

Algorithms for Tactile Rendering in Compliant Environments

Camilla K. L. Lau
Harvard University
lau@fas.harvard.edu

Christopher R. Wagner
Harvard University
cwagner@fas.harvard.edu

Robert D. Howe
Harvard University
howe@deas.harvard.edu

Abstract

Relaying spatially distributed forces promises to enhance the performance of telemanipulation systems. However, the proper way to render tactile information from the sensor to the display is not clear for current displays, where the user applies a constant contact force. We present a simple approach to rendering tactile information to improve performance in a lump localization task in a compliant environment. The algorithms presented subtract uninformative background information from the tactile signal. We observed that subtracting a fixed background pressure frame reduced lump localization error by up to 20% while decreasing the time required to find the lump by up to 44%. Subtracting a background frame that depends on applied force further reduced lump localization error by another 17%.

1. Introduction

Through haptic feedback, telemanipulation systems attempt to simulate the sensation of direct contact with a remote environment. Most current systems, however, only transmit force information based on a single point of contact. These telemanipulators can be enhanced with the transmission of spatially distributed force information, sometimes denoted “teletaction” [1]. Spatial pressure or shape profiles at a remote location are sensed, and then a corresponding pressure or shape distribution is recreated against the operator’s fingerpad. The presence of spatially distributed forces has been shown to significantly enhance spatial acuity, orientation detection, and performance on lump detection tasks using palpation [2]. These tasks are important in remote medicine and minimally invasive surgery where the addition of tactile feedback can improve manipulation as well as diagnosis [3, 4].

While much work has investigated rendering [5] and transmission [6] techniques for point based forces, little work has examined the proper way to render spatially distributed forces. This applies to both virtual tactile environments as well as teletaction systems. This observation is somewhat surprising given the number of tactile displays and tactile sensor devices that have been developed [7-12]. Fundamental work has been presented by Moy et al., who developed techniques for the optimal

transmission of tactile information based upon linear mechanical models of the tactile sensor and the fingerpad [13]. However, this work targets a tactile display that is attached to the finger and can apply stimulus that varies continuously from zero. This does not represent current tactile display technology, which is typically fixed in position and requires a significant user-applied force to maintain contact between the finger and display, across all levels of tactile stimulus presented by the display.

In this paper, we investigate algorithms for use with teletaction systems for finding stiff features within compliant environments. These algorithms are applicable to medical and surgical tasks such as breast tumor localization [14] and liver palpation [15]. Direct transmission of information from a tactile sensor to a tactile display can be problematic in compliant environments because the tactile information due to the contact with the enveloping medium can obscure the relevant tactile information of an embedded feature. Our hypothesis was that if this constant background stiffness were subtracted from the tactile signal, the signal-to-noise ratio would improve, and performance would be enhanced in a lump finding task. However, background pressure amplitude and distribution varies with applied pressure on the tactile sensor, so we investigated both the case where a fixed pressure frame was subtracted and where a pressure frame that varied linearly with applied pressure was subtracted. These algorithms, along with the direct transmission case, were evaluated in a lump finding task across a range of compliant environment thicknesses. The goal was to find salient features of touch information used in teletaction so that these features could be the focus of future algorithm and teletaction system development.

2. Methods and Materials

2.1. Teletaction system

The teletaction system used for the experiment consists of a tactile sensor that measures a two dimensional pressure profile, a tactile display that recreates small scale shape profiles on the fingerpad, and the necessary signal processing algorithms that process the information from the sensor to the display (Figure 1). The tactile sensor (Pressure Profile Systems, Inc., Los Angeles, CA) measures pressure across a 16 x 16 square array of

pressure sensels (sensor elements) spaced 2 mm apart with a resolution of about 0.4 kPa (Figure 2a). We minimize the effect of the curvature and achieve a one to one mapping with the tactile display by only using the center 6 x 6 sensels. When in contact with an object, the sensor can obtain a pressure profile reflective of the object's tactile qualities such as shape or stiffness at a bandwidth of up to 10 Hz.




Figure 1. Comparison between the flow of tactile information in direct contact and in the teletaction system

The tactile sensor is mounted in series with a force sensor (Nano43; force resolution: 0.031 N; ATI Industrial Automation, Inc; Apex, NC) to acquire the overall exploration force applied to the sensor through a handle (Figure 2b). In addition, a magnetic position tracker with a positional resolution of 0.51 mm (MiniBird, Ascension Technology Corporation; Milton, VT) is used to determine the position of the tactile sensor using position and orientation information. The magnetic tracker is mounted to the opposite end of the handle to minimize the interference of the sensor's steel base.




Figure 2. (a) Schematic of tactile sensor array (b) Tactile sensor (i), force sensor (ii), handle (iii), and position sensor (iv)

The tactile display is an array of mechanical pins actuated by RC servomotors [11] (Figure 3). The pins have diameters of 1 mm and an interpin spacing of 2 mm. They have a maximum displacement of 2 mm and a resolution of 0.1 mm. A 2 mm thick piece of silicone rubber (HSII RTV Base and Colored Catalyst, Dow Corning) was placed on the pins of the tactile display as a spatial low-pass filter [16]. Although the tactile display can run up to 25 Hz for small pin movements, the tactile sensor limited the bandwidth of the teletaction system to 10 Hz.




Figure 3. (a) Tactile display with rubber low-pass filter (b) Close-up of tactile display pins

2.2. Algorithms and signal processing

There are two filters applied to the tactile sensor data before the pressure profile is recreated on the tactile display. The first filter is one of three rendering algorithms that attempt to correct for a constant background stiffness. The second filter is a weighted average that reduces high spatial frequency noise in the sensor data. For a given sensel position, the corresponding sensel was weighted at 50%, while the eight surrounding sensels together were weighted at 50%.

Three rendering algorithms were used in the experiment. The first algorithm set the tactile display pin heights proportional to the pressure data from the tactile sensor. For the second algorithm, a fixed pressure frame is subtracted from each pressure frame received from the tactile sensor. Similar to the second algorithm, the third also subtracts a pressure frame from each frame collected, but the frame that is subtracted depends linearly on the current pressure the user applies. Typical pressure frames along the centerline for a range of applied pressures are shown in Figure 4. The full signal processing (algorithms and averaging) is expressed in Table 1, where D is the 6 x 6 array of tactile display pin heights, P is the 6 x 6 array of pressure data from the tactile sensor, \tilde{P}_2 is a fixed 6 x 6 array of background pressures, \tilde{P}_3 is a 6 x 6 array of background pressures that varies with f , the exploration

force in the dimension normal to the plane, $k = 1.4$ is a constant that scales sensor data to display pin height, $*$ denotes convolution of the pressure with A , a noise filtering averaging kernel

$$A = \begin{bmatrix} 0.0625 & 0.0625 & 0.0625 \\ 0.0625 & 0.5 & 0.0625 \\ 0.0625 & 0.0625 & 0.0625 \end{bmatrix}$$

Each element is then limited to values between 0 and 2 mm to reflect the capability of the tactile display.

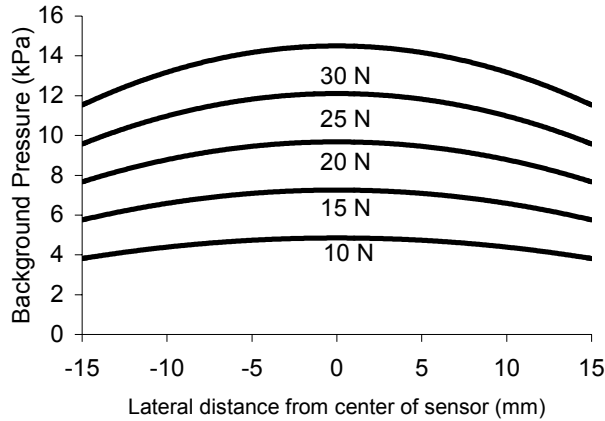


Figure 4. Background pressures for different applied forces

Table 1. Rendering algorithms

Alg.	Description	Full Expression
1	Pressure information directly transmitted	$D_{i,j} = k[P * A]_{i,j}; i, j = 1, \dots, 6$
2	Fixed frame subtracted	$D_{i,j} = k[(P - \tilde{P}_2) * A]_{i,j}; i, j = 1, \dots, 6$
3	A frame subtracted dependent on force	$D_{i,j} = k[(P - \tilde{P}_3 f) * A]_{i,j}; i, j = 1, \dots, 6$

2.3. Embedded lump models

The three algorithms were compared by having subjects locate rigid lumps embedded in constant stiffness elastic models. We constructed twelve models with a 1.90 cm rigid acrylic ball on the bottom of a 30.5 x 7.8 cm x 5.1 cm container filled with silicone rubber (GE Silicones RTV6166) with a Young's modulus of approximately 15 kPa (Figure 5). The silicone thicknesses were chosen to be 1.50, 1.75, and 2.00 times the diameter of the ball, or 2.86 cm, 3.33 cm, and 3.81 cm. The thinnest model

allowed users, through the teletaction system, to readily detect a lump with a reasonable exploration force (~20 N) while the thickest model required users to exert a substantially higher force to find the lump (~30 N). For each thickness of silicone, the ball was glued to the bottom of the container in one of four locations, spaced approximately five centimeters apart to minimize memorization of the location of the lump. The width of the container was chosen so that the tactile sensor would fit well with less than a centimeter on either side. In this way, subjects would only need to search in one dimension. A thin layer of latex was laid across the top of each model in order to protect the silicone from damage, and hand lotion was used to lubricate the surface so that the sensor could slide across the model with low friction.




Figure 5. Cross-section of silicone models

2.4. Experimental design

This experiment used a two-factor, within-subject repeated measures design with algorithms (three levels) and model thickness (three levels). Each combination of algorithm and model thickness was presented ten times in random order to each subject, for a total of 90 trials per subject. The order in which the various combinations of algorithm and model thickness were presented was counterbalanced across trials and across subjects. Fifteen subjects, ranging from ages 20 to 33, volunteered for the experiment for monetary compensation. All subjects described themselves as right-handed and reported no hand injury to either hand.

2.5. Procedure

Subjects were told to locate hard lumps in soft tissue models by exploring the models using the tactile sensor with their right hand and feeling the tactile display with their index finger of their left hand (Figure 6). For each trial, the subject began with the sensor touching the left side of the model and then scanned the sensor back and forth across the one-dimensional model. The subjects' primary goal was to accurately center the tactile sensor directly above the hard lump. Given that they could achieve the primary goal, the secondary goal was to complete this task as fast as possible. The trial was stopped when the subject verbally announced that he or she had found the lump.




Figure 6. Experimental setup for teletaction system

Although subjects had no time restraints, they were informed that each trial could take up to a minute, although the average trial length would be less than thirty seconds. They were given a short break every nine trials, and total experiment length was about an hour. Subjects were blindfolded so they were unable to see the ball in the model, and their exploration technique would not be affected by observing the thickness of the silicone model. To mask audio cues, the subjects wore earplugs and headphones that played noise in the frequency range of sounds made by the tactile display.

Before the experiment, subjects were given a training session in which they were taught how to use both devices and in which they could gain experience in practice trials. Subjects were always trained with a model of medium thickness, but the algorithm on which they were trained was chosen randomly. We also trained the subjects to recognize the largest force needed in order to avoid potentially damaging higher forces to the force sensor.

2.6. Analysis

We defined outlier trials to be those in which the final lateral position error was greater than 2.5 cm, approximately half the distance between the placement of the balls. These trials were discarded from further analysis and noted as trials where the subject did not find the lump.

The static background pressure frame P used for Algorithm 2 was determined specific to each model thickness prior to the experiment. Pilot studies showed that acquiring this typical pressure frame using a force of 20 N allows subjects to gain useful tactile information at typical forces of about 25-30 N. The background frame used for Algorithm 3 was determined by taking five sample frames at various forces up to 40 N. Previous tests showed us that each sensel's signal is approximately a linear function of the applied force. Using five data points of force applied and corresponding background frames, we calculated a linear fit using the method of least squares for each sensel.

3. Results

Increasing algorithm index (1: no subtraction, 2: fixed pressure frame subtraction, 3: force-dependent pressure frame subtraction) resulted in a decrease in lateral position. Within each algorithm, increased model thickness increased lateral error. A two-factor, within-subject ANOVA was performed on the lateral position errors. The two factors were algorithm (three levels) and model thickness (three levels). Both algorithm and thickness were found to be significant with respect to the primary performance metric of lateral error; ($F(2,248) = 27.0, p < 0.001$ and $F(2,248) = 17.7, p < 0.001$, respectively) (Figure 7a). Pairwise comparisons show that all three algorithms are significantly distinct from one another with respect to lateral error. However, the thinner models (2.86 and 3.33 cm) are not significantly distinct ($p > 0.10$).





Figure 7. (a) Average lateral error and (b) exploration time across all subjects' trials. Error bars show standard error.

Increasing algorithm number and decreasing model thickness each served to significantly decrease the time required to find the lump ($F(2,248) = 31.6, p < 0.001$ and $F(2,248) = 50.6, p < 0.001$, respectively) (Figure 7b). Pairwise comparisons show that Algorithm 2 and 3 are not significantly distinct with respect to time per trial ($p > 0.9$). However, all levels of thickness are distinct with respect to time per trial. The interaction effect between algorithm and thickness did not reach significance for either the lateral position error nor the time per trial ($F(4,496) = 1.43, p > 0.2$ and $F(4,496) = 0.738, p > 0.50$).





Figure 8. Outliers for all subjects, categorized by algorithm and thickness. Error bars show standard error.

Subjects were unable to find the lump in 4.1% of the trials, where increasing algorithm index resulted in a significant decrease in the number of times the subject missed the lump ($F(2,28)=5.84, p < 0.01$). Figure 8 shows the distribution of outliers for all subjects, categorized by algorithm and model thickness. Varying the thickness had no significant effect ($F(2,28)=2.51, p > 0.10$).


Average forces were calculated for each trial and averaged across subjects (Figure 9a). Subjects, when using Algorithm 3, applied significantly more force than when using Algorithms 1 or 2 and subjects applied significantly less force when using the thinnest model ($F(2,266) = 17.5, p < 0.001$ and $F(2,266) = 10.5, p < 0.001$). We also examined forces one second before the subject found the lump. Again, subjects, when using Algorithm 3, applied significantly more force than when using Algorithms 1 or 2 and increased model thickness results in increased applied force ($F(2,266) = 18.4, p < 0.001$ and $F(2,266) = 37.1, p < 0.001$) (Figure 9b).

For each combination of algorithm and thickness, the average absolute exploration speed was determined by dividing the distance moved between consecutive data points by the time elapsed and averaging across each trial (Figure 10). While a two-factor ANOVA shows that both algorithm and model thickness are statistically significant

($F(2,266) = 5.66, p < 0.005$ and $F(2,266) = 3.72, p < 0.05$), pairwise comparisons show that Algorithm 1 and 2 are the only significantly distinct groups with respect to exploration speed ($p < 0.01$), and that none of the thickness levels are statistically distinct.



(a)



(b)

Figure 9. (a) Average exploration force in a trial. (b) Average force applied one second before the end of the trial. Error bars show standard error.




Figure 10. Average absolute speed across all trials. Error bars show standard error.

4. Discussion

Our goal was to examine the performance benefits of using simple rendering algorithms when using a tactile system to locate rigid features embedded in compliant environments. We found that the subtraction of a fixed background frame (Algorithm 2) significantly improves the ability to localize the lump and reduces the time needed to find the lump across all levels of thickness. Average lateral error decreased from 0.61 cm to 0.49 cm for the thinnest model (20% error reduction) and from 0.81 to 0.68 cm for the thickest (16% error reduction). Average time required to find the lump decreased from 10.0 s to 5.6 s for the thinnest model (44% reduction in time) and from 16.0 s to 10.4 s (35% reduction in time). The subtraction of a fixed background frame also significantly reduces the chance that a user will miss a lump entirely, reducing the average number of outliers per 10 trials from 0.6 to 0.2 for the thinnest model (67% reduction in misses) and from 1.1 to 0.5 for the thinnest model (55% reduction in misses).

Using a force dependent background subtraction algorithm (Algorithm 3) further improves lump localization accuracy while not significantly changing the time needed to find the lump. Average lateral error decreased another 0.1 cm for the thinnest model (a 36% total decrease in lateral error) and an additional 0.18 cm for the thickest model (a 37% total decrease in lateral error). The chance that a user will miss the lump entirely is similarly reduced, with the average number of outliers per ten trials at 0.0 for the thinnest model (a 100% total reduction in misses) and 0.1 for the thickest (an 80% total reduction in misses).

These results are interesting due to the simple nature of the algorithms. Even though background tactile information varies with applied pressure on the sensor, a significant performance benefit is still observed when subtracting off a *fixed* background frame. The benefit is also observed across all levels of thickness, showing that the performance increase is robust across environments. This benefit can be derived without the use of an extra force sensor. The full benefits of Algorithm 3 can also be derived without a force sensor if the active area of the tactile sensor covers the entire probe contact area; summing each individual pressure and multiplying by the area will give the contact force.

Both background subtraction algorithms can be considered model-based algorithms. They both are making an assumption on the nature of the compliant environment. Based upon that model, the *expected* uninformative background tactile information is determined and subtracted off. The algorithms are also task based. Removing the background information is useful when localizing features that are dissimilar to the

background. Note that this background should be homogenous and elastic for the algorithms to enhance performance. If the task were, however, determining the stiffness of the compliant environment, subtracting off the assumed background information would be counterproductive. Thus, increased knowledge about the task and environment can lead to more complex models, better algorithms and presumably increased performance in teletaction tasks.

From these results we hope to examine possible conclusions as to the nature of the benefit of the background subtraction algorithms. One comment is that the algorithms only pass the relevant information for a lump finding task. A naturalness of stimulation is still maintained due to the offset pressure encountered when using a table mounted (as opposed to finger mounted) tactile display. In other words, people are accustomed to feeling a background pressure along with a stimulus when in direct contact. The removal of the background pressure compensates for the offset pressure of the user pressing his or her finger into the tactile display. This offset pressure is always present in table mounted tactile displays; a user cannot apply zero pressure nor is the user's finger flat and flush with the display at zero pressure.

We hypothesize that similar performance benefits would still be observed when using background subtraction algorithms on a teletaction system with a finger mounted tactile display. Because these algorithms target a specific task, the user would again only feel the informative parts of the tactile signal. However, the system may trade off task performance for naturalness of stimulation. The user may be able to localize a lump just as well, but the user will not feel any background tactile information that would normally be present in direct contact.

The mechanism by which background subtraction increases localization ability is unclear. One hypothesis is that the background subtraction provides a threshold such that, at a given exploration force, only information relevant to features is transmitted to the tactile display. In this case, the center of the tactile display would show the lump with some noise, while the rest of the display would be zeroed. This is contrasted with the direct transmission case where the entire display is showing a noisy signal with the lump information only slightly more prevalent.

The question remains how the background subtracting algorithms result in a decrease in lump localization time. First, we do not observe a significant speed difference between Algorithms 2 and 3, showing users are finding the lumps sooner without needing to explore faster. Also, we would not expect a speed difference because subjects moved at a speed such that they densely sampled the environment considering the update rate of the teletaction system. An update rate of 10 Hz with an average

exploration speed of 5 cm/s would give the user tactile pressure frames spaced at 5 mm. Since each tactile frame is 1.2 cm wide, users receive information about each spatial location multiple times. Thus, subjects would not pass over lumps due to low system bandwidth, only due to lack of discernable information in the tactile frames. Note that this exploration speed is twice what has been observed with single finger, direct contact lump finding tasks [4]. However, this task is a one dimensional localization task, while Peine conducted a two dimensional task.

We can hypothesize how background subtraction helps reduce lump finding time by first examining lump finding strategies. Most subjects explored the one-dimensional environment in two phases: a lump detection phase, and a lump localization phase. During the detection phase, subjects were observed to steadily scan the sensor across the environment. Once the lump was found, subjects would attempt to center the sensor over the lump in the localization phase with repeated small side-to-side motions. If the subject did not detect the lump on the first pass, however, subjects would continue to scan the environment. We hypothesize that the background subtraction reduces exploration time – with the uninformative background information removed from the tactile signal, the subjects were more likely to detect the lumps with a smaller number of scan passes.

Examination of previous research on lump finding allows a comparison of exploration strategy based on environment properties. Earlier direct contact, single finger lump detection work showed subjects explored the environment using a series of discreet palpations where the subject would indent a finger into the compliant environment then circularly probe the environment with the sides of the finger while maintaining contact between the fingerpad and surface. In the present study, subjects explored using a continuous scan, where there is constant relative motion between the tactile sensor and the surface. The difference may be due to the differences in contact friction in the two cases. In the previous work, the surface was sticky, while in our experiment, the model was covered with a thin layer of latex along with a lubricant to reduce the friction between sensor and model.

These experiments have specific relevance to teletaction tasks in compliant environments. While the specific algorithms used here seem to be well suited to table based tactile displays, algorithms based on task and environment models should improve task performance for any type of tactile display. Another benefit of the described algorithms is that they are simple and cost effective to implement – a force sensor is not needed to derive most of the benefits of background subtraction.

5. References

- [1] R. S. Fearing, G. Moy, and E. Tan, "Some basic issues in teletaction," Proceedings of International Conference on Robotics and Automation, New York, NY, USA, 1997.
- [2] S. J. Lederman and R. L. Klatzky, "Sensing and displaying spatially distributed fingertip forces in haptic interfaces for teleoperator and virtual environment systems," *Presence*, vol. 8, pp. 86-103, 1999.
- [3] R. D. Howe, W. J. Peine, D. A. Kantarinis, and J. S. Son, "Remote palpation technology," *IEEE Engineering in Medicine and Biology Magazine*, vol. 14, pp. 318-23, 1995.
- [4] W. J. Peine, "Remote Palpation Instruments for Minimally Invasive Surgery," in *Division of Engineering and Applied Sciences*. Cambridge, MA: Harvard University, 1998, pp. 161.
- [5] C.-H. Ho, C. Basdogan, and M. A. Srinivasan, "Efficient Point-Based Rendering Techniques for Haptic Display of Virtual Objects," *Presence*, vol. 8, pp. 477-491, 1999.
- [6] T. B. Sheridan, *Teletrobotics, automation, and human supervisory control*. Cambridge, Mass.: MIT Press, 1992.
- [7] B. L. Gray and R. S. Fearing, "A surface micromachined microtactile sensor array," Proceedings of IEEE International Conference on Robotics and Automation, Minneapolis, MN, USA, 1996.
- [8] J. S. Son, E. A. Monteverde, and R. D. Howe, "A tactile sensor for localizing transient events in manipulation," Proceedings of the 1994 IEEE International Conference on Robotics and Automation, San Diego, CA, USA, 1994.
- [9] D. G. Caldwell, N. Tsagarakis, and C. Giesler, "An integrated tactile/shear feedback array for stimulation of finger mechanoreceptor," Proceedings of International Conference on Robotics and Automation, Piscataway, NJ, USA, 1999.
- [10] W.-T. Lo, Y. Shen, and Y.-H. Liu, "An integrated tactile feedback system for multifingered robot hands," Proceedings of RSJ/IEEE International Conference on Intelligent Robots and Systems, Maui, HI, USA, 2001.
- [11] C. R. Wagner, S. J. Lederman, and R. D. Howe, "A tactile shape display using RC servomotors," Proceedings 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2002, Orlando, FL, USA, 2002.
- [12] G. Moy, C. Wagner, and R. S. Fearing, "A compliant tactile display for teletaction," Proceedings 2000 ICRA. IEEE International Conference on Robotics and Automation, Piscataway, NJ, USA, 2000.
- [13] G. Moy, U. Singh, E. Tan, and R. S. Fearing, "Human Psychophysics for Teletaction System Design," *Haptics-e*, vol. 1, 2000.
- [14] H. S. Pennypacker and C. A. Pilgrim, "Achieving competence in clinical breast examination," *Nurse Pract Forum*, vol. 4, pp. 85-90, 1993.
- [15] C. D. Naylor, "Physical examination of the liver," *JAMA*, vol. 271, pp. 1859-65, 1994.
- [16] J. M. Lee, C. R. Wagner, S. J. Lederman, and R. D. Howe, "Spatial Low Pass Filters for Pin Actuated Tactile Displays," 11th Annual International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Los Angeles, CA, 2003.