Ultrasound Imaging For Material Characterization

Tamas Aujeszky[1,], Georgios Korres[1,], Mohamad Eid[1,]

Department of Engineering, New York University Abu Dhabi, Saadiyat Island, United Arab Emirates

Abstract

This paper presents a study to classify families of materials using ultrasound reflection, with the aim of developing a nondestructive, contactless method to extract haptic properties of materials. A range of Sorbothane samples are subjected to ultrasound stimulation. Reflected data is captured, processed and subjected to support vector machine-based one-vs-all classification. Results show a high correlation between the classified data and the sample classes, when a part of the data in the same session is used as training. The high classification accuracy is retained when multiple data sessions are mixed. However, the classification accuracy drops when samples from new (untrained) sessions are introduced. It is suggested that a wide range of training data would provide an adequate basis for accurate classification of sessions.

Keywords: Haptic scanner, Ultrasound stimulation, Machine learning

1. Introduction

Development of sensors that are capable of recording the same sensations that humans can feel is a topic that has long been explored. Not only is this goal fueled by the aspiration of building humanlike machines, but also by the idea to record and store these sensations with the aim of reproducing the objects that generated them in high detail and quality.

The advent and spread of color cameras provided us with the ability to record visual data. The appearance of the Microsoft Kinect [1] sensor brought about affordable sensors for depth data. We are at a point where we are able

Email addresses: `tamas.aujeszky@nyu.edu` (Tamas Aujeszky), `george.korres@nyu.edu` (Georgios Korres), `mohamad.eid@nyu.edu` (Mohamad Eid)

Preprint submitted to International Journal of Computing and Information Sciences September 3, 2016

to scan and transmit the exact shape and color of an object with great fidelity. This, however, is not sufficient to completely describe the object, as other information is needed related to properties that humans deduce by relying on their sense of touch. Following this principle, several contact-based systems have been introduced to measure haptic surface properties. This approach, however, has certain fundamental limitations regarding automation as well as preservation of samples.

In this paper we present a nondestructive, contactless haptic scanning approach based on ultrasonic excitation and reflection. The rest of the paper is organized as follows: Section 2 gives an overview of existing approaches. Section 3 covers the experimental setup and method for our solution. The collected data is presented, interpreted and discussed in Section 4, while conclusions are drawn in Section 5.

2. Related Work

2.1. Related Work in Haptics

Most of the related research in Haptics is conducted in a direction that involves a contact-based sensor array in the shape of a pen, which is the most convenient arrangement to use when manually guiding the sensor over the surface.

One of the earliest proposals [2] goes back to 2003: it presents a device named WHaT that employs a pair of 2-axis accelerometers and a piezoresistive force sensor chip. Preliminary results indicated a clear difference between different test materials. However, this device can only sense force along a single axis, and is therefore not capable of measuring Coulomb friction. Andrews and Lang [3] improved on the WHaT in 2007 using advanced filtering and calibration (at the expense of mobility), and achieved recognition of surface patterns as well as giving a realistic estimate to surface compliance coefficients.

In 2011, Kuchenberger et al. introduced a system named haptography that employs linear prediction of acceleration signals to present a new texture modeling and synthesis method [4]. Their handheld tool measures translation, rotation, force, torque, and high frequency accelerations in 3 dimensions, with a high emphasis placed upon the latter. The system is capable of generating texture models that, when fed to a Phantom Omni system, produce textures that were similar to the original surface according to test subjects. Culbertson et al. proposed an improved version of this system in 2014 [5].

The hardware in this system is capable of recording position, orientation, force, and high frequency acceleration, for which it only relies on a pair of 2-axis accelerometers. A Wacom tablet and a Haptuator are used for realizing the generated virtual texture models. On the software side, the input signal is segmented and presented as a piecewise autoregressive process, allowing the generation of a set of localized texture models that make up a realistic virtual texture model. A detailed usability study showed that test subjects found the virtual recreation to be highly similar to the original in terms of roughness, but not with respect to hardness or slipperiness.

The contact-based haptic surface property acquisition systems discussed above have undergone a great amount of improvement over the last decade. However, their fundamental limitations of 1-dimensional texture mapping and short measurement range (arising from the need for contact) still keep them from being considered as an option for a haptic vision system. This is why we are applying ultrasound imaging to this problem, as it does not require physical contact in all cases. Hence, ultrasound has the potential to be used as a basis of a nondestructive and noncontact evaluation method.

2.2. Related Work in Ultrasound Imaging

Ultrasound imaging rose to prominence through its application in medicine. The basis of this imaging method is exposing the object under investigation to ultrasound beams and capturing the corresponding reflections. An early study [6] presented a 3D ultrasound system that combines 2D B-mode scans into a 3D model. The system is characterized by relatively good accuracy and precision (97.4% and 97.5%, respectively), as well as acceptable intra / interobserver variability (5.1% / 11.4%). Many other medical publications related to ultrasonic imaging followed suit in the following years, including use cases for prostate cancer detection [7], kidney stone detection [8], chronic kidney disease classification [9], and thyroid cancer detection [10].

One of the earliest manifestations of the idea to utilize ultrasound for measuring surface properties of materials comes from Murayama et al., who presented an ultrasound based remote sensing system in 2005 [11]. This system consisted of a piezoelectric transducer and a feedback circuit. The sought after properties were derived from the phase-shifted values of the reflected signal. Ultrasonic reflection is analogous to that of light, and Snell's law can be applied to it, albeit with acoustic impedance as the characteristic property of the materials. This technology is capable of differentiating between various metals (aluminum, copper, iron and even silicon), although the semi-logarithmic correspondence found by the researchers might not be generalizable.

In 2007, Park et al. proposed using a high frame rate (85 fps) ultrasound imaging system to determine the reflection distribution of a sample with known material composition [12]. They created focal points slightly beneath the surface of the object and captured the reflections. The strain image and elasticity map they acquired corresponded reasonably well to traditional imaging results, as well as their own analytical model, although the signal-to-noise ratio was slightly lower than for conventional imaging.

As ultrasound for elasticity imaging gained prominence, it was soon adopted by the medical field, to which [13] is an example, where comparable measurement results were achieved compared to MRI and CT. Yordanov et al. [14] also proposed a system for noncontact ultrasonic measurements of both fluids and solids that can be embedded into automated manufacturing systems. The system contained the conventional elements (piezoelectric transducers, amplification circuit, and software-based processing) and managed to differentiate between fluids with different levels of alcohol content, as well as between iron and cast iron.

In 2009, Urban et al. published a computational model as well as an experimental approach to analyze the errors in the measurements of shear wave velocity and material properties, when shear wave dispersion ultrasound vibrometry (SDUV) is performed [15]. They found that these values are most overestimated for materials with low shear viscosity and high viscosity. The opposite takes place with materials possessing high shear elasticity and low (but nonzero) viscosity values. Amador et al. then performed SDUV on swine kidney in 2011 [16], examining a set of eight female kidneys in vitro. They measured the elasticity and viscosity of the renal cortex and found that the former does not change significantly over time, though the latter does. They also concluded that the renal cortex is anisotropic.

Using ultrasound is now a widely used form of imaging in the medical field, thanks to its relatively low cost and high speed (real time imaging), as well as its noncontact nature. These properties make it an ideal candidate to be used in a nondestructive evaluation tool such as the one proposed in this article. It is a novel approach meant to replace current contact-based haptic surface property measurement tools. This paper is dedicated to exploring the feasibility of this concept by carrying out an experiment to classify different materials using their measured reflection of ultrasound.

3. Experimental Setup and Method

The experimental setup consists of the following modules: the experimental module and the control and acquisition module. The experimental module comprises two 44kHz piezoelectric ultrasonic transducers mounted inside the end of a 69.3cm long PVC tube as well as a sample placed at the other end of the same tube. The transducers are placed in a way that one of them occupies the center of the cross-sectional area of the end of the tube, while the other is next to it, being slightly off-center. A photograph of the experimental module can be seen in Figure 1. The samples (visible in Figure 2) were all Sorbothane samples, with hardness values of 30, 40, 50, 60, and 70 on the Shore 00 scale.

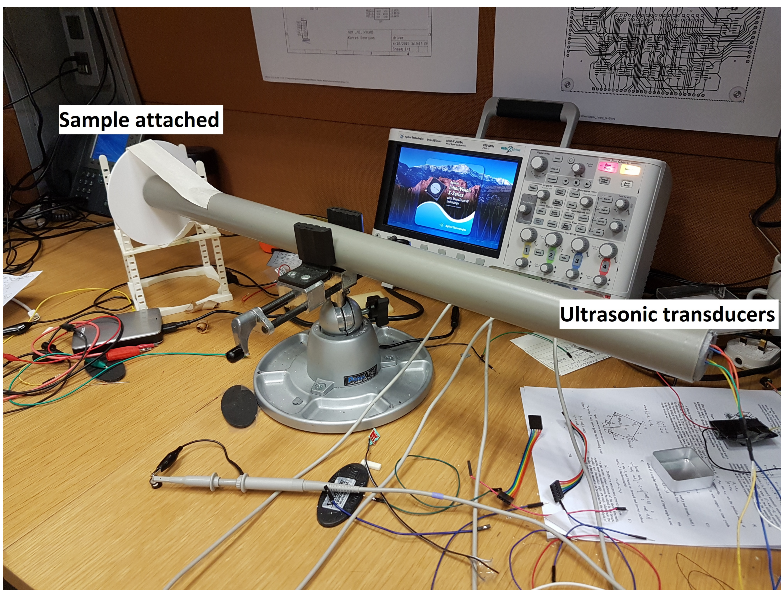

Figure 1: The experimental module. In this image, a piece of acrylic glass is examined, with the Sorbothane samples located near the tube.

The control and acquisition module consists of a control circuit for the transducer, an Agilent InfiniiVision MSO-X 2024A Mixed Signal Oscilloscope with 1GSa/s maximum sampling rate, and a computer that controls the operation of the above two elements in the MATLAB environment.

The experimental procedure is as follows. Over a single experimental session, each of the five Sorbothane samples are placed at the end of the tube that is opposite to the ultrasonic transducers at separate times. During each

(a) Labeled side of the Sorbothane samples.

(b) Backside of the Sorbothane samples. This side was subjected to the ultrasonic radiation.

Figure 2: The experimental samples.

of these times, a series of a thousand pulses are emitted from the centered transducer, with a 100ms delay between them. The second transducer acts as a receiver and converts the reflected ultrasound waves to an analog signal that is acquired by the oscilloscope. Figure 3 shows the relationship between the pulse and its echo waveform (in this instance, a Sorbothane 30 sample was placed at the examination end of the PVC tube). Once the time domain data for a sample is collected, it is replaced by the next sample and the thousand pulse-echo data is recorded again for that sample.

When all the time domain data for each of the samples is acquired, the data acquisition session is over and the data is processed. First, a Fourier-transform is performed on the time-domain data. Since each pulse-echo period is recorded for 100ms, the frequency-domain resolution of the transformed data will be 10Hz. Given that the time-domain data is sampled every 10 microseconds, the single-sided Fourier-data ranges from 0 to 100kHz. In order to increase the relevance of smaller frequency bands, the DC compo-

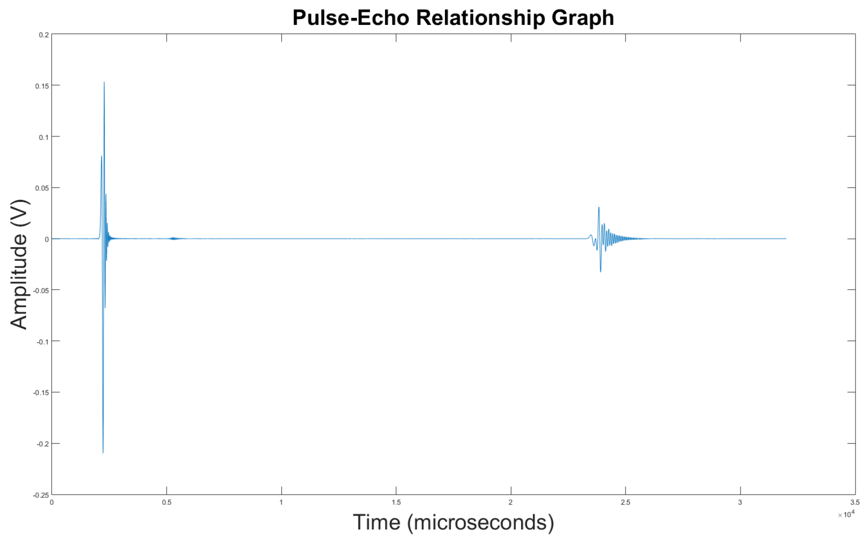

Figure 3: The Pulse-Echo waveform graph with a Sorbothane 30 sample under examination.

nents are removed from the spectra. Figure 4 shows the frequency-domain data for a thousand pulse-echo waveforms performed on a Sorbothane 50 sample, averaged.

As a final step of evaluating data, the previously described frequency-domain data entries are fed into a Support Vector Machine that performs fivefold cross-validation on them. This is accomplished in the following fashion. Firstly, the order of the 1000 pulse-echo frequency entries is randomized. Then, they are separated into five subgroups of 200 entries for each sample. Finally, five groups are formed, each of which contains a subgroup of each of the five samples (Sorbothane 30/40/50/60/70). During the fivefold cross-validation process, four of these large groups (4000 frequency domain entries in total) are used to train a linear Support Vector Machine (SVM) classifier and the fifth group is used to validate the classifier. Depending on which large group to use as the validating data set (with the rest being used for training), five scenarios arise, each with a possibly somewhat different classifier. This is what the "fivefold" term refers to. Figure 5 shows a schematic for how the cross-validation is performed.

Training a five-way classifier is accomplished by creating a series of one-vs-all classifications. This means first training a classifier where the Sorbothane 30 entries are labeled as '30', while each other entry is labeled as 'not 30' (or

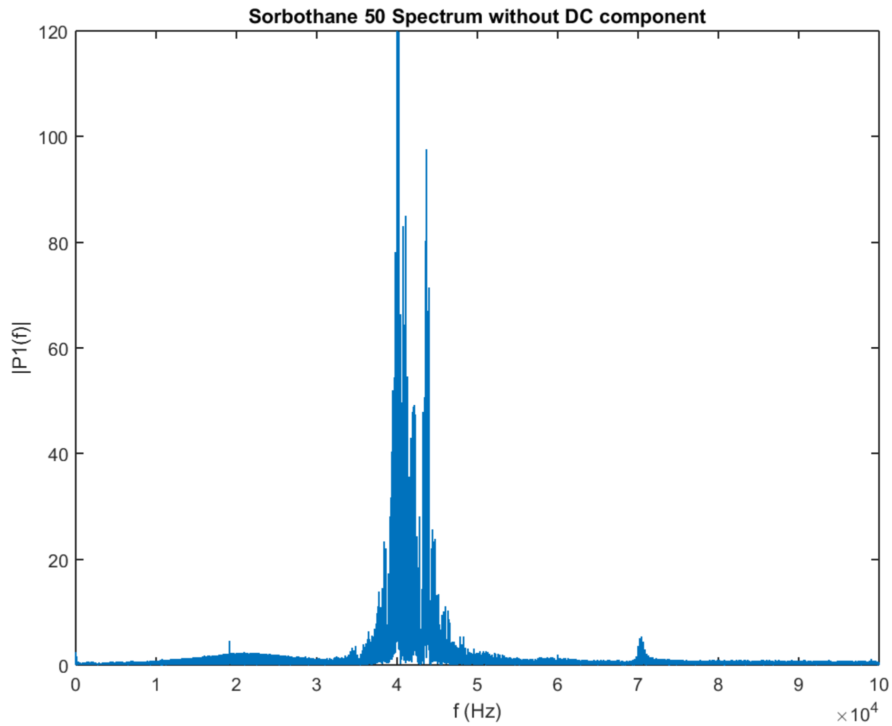

Figure 4: Fourier transform of the Pulse-Echo waveform graph of a Sorbothane 50 sample with the DC component removed.

'-30' for simplicity). Next, another classifier is trained where the Sorbothane 40 entries gain a label of '40' while all other entries have a label of 'not 40'. This is repeated for the Sorbothane 50/60/70 classifiers as well. When it comes to validation, the classification scores, as well as the classified labels, are recorded for every validation entry in each step. Out of all these labels, the one with the highest classification score is selected as the actual label. This is the answer that we ideally desire to get as to which sample the entry could come from.

Two extensions of the above mentioned fivefold cross-validation process were also investigated. In the first extension, data from multiple data collecting sessions (with the exact same samples) were combined into both the training and the validating data set, conforming to the fivefold cross-validation protocol. In the second extension, data from five different sessions were combined to form the training data set, while data from a sixth session was used

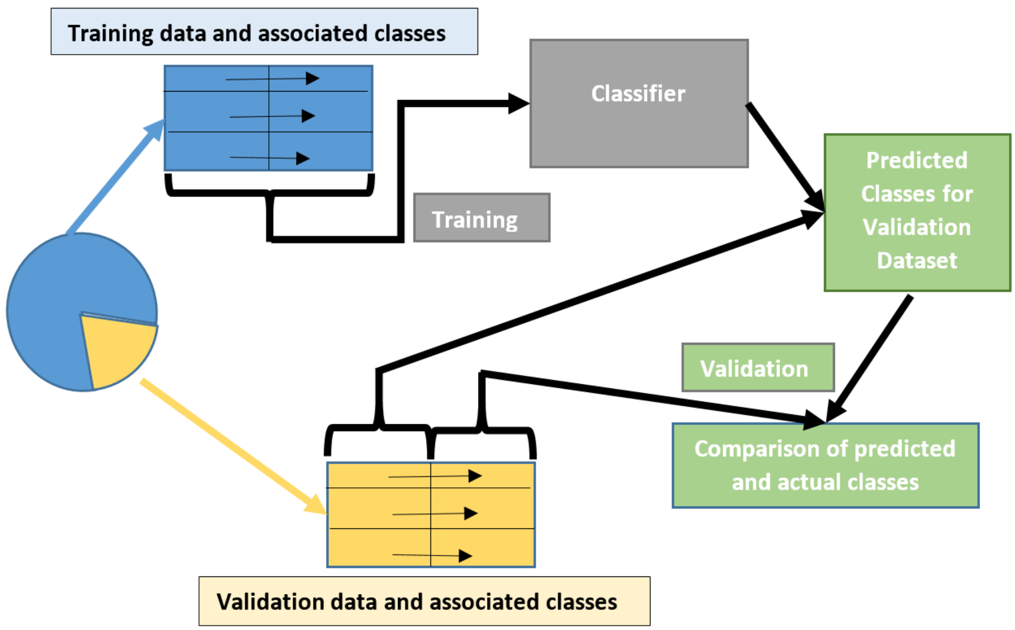

Figure 5: Schematic diagram representing cross-validation.

for validation.

4. Results

Figure 6 shows fivefold classification performed on one of the data sets. The training data set consisted of 800 samples of Sorbothane 30/40/50/60/70 each, while the validation data set contained 200 entries of each category. The order of the validation entries does not matter, as they are categorized individually based on training data. Therefore, for the best visibility, the validation entries were ordered in an increasing fashion with respect to their durometer values (i.e. 200 entries of Sorbothane 30 followed by 200 entries of Sorbothane 40, and so on). This means that the ideal result is a staircase function with the first 200 entries corresponding to the label '30', the next 200 corresponding to '40' and so on. Actually, this is exactly what is present in Figure 6. It can be seen from the image that the Sorbothane 50 samples were classified with the highest confidence, as their associated classification score is significantly higher than those corresponding to the other samples.

Figure 7 shows an extension to the previous case, where both the training and the validation data were acquired over 3 different data acquisition ses-

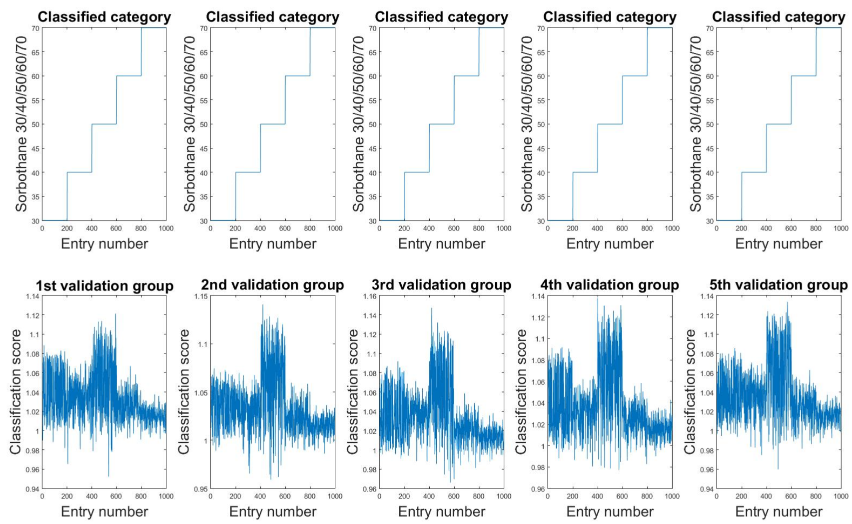

Figure 6: Fivefold cross-validation results for session Y. Top: classified categories. Bottom: classification scores.

sions (corresponding to batches H, I and J). We can, once again, see that the classification is carried out as desired with correct results. A difference compared to the previous case is that the classification scores are less significantly different with respect to the different Sorbothane samples.

Figures 8 and 9 show the results for the second extension, where the training and validation entries were acquired over different data acquisition sessions. For this experiment, four out of five of the Sorbothane test samples were replaced with vastly different materials: cardboard, wood, acrylic glass, and steel. This change was meant to ease the classification task, as these materials have physical properties that are much more different from each other than the amount of diversity between the Sorbothane samples. Figure 8 shows perfect cross-validation results for sessions A, B, C, D and E combined. However, when these data sets are used entirely as training entries and another (F) data set is used for validation, as Figure 9 shows, the classified results fail to correspond to the actual situation.

As we have seen from the above results, the fivefold cross-validation method provides a remarkable accuracy when it comes to classifying the Sorbothane materials. Since in each case, no entries were used for both training and validation, it means that the results are indicative of an underlying difference in the way the different samples reflect ultrasound.

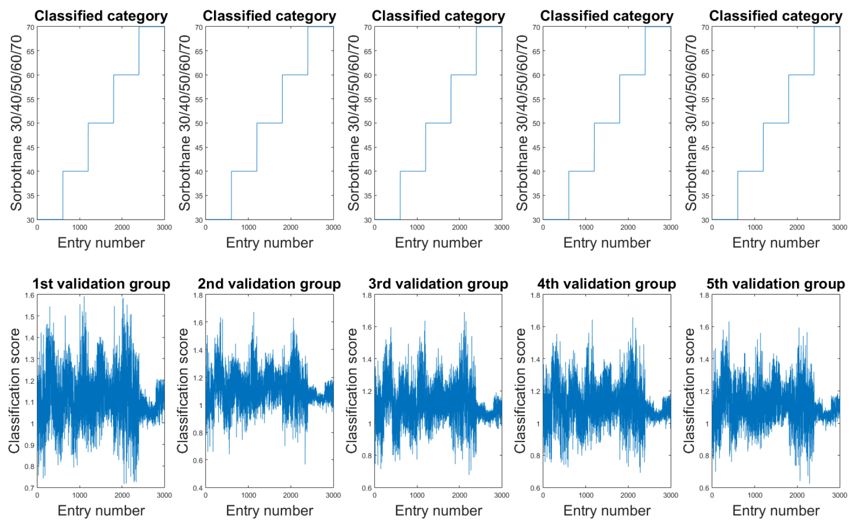

Figure 7: Combined fivefold cross-validation results for sessions H, I and J. Top: classified categories. Bottom: classification scores.

Although concerns of overfitting are always worthy of consideration when it comes to machine learning, the multi-session cross-validation results (such as Figure 7 or Figure 8 show that even when the training sets encompass different sessions (negating conditions that could change from session to session, such as temperature while potentially keeping features that are related to the samples themselves), the trained models provide accurate classification.

The negative side of the results is visible in Figure 9, where classification accuracy is very low, despite the fundamentally different materials used. This is the scenario where the classifiers were trained on data sets A, B, C, D, and E, and validated against data set F. What this means is that even though the first five data set provided a good classification base for entries that were acquired over the same sessions (even though different entries were used for training and validation), the classification is still not accurate enough to work well for a different session. The reason for this could be that environmental effects that vary from session to session are still not negligible even when the training is done on a data set containing five different sessions. One way to overcome this is to include many more sessions in the training process. This could come at a cost of additional training time, although the amount of training entries per session could be reduced in turn, in order to achieve a good trade-off between training time and training accuracy.

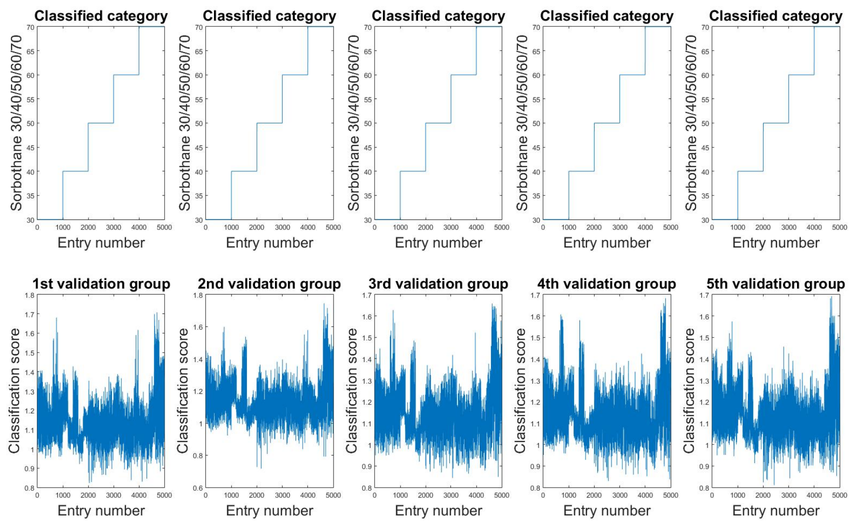

Figure 8: Cross-validation results on sessions A, B, C, D, and E.

5. Conclusion

This paper presented a material classification approach based on machine learning to distinguishing between materials by examining the way they reflect an ultrasonic pulse. Results show that classification is possible with high accuracy when training data is available for similar physical conditions. However the range of such conditions might warrant a relatively large training data set to cover most of the possible data acquisition scenarios.

[1] D. Whitworth, Kinect gets uk release date, BBC Newsbeat (http://www.bbc.co.uk/newsbeat/article/10996389/kinect-gets-uk-release-date) Retrieved: November 15, 2015.

[2] D. K. Pai, P. Rizun, The what: a wireless haptic texture sensor, 11th Symposium on Haptics (2003) 3–9.

[3] S. Andrews, J. Lang, Interactive scanning of haptic textures and surface compliance, Sixth International Conference on 3D Digital Imaging and Modeling (3DIM) (2007) 99–106.

[4] K. J. Kuchenbecker, J. Romano, W. McMahan, Haptography: Capturing and recreating the rich feel of real surfaces, Springer Tracts in Advanced Robotics 70 (2011) 245–260.

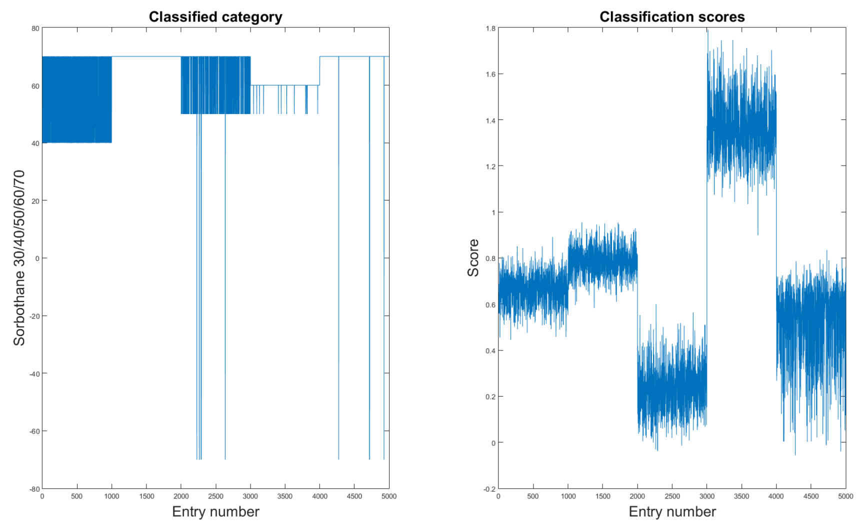

Figure 9: Inter-dataset validation: Trained on A, B, C, D, and E, validated on F.

[5] H. Culbertson, J. Urwin, K. J. Kuchenbecker, Modeling and rendering realistic textures from unconstrained tool-surface interactions, IEEE Transactions on Haptics 7 (2014) 381–393.

[6] A. Fenster, S. Tong, Three-dimensional ultrasound imaging system for prostate cancer diagnosis and treatment, IEEE Transactions on Instrumentation and Measurement 47 (1998) 1439–1447.

[7] M. P. J. Kuenen, M. Mischi, H. Wijkstra, Coherence-based contrast ultrasound diffusion imaging for prostate cancer detection, IEEE International Ultrasonics Symposium Proceedings (2010) 1936–1939.

[8] K. Viswanath, R. Gunasundari, S. A. Hussain, Analysis of kidney stone detection by reaction diffusion level set segmentation and xilinx system generator, Proceedings of the International Conference on Advanced Research in Computer Science Engineering and Technology (ICARSET 2015) 36 (2015) 1–10.

[9] J. W. Hsieh, C. H. Lee, Stage classification in chronic kidney disease by ultrasound image, Proceedings of the 29th International Conference on Image and Vision Computing New Zealand (IVCNZ '14) (2014) 271–276.

[10] M. L. Oelze, Quantitative ultrasound techniques and improvements to diagnostic ultrasound imaging, IEEE International Ultrasonics Symposium Proceedings (2012) 232–239.

[11] Y. Murayama, C. E. Constantinou, S. Omata, Remote sensing of mechanical properties of materials using a novel ultrasound transducer and signal processing, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 52 (2005) 439–444.

[12] S. Park, S. R. Aglyamov, S. Y. Emelianov, Elasticity imaging using conventional and high-frame rate ultrasound imaging: Experimental study, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 54 (2007) 2246–2256.

[13] H. Xie, A. T. Fernandez, Real-time ultrasound elasticity imaging for liver rf ablation assessment: Preliminary ex vivo and in vivo animal studies, IEEE International Ultrasonics Symposium Proceedings (2009) 139–142.

[14] S. Yordanov, R. Ilarionov, System for non-contact ultrasonic study of mediums and materials intended for embedding into automated manufacturing systems, International Conference on Computer Systems and Technologies (CompSysTech '10) (2010) 353–358.

[15] M. W. Urban, S. Chen, J. F. Greenleaf, Error in estimates of tissue material properties from shear wave dispersion ultrasound vibrometry, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 56 (2009) 748–758.

[16] C. Amador, M. W. Urban, Shearwave dispersion ultrasound vibrometry (sduv) on swine kidney, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 58 (2011) 2608–2619.