of 0.61 for unseen garments, for 5 types of fabric property. The results also highlight the possibility of additionally recognizing the consumer’s subjective rating of the fabric when the property being rated is known, mean F1 score of 0.97 for unseen subjects, for 3 rating levels.

Index Terms—Dataset, fabric, gesture recognition, movement, touch, muscle activity, multimodal

I. INTRODUCTION

Clothing is central to living in the human society. From birth, there are few settings in our lives where clothing in some form is not relevant. Further, clothing not only has cultural, economic, and political significance [1], [2], but clothes are also intimate objects as they are worn on the body or otherwise stay closer to the body than most other objects that we interact with in everyday life [3]. Touch is the most basic sense used to interact with clothing as touch is critical in wearing, caring for, as well as evaluating, e.g. while shopping, the properties (such as softness) of the fabric of clothes [3]. Our work in this paper contributes a novel dataset called *FabricTouch*, containing fabric assessment touch gestures to support the development of applications that encourage more sustainable fashion consumption and usage. We further use the dataset to investigate automatic detection of the fabric properties being explored by a person and their subjective rating of the property, based on their touch behaviour.

The motivation behind our work is the urgent need to foster and support more careful and reflective purchasing of clothes.

Such consumer behaviour has been highlighted as one of the elements through which society can combat fast fashion and its negative impact on the environment and the humans who live in it [4]. Careful and reflective consumers would pay attention to quality (higher quality translating to longer durability), select items that are valued (e.g. for their comfort, and so more likely to be used for longer), and buy less. Touch is significant in informing this attentive purchase practice [3]. As Stanes proverbially puts it, “The eye shortlists the possibilities, but the hand has the power of veto.” [3](p. 232). In fact, tactile experience or evaluation that influences consumer preferences has long had an important place in shaping textile engineering [5]. Clothing designers also employ tactile evaluation in their design process [6]. For example, they explore fabric with their hand(s) to query its properties and its behaviour in response to touch, e.g. using fingers to hold the fabric by one of its corners to perceive how it drapes, or pulling the fabric between two hands to perceive its stretchability [6]. They also evaluate tactile experience with other parts of the body, usually the forearm, that function as ‘passive’ recipients of the experience while the (other) hand acts as the ‘active tool-hand’ [6]. Here, their touching the fabric is not at the forefront of exploration so much as the fabric touching them [6].

Although objective methods, such as textile surface testers that measure the roughness-smoothness of fabric, have been used to capture relevant fabric properties, tactile experience of fabric is understood to be largely subjective [3], [5], [7]. In this paper, we focus on the two subjective dimensions of tactile evaluation, tactile sensation and affect [7]. We capture sensory assessment of warmth, thickness, smoothness, softness, and flexibility, which tactile sensation is particularly sensitive to [5], [7], in our *FabricTouch* dataset. Findings of pilot studies carried out with the aim of understanding tactile gestures used by consumers in exploring clothing fabric support this choice. Our dataset further contains assessment of tactile pleasure (enjoyment), informed by the pilot studies.

The current paper makes the following contributions:

- 1) *FabricTouch* dataset: The first dataset on touch gestures for fabric assessment, it consists of hand muscle activity and arm movement data captured from 15 participants while exploring several different types of garments and using wearable sensor bracelets. It further contains RGB

images of each garment, labels of the fabric properties explored by a participant for each garment, the participant's ratings of the fabric with respect to each of the 5 properties, and the level of pleasure that they experienced during exploration of the fabric.

- 2) Investigation of the possibility of automatic detection of two different fabric exploration elements based on hand muscle activity and/or movement data: (1) automatic identification of the type of fabric property that a person is exploring (out of 5 property classes); and (2) automatic recognition of subjective rating (on a scale of 3 levels) of a fabric with respect to a given property.

II. RELATED WORK

A. Tactile Hand Gestures Datasets

Although a large number of hand gesture datasets exist, many are based on contactless gestures such as mid-air zoom [8], wave [9], or sign language gestures [10], [11]. Only a limited collection cover tactile hand gestures during human-object interactions. One category of these datasets use smart surfaces (i.e. surfaces embedded with sensors such as pressure sensors or touchscreens) to capture gestural data. For example, in [12] and [13], typing interactions are captured during the use of smartphones. Similarly, in [14], petting gestures (e.g. stroke, scratch, squeeze) on a zoomorphic object were collected.

The other category of tactile hand gesture datasets is based on passive objects, i.e. objects without sensors that capture gestures used in human-object interactions. One type of datasets in this category use wearable sensors. An example is the Ninapro dataset [15] captured during instructed hand gestures, with only a portion of the gestures involving interaction with an object (e.g. labels such as 'large diameter grasp'). Data was collected using surface electromyography (sEMG), i.e. muscle activity sensors, and accelerometers (with 10 electrodes around the muscle belly of the forearm; 1 on the wrist; and 2 on the biceps and triceps brachii), a data glove with 22 angular data channels, and an inclinometer placed on the wrist. Similarly, the csl-hdemg dataset has 168-channel sEMG data from the muscle belly of the forearm captured during gestures including controlled finger tapping while the hand and forearm are laid flat on a table. The use of data gloves that prevent haptic feedback during interaction with objects in both datasets limits their representation of tactile gestures. Indeed, gestures are shaped by the tactile experience of the person interacting with an object (amongst other factors), and so it is important to include natural haptic feedback while recording touch gestures. This is particularly critical in the context of clothing where changes in gestures may indicate how fabric feels [16] as it responds to touch exploration.

A different type of datasets use head-mounted cameras. For instance, the EPIC-KITCHENS dataset [17], collected from 32 participants during everyday kitchen activities in their homes, consists of videos which show hand gestures and objects. The dataset was annotated for actions such as 'pick up spoon' and also includes separate labels for objects (e.g. 'vegetables') and actions (e.g. 'spray'). EPIC-TENT [18] is a

similar dataset captured from 24 participants during outdoor camping tent assembly and contains both video and eye gaze data. It includes action labels such as 'pickup/open tent bag'. Another is MECCANO [19], which includes videos captured from 20 participants while they assembled a motorbike toy model, with both action (e.g. 'screw') and object (e.g. 'wheel rim') labels. A challenge with vision-based methods is the problem of occlusion that is particularly relevant to exploration of clothes where the hand can be placed under or inside a clothing item while exploring it. Further, unlike muscle activity sensors, video does not capture the effort of movement.

We contribute to the space of tactile hand gesture datasets by introducing a new dataset captured during natural gestures used while exploring clothes through touch and using sEMG and inertia sensors that do not interfere with natural tactile experience. We have chosen to use wearable sensors rather than video cameras to be able to detect both gesture shapes and effort rather than just their shapes. Wearable sensors further allow for more ubiquitous sensing, and similar types of sensors are increasingly being integrated in accessories (e.g. smartwatches) and on the body (e.g. skin tattoo sensors [20]). Nevertheless, insights from studies based on our dataset could inform design of fabric touch assessment applications that employ cameras. The dataset aims to contribute an understanding of how hand gestures reflect a person's experience of clothing fabric through touch. No other dataset is currently available for this purpose. This work is part of a larger project, Textiles Circularity Centre, aimed at designing technology to support circular and sustainable fashion behaviour.

B. Hand Gesture Recognition Methods with Wearable Sensors

Most works in this area have been based on data captured using sEMG sensors. For example, [15] used sEMG data from the Ninapro (DB1) dataset and [21]–[25] similarly used data captured using a commercial sEMG armband with dry electrodes placed around the muscle belly of the forearm. Some other studies [26], [27] used both sEMG data and data from inertia measurement units (IMUs), i.e. accelerometers and/or gyroscopes. A few used IMU data alone, e.g. [28].

Of these studies, several have been based on convolutional neural networks (CNNs). For instance, [15] used an architecture of CNNs and recurrent neural networks (RNNs) with attention applied along the time dimension on image data created from sEMG signals. They obtained 87% accuracy in discriminating between 50 gestures in the NinaPro dataset with hold-out validation. [29] obtained 88.2% accuracy for classification of 53 gestures from the same dataset based on a CNN, with leave-one-subject-out cross-validation. Other relevant studies are [21], [25] where CNNs were used with image spectrograms of sEMG signals. [21] obtained 97.8% accuracy respectively for classification of 7 simple, mid-air gestures, e.g. wrist flexion, based on hold-out validation. [24] similarly used a multilayer perceptron on hand-crafted features (e.g. mean absolute, gradient direction change) extracted from sEMG signals. They obtained 98.7% accuracy for differentiation of 6 mid-air gestures (e.g. double finger tap). Another related study

is [28] in which a restricted coulomb energy neural network and accelerometer data were used to differentiate between 10 mid-air digit writing gestures. They obtained 98.6% with 5-fold subject-dependent cross-validation.

Traditional learning algorithms, particularly support vector machines (SVMs), have also been widely used. [22], for example, used SVMs with hand-crafted features extracted from both time and frequency domains of sEMG signals (e.g. mean absolute, amplitude spectrum). Based on subject-dependent leave-one-out cross-validation, they obtained 16.4% accuracy for the classification of 40 instructed hand and finger gestures such as mid-air pinching. [30] also used SVMs and with mean sEMG as features, they achieved 90% accuracy for 4 mid-air gestures (e.g. hand close vs open). In [26], second order polyfit features were extracted from both sEMG and Euler angle data with accuracy between 88% and 93% for 7 mid-air gestures (such as clockwise hand tilt) based on subject-dependent hold-out validation. SVMs were compared with k-nearest neighbour, discriminant analysis, Naives Bayes, and random forest models in [23]. The authors used features such as mean absolute, gradient direction change, and waveform length for sEMG signals. The best performance of 96.4% (standard deviation=4.5%) for 3 simple mid-air gestures, e.g. fist, was achieved with the random forest based on cross-validation with separate subjects in the training and test sets.

The studies above highlight the feasibility of automatic recognition of hand gestures based on sEMG and/or IMU data using either deep learning or traditional machine learning methods. However, there is still limited understanding of how more complex gestures like rubbing to explore the roughness of fabric can be modelled from such data. Real world gestures like this can involve the two hands playing different roles [6] and include interrupts (e.g. scratching the nose) in between, unlike the instructed single-hand gestures and more controlled settings in the studies discussed above. Our new dataset includes sEMG and IMU data captured from both hands while participants explore fabric properties without instruction about what gestures to use in their exploration. This represents a much more complex classification task than existing studies.

III. THE FABRICTOUCH DATASET

This study was approved by the local research ethics committee, and participants gave informed consent that covers sharing of pseudonymised data with the research community. Researchers can access the dataset by emailing the last author.

Data was captured using 2 gForcePro+ armbands (see Figure 1) [31], each of which has 8 sEMG electrodes and a 9-axis IMU sensor. In this study, the armbands were worn on each arm such that the placement on one arm was a vertical flip of the placement on the other. We developed a custom mobile application to enable participants to perform data capture on their own. The app allowed the researcher to set up data capture sessions that the participants would implement. The researcher could set a duration for all exploration sessions, the fabric properties to explore and rate, and the rating scale to use for each property. For this study, the participants explored




Fig. 1: Top - Left to right: Sensor placement for our *FabricTouch* dataset; examples from Phase II collection showing assessment of smoothness (finger caress), warmth (hand inside sock), and flexibility (pulling opposite side) respectively. Bottom - Examples of garments from Phase I collection.

5 properties (warmth, thickness, smoothness, softness, and flexibility) and rated each on a 7-item likert scale, e.g. from ‘hard’ to ‘very soft’ for softness. In addition, they rated the level of pleasure on the same scale.

A. Collection Phases

There were two data collection phases.

1) Phase I Collection - Self-Selected Clothes at Home:

Lockdown and social distancing restrictions forced us to move data capture out of lab settings into participants’ homes. However, data based on clothes selected by participants from their wardrobes and captured in their homes have an advantage as we expect it to have fostered gestures more representative of natural exploration of clothing. Participants could select any clothing from their wardrobe (typically jeans, jumpers, coats, and skirts). The only constraint given was that the participant had to choose as wide a variety of material types as possible. This was to ensure that many different fabric properties were represented across a variety of garment and material types, to capture the real-life variability of clothing handling gestures.

In this data collection phase, data was captured by the participants themselves. For each participant, the researcher was present via videoconferencing in the first round of capture to guide the participant in case they experienced any problems; the participant recorded data on their own in the second round, which took place on a different day. For the two rounds, data capture was completed with the participant seated, and with the garment on a table in front of them. Participants explored at least 6 garments in each round.

9 participants (8 female and 1 male) took part in this phase, all students with ages between 20 and 27. They were rewarded with £35 gift vouchers for their participation.

2) Phase II Collection - Pairs of Socks in the Lab: We

extended the Phase I data with gestures based on exploration of socks, to capture both a clothing item not covered in the initial

phase as well as between-subject variations for the same item. We invited a different set of 6 participants to our lab (3 female and 3 male), 5 students and 1 researcher; they explored six pairs of socks of different material composition and thickness. They received £15 gift vouchers for their participation.

B. Collection Procedure

At the start of each data capture round in either phase, the participant recorded two EMG baselines: relaxed hands and tight fists. The purpose of the baseline was twofold. First, it was for use to check that the armbands were placed correctly, e.g. fitted to the forearm with complete contact, and ascertain the validity of each capture round. Second, it was to enable normalisation of the EMG sensor data to account for differences, e.g. in forearm morphology, between participants.

Next, for each garment, the participant explored each of the 5 fabric properties in turn. As this was not a controlled study, the property exploration was done in fixed order of smoothness, thickness, warmth, flexibility, and softness in the Phase I collection, rather than being randomised. Smoothness and softness were set as far from one another as possible because previous unpublished studies suggested that consumers confuse the two constructs. Nevertheless, in Phase II, both the order of presentation of the fabric and exploration of the properties were randomised. In the Phase I collection, the garment fabric was explored through touch for 20 seconds, which previous studies found to be adequate. This was lowered to 15 seconds in Phase II based on the discovery that participants found 20 seconds too long a period to explore a single fabric property. No instructions were given to participants on how to touch fabric to explore any of the 5 properties of interest. Rather, they were free to touch the garment in any way that they considered suitable to perceive the given property. 114 different garments were explored in total.

After exploration of a given property for a given garment, the participant would rate the garment on the scale for that property, e.g. scale from 'rough' to 'smooth' for the smoothness property. After exploring and rating all 5 properties for a garment, the participants would additionally rate how much they enjoyed touching the garment fabric, on a 7-item Likert scale from 'not at all' to 'very much'.

C. Data Summary

There were a total of 888 exploration instances, 672 and 216 in Phases I and II respectively. For each instance, the following sensor data were recorded for each of two hands: 8-channel sEMG; 3D Euler angle; 4D quaternion; 3D acceleration; 3D angular velocity; and 3D magnetic field strength and direction.

For the Phase I instances, the ratings for each of the 5 properties ranged from 1 to 7 with mean (and standard deviation) of 4.20 (1.8), 3.12 (1.7), 3.95 (1.6), 3.37 (1.9), and 4.26 (1.6) for smoothness, thickness, warmth, flexibility, and softness respectively. The rating for enjoyment ranged between 1 and 7, mean=4.77, standard deviation=1.4.

For the Phase II instances, the ratings for the properties also ranged from 1 to 7 except for flexibility whose minimum was

2. The mean (and standard deviation) for smoothness, thickness, warmth, flexibility, and softness was 3.92 (1.64), 4.14 (1.53), 4.97 (1.83), 4.83 (1.52), and 4.44 (1.9) respectively. The enjoyment rating ranged from 1 to 7 with mean of 4.44 and standard deviation of 1.8.

IV. AUTOMATIC RECOGNITION OF EXPLORED FABRIC PROPERTY

To understand the problem space and set a baseline for the community, we investigated the possibility of automatic detection of fabric property being explored based on our novel dataset. As a first step, we used random forests [32] of 100 trees for the task. Since this work focused on the properties and did not include the enjoyment exploration, there were only 480 exploration instances for Phase I (data for one participant was additionally excluded due to error in the signals recorded) and 180 exploration instances for Phase II.

A. Preprocessing and Feature Extraction

For the Phase I data, only the last 15 seconds of each exploration instance was used, for consistency with the Phase II data where each exploration was for 15 seconds. Each 15-second exploration was then segmented into 9 non-overlapping segments. This led to a total of 4,320 segments for the Phase I data and 1,620 segments for the Phase II data. For each segment, we used data from the sEMG and IMU sensors.

For the sEMG, we first performed full-wave rectification as is standard. We then normalized each signal by dividing by its maximum, done separately for each channel, hand, and exploration instance. For the IMU, we focused on the acceleration and quaternion signals. We computed instantaneous velocity and jerk from the acceleration and computed instantaneous angular velocity, acceleration, and jerk from filtered quaternion signals. For each of the angular velocity, acceleration, and jerk, in addition to the values for each axis, we also computed a total across axes. We first filtered the quaternion signals to remove outliers outside the range of unit quaternions, [-1, 1]. For each of the 33 resulting signals (8 sEMG, 4 quaternion, 3 linear acceleration, 3 linear velocity, 3 linear jerk, 4 angular velocity, 4 angular acceleration, 4 angular jerk), we extracted the maximum, mean, and standard deviation in each segment resulting in 99 features for each hand and 198 features in total per segment (48 sEMG features and 150 IMU features).

We replaced missing values, e.g. due to error in the recording individual signals, with the mean across segments and then normalized the resulting feature table using z-score scaling to zero mean and unit standard deviation.

B. Results and Discussion

1) *Can the property assessed be recognized from sEMG/IMU data?*: The top half of Table I shows the results of automatic recognition of explored fabric property based on leave-one-clothing-out cross-validation (LOCOCV). We compared use of all features with use of sEMG, quaternion-based IMU, and acceleration-based IMU features separately.

TABLE I: Explored fabric property recognition results (LOCOCV & LOSOCV) - Phase I data only

Features	F1 score - LOCOCV				
	Smoothness	Thickness	Warmth	Flexibility	Softness
sEMG	0.58	0.51	0.53	0.65	0.50
IMU (quaternion-based)	0.56	0.50	0.50	0.54	0.48
IMU (acceleration-based)	0.56	0.50	0.46	0.58	0.39
All	0.64	0.60	0.56	0.68	0.56

Features	F1 score - LOSOCV				
	Smoothness	Thickness	Warmth	Flexibility	Softness
sEMG	0.46	0.42	0.25	0.53	0.34
IMU (quaternion-based)	0.45	0.35	0.14	0.29	0.32
IMU (acceleration-based)	0.44	0.43	0.19	0.46	0.24
All	0.55	0.48	0.24	0.57	0.43

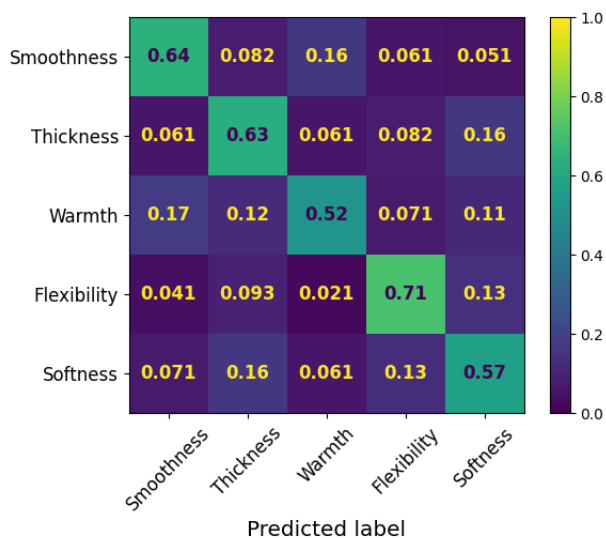


Fig. 2: Confusion matrix for automatic recognition of explored fabric property (LOCOCV, Phase I data - all features).

The results show performance well above chance-level classification (F1 score = 0.2) for all 5 properties and for each feature set. The best performance was obtained when both sEMG and IMU features were combined (F1 score = 0.61, mean across properties). The use of the sEMG features alone was better than either of the IMU feature sets alone especially for *warmth* and *softness*, and the quaternion-based features led to better performance than the acceleration-based features for those same two properties. The acceleration-based features were slightly better for *flexibility* than the quaternion-based features. It is not surprising that the sEMG features are superior to the IMU features given that the sEMG data is able to capture hand- and finger-level movements whereas the IMU data we collected can only represent whole-arm movements. In addition, sEMG data captures movement effort which is related to affect. The better performance of quaternion-based features compared to acceleration-based features for *warmth* and *softness* further suggests that assessment of rotational

arm movements is more helpful for recognizing these than assessment of translational arm movements.

Flexibility was the easiest exploration to recognize (except when the quaternion-based features were used). The higher performance for *flexibility* is likely a result of the distinctiveness of the prevalent gesture, pulling the fabric with both hands, used by the participants to explore this property. *Softness* consistently had the worst performance. As the confusion matrix in Figure 2 shows, it was typically confused with *thickness* and *flexibility*. Although the confusion with exploration of thickness could be explained with participants' use of two-/three-finger rubbing gestures for both properties, the overlap with exploration of flexibility is not clear.

2) *Is generalization to unseen participants feasible?*: We sought to understand how well automatic detection of explored fabric properties based on sEMG and IMU data generalized to unseen participants, and so we repeated the experiments above but using leave-one-subject-out cross-validation (LOSOCV). The bottom half of Table I shows the results.

As can be seen in the table, although the use of sEMG features alone or together with the other features still had better than chance level performance, as expected generalization to unseen subjects was worse than generalization to unseen garments. Recognition of *warmth* exploration was only marginally better than chance level for both the sEMG features alone and all features, and it was lower than chance level when only IMU features were used. As with LOCOCV, exploration of *flexibility* was the best recognized except when quaternion-based features were used, whereas recognition of *smoothness* exploration was consistently much better than chance level.

The noticeably poorer performance for *warmth* in unseen subjects versus in unseen garments, compared with the other properties, suggests stronger idiosyncrasies for this property assessment. This could be due to variations in the interpretation of its definition. For example, it is assessed with respect to temperature by some people, but in terms of comfort and coziness by others. In the first type of interpretation, relevant gestures could be touching and holding the garment or otherwise wrapping the garment around the hand or inserting the hand inside the garment. The second leads to different

TABLE II: Explored fabric property recognition results (All features) - Phase I & II data

Cross-validation strategy	F1 score				
	Smoothness	Thickness	Warmth	Flexibility	Softness
LOSOCV	0.44	0.34	0.29	0.38	0.22
LOCOCV	0.46	0.39	0.36	0.49	0.27
Leave out one subject-cloth combination	0.60	0.53	0.50	0.61	0.45

gestures, e.g. grabbing the garment and holding it against oneself. Warmth could also be associated with garments that a person has strong emotional attachment to, e.g. due to fond memories. Such interpretation differences highlight affective meanings of fabric property constructs and are also in line with the knowledge that bodily gestures provide a way to predict and confirm expected interoceptive and exteroceptive sensations [33]–[35]. It is hence important that datasets represent the different meanings assigned to constructs, to reflect the complexity of the language of (affective) touch of clothing.

3) *How does inclusion of a different type of garment affect performance?*: The main difference between the Phase I and II data collection is the type of garments explored. In Phase I, participants explored a variety of garments such as jeans, jumpers, coats, and skirts, while Phase II was focused on socks. Compared to most other types of garments, for example those designed to cover the upper and/or lower parts of the body (e.g. shirts, dresses), socks have a peculiar form factor that can invite a different type of exploration. For instance, socks can fit on the hands like a glove and can also easily be turned inside out while exploring their inner layer properties. Further, their size allows them to be bunched up as a whole and fit into the palm, e.g. in a crunching gesture. Thus, it is valuable to investigate how combining gestures of socks exploration (Phase II data) with gestures exploring other types of garment (Phase I data) would affect performance.

Similar to the evaluation carried out on the Phase I data above, we performed LOSOCV and LOCOCV. As can be seen in Table II, generalization to unseen subjects or garments was worse when both Phase I and II data were used together. *Softness* was the property with the largest decrease in recognition performance. In the Phase I data, each garment explored was unique to the respective participant unlike the Phase II data where all the participants explored the same pairs of socks. So, beyond the LOCOCV done here, we also performed leave-one-subject-clothing-combination-out CV, i.e. a variant of LOCOCV where only explorations of a given pair of socks by a single participant in the Phase II data instances were held out at a time during cross-validation. In the ordinary LOCOCV, explorations of the given pair of socks by all participants are held out together in one fold. The adapted LOCOCV, i.e. leave-one-subject-clothing-combination-out, led to much better performance across all properties although performance was still lower than using Phase I data alone.

These findings suggest that generalization is better with both garment types and subjects seen during training. We have already discussed the difficulty of generalizing to unseen subjects in Section IV-B2. With unseen garment types, beyond

the differences between types of garments (e.g. skirts vs socks) in terms of their shapes and sizes, differences in tactile experience due to the properties (e.g. dryness, lightness) and composition (e.g. wool vs nylon) of their fabric itself can lead to additional differences in tactile exploration of any single property. Empirically, we found statistically significant differences in perceived fabric properties between the two data collection phases based on linear mixed model analysis, i.e. accounting for repeated measures from the same subjects. In particular, differences were significant for *warmth* ($p=0.018$, $t=-2.66$, $n=145$), *thickness* ($p=0.033$, $t=-2.36$, $n=145$), and *flexibility* ($p=0.02$, $t=-3.61$, $n=143$). No statistical difference was found for *smoothness* and *softness*.

V. AUTOMATIC RECOGNITION OF SUBJECTIVE RATING OF FABRIC PROPERTIES

We further explored the possibility of automatic detection of the subjective rating of a given fabric property for a baseline. Preliminary experiments using random forest showed poor discrimination between the ratings (e.g. average F1 score of 0.36 for 3 levels). Thus, we explored the use of RNNs that allow capture of temporal information, in particular long short-term memory neural networks (LSTMNNs) [36], [37]. We used a simple model with a single LSTM layer followed by a 3-layer multilayer perceptron (MLP) with hidden units of sizes 20 and 10 respectively. The LSTM layer was shared by each of the two hands and the encodings for the two hands were then concatenated before being processed by the MLP. We compared this LSTM-based model with a 3-layer MLP without the LSTM. We additionally compared with random forest. We evaluated the models using LOSOCV and LOCOCV.

A. Preprocessing and Feature Extraction

We used both Phase I and II data for the experiments described in this section. The same processing methods described in Section IV-A were applied except that instead of segmentation of each exploration instance (duration = 15 seconds) into 9 segments, segmentation into 15 was used. For each segment, we extracted the same types of features as in Section IV-A, but we did not include the total across axes for angular velocity, acceleration, and jerk. Thus, there were only 30 signals used here. With 3 features per signal (maximum, mean, standard deviation), there were a total of 90 features for each hand. To explore temporal characteristics, for the LSTMNN and MLP models, we extracted these features from three non-overlapping slices within each segment, making 540 features for these two models. In addition to these sensor-based features, we included 5 categorical features to represent

TABLE III: F1 scores for automatic recognition of subjective ratings of explored fabric properties

Model	F1 scores - LOCOCV			F1 scores - LOSOCV		
	Low	Medium	High	Low	Medium	High
Random forest	0.25	0.62	0.37	0.18	0.64	0.14
LSTMNN	0.77	0.86	0.79	0.63	0.79	0.67
MLP	0.97	0.98	0.97	0.97	0.98	0.97

the explored fabric property such that each was a ‘1’ for the property explored in a given exploration segment and ‘0’ otherwise. We used the ground truth here, but in future work, the explored fabric property information will come from automatic recognition such as presented in Section IV.

Rather than discrimination between all 7 property intensities captured in our dataset, we focused on detection of 3 levels as a first step. This lower level of granularity has the advantage of reducing the effect of between-subject variations in the use of the rating scale. There is also the additional advantage of having more data per class. We re-coded ratings of 1 and 2 on the original scale as low level, 3 to 5 as medium level, and 6 and 7 as high level. There were 1,500, 3,210, and 1,590 instances for the low, medium, and high classes.

B. Results and Discussion: Can subjective rating of fabric properties be recognized?

Table III shows the results obtained comparing the 3 models based on both LOCOCV and LOSOCV.

The random forest model performs very poorly with worse than chance level recognition for the two minority classes (low and high levels) and has mean F1 scores of 0.41 and 0.32 with LOCOCV and LOSOCV respectively.

Both the MLP and LSTMNN perform far better and well above chance level classification with mean F1 scores of 0.97 and 0.81 respectively based on LOCOCV and mean F1 scores of 0.97 and 0.70 respectively for LOSOCV. As can be seen, the MLP model has considerably higher performance than the LSTMNN, both with LOCOCV and LOSOCV. The larger number of parameters for the LSTMNN (which includes a LSTM layer in addition to the same number of fully connected layers as the MLP model), given the size of the dataset ($n = 6,300$), may have contributed to its lower performance. Another critical difference between the two models is where fusion of the property information features and the sensor-based features occur. For the MLP model, the fusion is done at the input level, whereas for the LSTMNN it occurs mid level after the LSTM layer encoding of the sensor-based features. Given the importance of the property information, this difference may have influenced the performance of the models. The results also suggest that static touch postures are more informative for recognition of the rating levels than the temporal relationships between postures.

VI. ETHICAL IMPACT STATEMENT AND CONCLUSION

We present a novel multimodal hand gestures dataset of naturalistic fabric touch assessment interactions, *FabricTouch*. The dataset includes 8-channel sEMG data and 16-channel

IMU data from each of the two forearms for 15 people while they explored a wide variety of garments in their homes or in lab settings. The dataset further contains labels for garment fabric properties explored in each interaction, subjective ratings of the given property for the given garment, and the level of pleasure experienced in exploring the fabric through touch.

The investigation of automatic detection of details of such tactile fabric interaction based on this dataset reveals the possibility of equipping technology with the capability of automatic recognition of the fabric property being explored, mean F1 score of 0.61 for 5 classes of properties being explored for garments not included in the training data. For automatic identification of 3 levels of the consumer’s rating of the property based on their touch behaviour and given that the property being explored is known, we obtained mean F1 score of 0.97 for people that the system was not trained on.

Our work provides groundwork for addressing sustainability in the fashion industry through the consumer. Fast fashion is one of the world’s most wasteful and polluting industries, driving impulsive purchase of clothes that are quickly disposed of by consumers, returned (then disposed of by shops) or unused [39]. Our dataset was created as a starting point in developing a fabric touch databank that enables the creation of affective-touch-aware technology which invites tactile exploration to foster reflection and awareness during purchase or in caring for clothes. Further, fostering engagement in touch interactions with clothes could increase bonding with the clothes [40]. The larger project on textile circularity that this work sits within covers the broader textile value chain.

The work is currently limited with respect to diversity and inclusivity as the dataset does not currently capture the diversity of touch abilities. For example, it does not capture touch exploration by people with only one hand or no hands, people with upper limb movement disorders such as cerebral palsy or stroke, or people with sensory disorders. It is expected that people with such conditions may have alternative gestures for exploring clothing. Further, although our data was captured from two different countries (China, UK), it is important to note that our work is largely based on understanding of consumer behaviour and fabric handling for limited geographies and cultures. As noted in our introduction, clothing has cultural, economic, and political significance that may be reflected in how people purchase, use, experience, care for, and discard clothes in different countries and regions. It is, for instance, unknown if (and to what extent or how) the problem that we address is relevant in regions beyond rich, industrialized countries. Our long-term aim is to gradually build a more inclusive dataset for the UK and in parallel

work with other researchers in the area to develop a more geographically inclusive understanding of clothing handling, purchase, and use behaviours.

The dataset and findings of the paper are significant contributions given that research on touch gestures is still limited in comparison to other modalities, more so for fabric experience. They invite the community to build recognition models for tactile exploration of clothing fabrics and lay the groundwork for further studies to capture and support a consumer's sensory experience during clothing exploration.

ACKNOWLEDGMENT

This work was supported by the EU Future and Emerging Technologies Proactive H2020 programme (Grant 824160: EnTimeMent) and the UKRI National Interdisciplinary Circular Economy Research programme (Grant EP/V011766/1: Textiles Circularity Centre) and Consumer Experience Digital Tools for Dematerialisation for the Circular Economy project (Grant EP/V042289/1). For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

REFERENCES

- [1] O. A. Ajani, Aso Ebi: The dynamics of fashion and cultural commodification in Nigeria. *J. Pan Afr. Stud.*, 5(6), pp. 108-118. 2012.
- [2] Z. Mir-Hosseini. The politics and hermeneutics of hijab in Iran: From confinement to choice. *Muslim World J. Hum. Rights*, 4(1). 2007.
- [3] E. Stanes. Clothes-in-Process: Touch, Texture, Time. *Textile*, 17(3), pp. 224-245. 2019.
- [4] R. Bick, E. Halsey, and C. C. Ekenga. The global environmental injustice of fast fashion. *Environmental Health*, 17(1), pp. 1-4. 2018.
- [5] I. Ciesielska-Wrobel and L. Van Langenhove. The hand of textiles—definitions, achievements, perspectives—a review. *Text. Res. J.*, 82(14), pp. 1457-1468. 2012.
- [6] B. Petreca, S. Baurley, and N. Bianchi-Berthouze. How do designers feel textiles?. In *Proc. Int. Conf. Affect. Comput. Intell. Interact.*, pp. 982-987. 2015.
- [7] D. Atkinson, S. Baurley, B. Petreca, N. Bianchi-Berthouze, and P. Watkins. The tactile triangle: a design research framework demonstrated through tactile comparisons of textile materials. *J. Des. Res.*, 14(2), pp. 142-170. 2016.
- [8] M. Schak and A. Gepperth. On Multi-modal Fusion for Freehand Gesture Recognition. In *Proc. Int. Conf. Neural Netw.*, pp. 862-873. 2020.
- [9] A. Gabriel, S. Coşar, N. Bellotto, and P. Baxter. A dataset for action recognition in the wild. In *Proceedings of Annual Conference Towards Autonomous Robotic Systems*, pp. 362-374. 2019.
- [10] O. Özdemir, A. Kındıroğlu, N. Camgöz, and L. Akarun. Bosphorus-sign22k sign language recognition dataset, arXiv. 2020.
- [11] J. Imran and B. Raman. Deep motion templates and extreme learning machine for sign language recognition. *The Visual Computer*, 36(6), pp.1233-1246. 2020.
- [12] M. Exposito, J. Hernandez, and R. Picard. Affective keys: towards unobtrusive stress sensing of smartphone users. In *Proc. Int. Conf. Human-Comput. Interact. Mobile Devices Services adjunct*, pp. 139-145. 2018.
- [13] S. Ghosh, N. Ganguly, B. Mitra, and P. De. Evaluating effectiveness of smartphone typing as an indicator of user emotion. In *Proc. Int. Conf. Affect. Comput. Intell. Interact.*, pp. 146-151. 2017.
- [14] A. Flagg and K. MacLean. Affective touch gesture recognition for a furry zoomorphic machine. In *Proceedings of International Conference on Tangible, Embedded and Embodied Interaction*, pp. 25-32. 2013.
- [15] M. Atzori, A. Gijssberts, C. Castellini et al. Electromyography data for non-invasive naturally-controlled robotic hand prostheses. *Scientific data*, 1(1), pp. 1-13. 2014.
- [16] B. Petreca, C. Saito, S. Baureley, D. Atkinson, X. Yu, and N. Berthouze. Radically relational tools: A design framework to explore materials through embodied processes. *Int. J. Des.*, 13(2), pp.7-20. 2019.
- [17] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. 2018. Scaling egocentric vision: The EPIC-KITCHENS dataset. In *Proc. European Conf. Comput. Vis.*, pp. 720-736. 2018.
- [18] Y. Jang, B. Sullivan, C. Ludwig et al. EPIC-Tent: An egocentric video dataset for camping tent assembly. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops*. 2019.
- [19] F. Ragusa, A. Furnari, S. Livatino, and G. Maria Farinella. The MEC-CANO dataset: Understanding human-object interactions from egocentric videos in an industrial-like domain." In *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, pp. 1569-1578. 2021.
- [20] D. Kireev, S. Ameri, A. Nederveld et al. Fabrication, characterization and applications of graphene electronic tattoos. *Nat. Protoc.* 16(5), pp. 2395–2417. 2021.
- [21] U. Côté-Aillard, C. Fall, A. Campeau-Lecours, C. Gosselin, F. Lavolette, and B. Gosselin. Transfer learning for sEMG hand gestures recognition using convolutional neural networks. In *Proc. IEEE Int. Conf. Systems Man Cybernetics*, pp. 1663-1668. 2017.
- [22] F. Kerber, M. Puhl, and A. Krüger. User-independent real-time hand gesture recognition based on surface electromyography. In *Proc. Int. Conf. Human-Comput. Interact. Mobile Devices Services*, pp. 1-7. 2017.
- [23] M. Wahid, R. Tafreshi, M. Al-Sowaidi, and R. Langari. Subject-independent hand gesture recognition using normalization and machine learning algorithms. *J. Comput. Sci.*, 27, pp 69-76. 2018.
- [24] Z. Zhang, K. Yang, J. Qian, and L. Zhang. Real-time surface EMG pattern recognition for hand gestures based on an artificial neural network. *Sensors*, 19(3170). 2019.
- [25] L. Chen, J. Fu, Y. Wu, H. Li, and B. Zheng. Hand gesture recognition using compact CNN via surface electromyography signals. *Sensors*, 20(672). 2020.
- [26] A. Kundu, O. Mazumder, P. Lenka, and S. Bhaumik. Hand gesture recognition based omnidirectional wheelchair control using IMU and EMG sensors. *J. Intell. Robot. Syst.*, 91(3), pp. 529-541. 2018.
- [27] S. Jiang, B. Lv, W. Guo et al. Feasibility of wrist-worn, real-time hand, and surface gesture recognition via sEMG and IMU sensing. *IEEE Trans. Industr. Inform.*, 14(8), pp. 3376-3385. 2018.
- [28] M. Kim, J. Cho, S. Lee, and Y. Jung. IMU sensor-based hand gesture recognition for human-machine interfaces. *Sensors*, 19(3827). 2019.
- [29] W. Wei, Q. Dai, Y. Wong, Y. Hu, M. Kankanhalli, and W. Geng. Surface-electromyography-based gesture recognition by multi-view deep learning. *IEEE Trans. Biomed. Eng.*, 66(10), pp. 2964-2973. 2019.
- [30] M. Tavakoli, C. Benussi, P. A. Lopes, L. B. Osorio, and A. T. de Almeida. Robust hand gesture recognition with a double channel surface EMG wearable armband and SVM classifier. *Biomedical Signal Processing and Control*, 46, pp. 121-130. 2018.
- [31] Oymotion Technologies Co. Ltd. gForcePro+ EMG Armband. <http://www.oymotion.com/en/product32/149>
- [32] L. Breiman. Random forests. *Machine learning*, 45(1), pp.5-32. 2001.
- [33] L. Barrett and M. Bar. See it with feeling: affective predictions during object perception. *Phil. Trans. R. Soc. B*, 364, pp. 1325-1334. 2009.
- [34] K. Hoemann, R. Wu, V. LoBue et al. Developing an understanding of emotion categories: Lessons from objects. *Trends Cog. Sc.*, 24(1), pp. 39-51. 2020.
- [35] K. Quigley, S. Kanoski, W. Grill, L. Barrett, and M. Tsakiris. Functions of interoception: From energy regulation to experience of the self. *Trends Neurosc.*, 44(1), 29-38. 2021.
- [36] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Comput.*, 9(8), 1735–1780. 1997.
- [37] F. A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: continual prediction with LSTM. In *Proceedings of the International Conference on Artificial Neural Networks*, pp. 850–855. 1999.
- [38] C. Wang, Y. Gao, A. Mathur et al. Leveraging Activity Recognition to Enable Protective Behavior Detection in Continuous Data. *Proc. ACM Interact. Mob. Wearable Ubiquitous*, 5(2), pp. 1-27. 2021
- [39] UN Environment Programme. Sustainable Fashion: communication strategy 2021-2024. <https://www.unep.org/resources/publication/sustainable-fashion-communication-strategy-2021-2024>. 2021.
- [40] B. Petreca, S. Baurley, K. Hesseldahl, A. Pollmann, and M. Obrist. The Compositor Tool: Investigating Consumer Experiences in the Circular Economy. *Multimodal Technol. Interact.* 6(24), pp. 1–22. 2022.