

Semi-Supervised Haptic Material Recognition for Robots using Generative Adversarial Networks

Zackory Erickson, Sonia Chernova, and Charles C. Kemp

Institute for Robotics and Intelligent Machines

Georgia Institute of Technology, United States

zackory@gatech.edu, chernova@cc.gatech.edu, charlie.kemp@bme.gatech.edu

Abstract: Material recognition enables robots to incorporate knowledge of material properties into their interactions with everyday objects. For example, material recognition opens up opportunities for clearer communication with a robot, such as “bring me the metal coffee mug”, and recognizing plastic versus metal is crucial when using a microwave or oven. However, collecting labeled training data with a robot is often more difficult than unlabeled data. We present a semi-supervised learning approach for material recognition that uses generative adversarial networks (GANs) with haptic features such as force, temperature, and vibration. Our approach achieves state-of-the-art results and enables a robot to estimate the material class of household objects with $\sim 90\%$ accuracy when 92% of the training data are unlabeled. We explore how well this approach can recognize the material of new objects and we discuss challenges facing generalization. To motivate learning from unlabeled training data, we also compare results against several common supervised learning classifiers. In addition, we have released the dataset used for this work which consists of time-series haptic measurements from a robot that conducted thousands of interactions with 72 household objects.

Keywords: Generative Adversarial Networks, Material Recognition

1 Introduction

The ability to recognize materials plays an important role in how we interpret our surroundings and interact with the environment. For example, touch-based material recognition can enable us to infer objects when reaching into occluded spaces. If we reach into a purse and feel metal, it is likely either keys or coins. In addition, recognizing materials allows us to decide if an object can be heated safely in a microwave. Yet, material recognition of real-world objects remains challenging for robots.

In comparison to material recognition, the field of computer vision has seen substantial improvements recently in part due to the availability of massive image and video datasets. Given this, one approach to learn material recognition is to collect a large dataset of labeled training data that covers most of the variation a robot will encounter in the real world. However, robots are bound to interact with new objects that differ significantly from their original training set. Oftentimes, there is also a high cost associated with collecting labeled data on a robot.

In this work, we present a semi-supervised learning approach with generative adversarial networks (GANs) that enables a robot to learn from unlabeled tactile sensory data from interactions with everyday objects¹. By leveraging unlabeled sensor data that are more abundant in unstructured environments, we mitigate the need for massive labeled training sets. For example, our robot can classify the materials of household objects with $\sim 90\%$ accuracy, when 92% of the training data are unlabeled. This generative model can recognize an object’s material with 90% accuracy given just half a second of contact. We also explore how well this approach can recognize the materials of new objects. We note that further work is necessary, and we discuss some current challenges facing generalization. Finally, we are releasing a dataset² of time-series tactile measurements from a robot that conducted thousands of interactions with 72 household objects, as seen in Figure 1.

¹All code can be found at: <https://github.com/healthcare-robotics/mr-gan>

²Dataset can be downloaded at: <http://healthcare-robotics.com/mr-gan>



Figure 1: A PR2 robot along with the 72 objects that it interacted with.

2 Related Work

2.1 Haptic Perception

Bell et al. [1] introduced the Materials in Context Database (MINC) for material classification of everyday images. However, there are many scenarios in which it is challenging to visually determine an object’s material. For example, objects can be colored to alter their material appearance, like non-metal objects with metallic paint, or an object could be visually occluded or surrounded by clutter.

Bhattacharjee et al. [2] presented how a 1-DoF robot can perform supervised material recognition using active temperature sensing over a short duration of contact. Their robot distinguished between 11 materials by interacting with flat blocks of each raw material. Using an SVM on 500 labeled interactions per material, they achieved a material classification accuracy of 84% with 0.5 seconds of contact and 98% with 1.5 seconds of contact. However, note that this accuracy may not transfer to real-world objects with curved or uneven surfaces [3]. In our work, we also used active temperature sensors, and we compare GAN performance to more classical approaches such as SVMs.

Kerr et al. [4] also explored thermal sensing for material classification. They used a BioTac sensor to interact with 6 materials and collected 15 trials per material. Using a neural network, they achieved a classification accuracy of 73% given 20 seconds of contact per trial. Decherchi et al. [5] used force sensing for material classification. They compared neural networks, regularized least squares, and SVMs. Their best model achieved a classification accuracy of 89.5% using 300 training examples across four materials. Chu et al. [6] used a PR2 with two BioTac sensors to estimate haptic adjectives (e.g. soft, slippery, or fuzzy) of 51 everyday objects. Sinapov et al. [7] used a robot with Barrett WAM arms for object recognition of 100 objects spanning 20 categories.

While many modern learning techniques have been relatively underexplored in haptic perception, there are a few works that have employed deep learning. Bell et al. [1] used visual features from the MINC dataset to fine-tune a convolutional neural network (CNN) model for material classification. Gao et al. [8] coupled this visual CNN model with a CNN for BioTac haptics data to estimate haptic adjectives. Yuan et al. [9] showed how long short-term memory and CNN networks can learn the hardness of an object given visual tactile features from a GelSight sensor. Note that all of these works rely on visual features and network models pretrained on large, labeled image datasets, such as VGG-16 [10]. This is due in part to the complexity of collecting labeled haptics data with a robot.

2.2 Semi-Supervised Learning

Several past works have considered semi-supervised learning for the task of image classification. This can be seen in ladder networks by Rasmus et al. [11] and deep generative networks by Kingma et al. [12]. We leverage a deep neural network architecture similar to these works, but we also use adversarial training in order to improve material classification results.

There is little research in robotics on semi-supervised learning for haptic perception. One exception is work by Luo et al. [13] who explored semi-supervised learning for object recognition via ensemble manifold regularization. They used a Phantom haptic device to interact with 12 objects. However, unlike our work, they focused on object recognition and interacted with each object only five times.

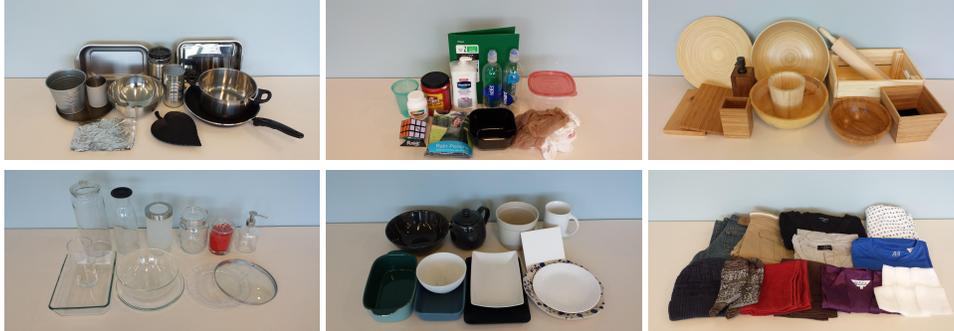


Figure 2: Each of the six material categories with 12 objects per category. From left to right, top to bottom: metal, plastic, wood, glass, ceramic, and fabric.

3 Material Recognition of Everyday Objects (MREO) Dataset

We are releasing a dataset collected from thousands of robot interactions with household objects to support new learning approaches for material recognition. We collected our dataset on a PR2 robot that interacted with 72 objects. We distributed objects across six material categories frequently found in household environments—metal, plastic, wood, glass, ceramic, and fabric. Unlike many past works in material recognition, we do not use flat blocks of materials. We instead use objects which can be commonly found in a household, such as cups, bowls, jars, and clothing. There are 12 uniquely shaped objects per material, as shown in Figure 2.

In order for the robot to interact with an object, we placed a round platter in the robot’s left gripper and then rigidly attached objects to this platter. The robot then performed nonprehensile manipulation with these objects by interacting with the side or top of an object through either horizontal or vertical end effector movements. Holding a platter allowed the robot to rotate objects and interact with most of each object’s exterior surface. This interaction setup is visually displayed in Figure 3. For each interaction, the robot’s end effector follows a linear trajectory using a PID controller. We also uniformly vary the robot’s right end effector velocity from 5 to 10 cm/s between interactions to incorporate variability in the amount of force applied to objects at collision. In total, the robot performed 100 interactions with each object, amounting to a dataset of 7,200 interactions. When compared to prior works in haptic perception for robots, the MREO dataset spans a wide range of household objects and it is currently one of the largest datasets publicly available.

During each interaction, we measured time-varying signals from three different types of sensors mounted on the robot’s right end effector. We used an actively heated point thermistor heated to 55 degrees Celsius prior to contact with an object. During interaction, this thermistor allows the robot to measure the rate of heat transfer into an object. Several previous works in haptics have proposed contact microphones for measuring vibrations during contact [14, 15]. Given this, we attached a contact microphone to the PR2’s fingertip to measure the vibrations that occurred at the moment of impact with an object. We also measured applied forces via two force sensors that are built into the robot’s fingertip. Figure 4 shows all of these sensors. We chose these three modalities—force, temperature, and vibration—since all three are found in the human somatosensory system. We use force and temperature sensing feedback to halt the robot’s movement when contact is made with an object. Contact was detected when either one of the force sensors exceeded a measured reading of 1 N, or the temperature sensor differed by more than 1 degree from its initial set point. During interaction, we sampled the temperature and force sensors at a rate of 100 Hz and 25 Hz, respectively. We sampled the contact microphone at ~ 35 kHz as higher sampling rates are advantageous for capturing vibrations upon impact between the end effector and object. Figure 5 displays examples of force, contact microphone, and temperature signals collected during interactions with an object.

4 Semi-Supervised Learning with GANs

A generative adversarial network (GAN) consists of a minimax adversarial game between a generator, G , and a discriminator, D [16]. The generator aims to output realistic data samples by mapping random noise \mathbf{z} from a noise distribution $p(\mathbf{z})$ to a point in the data distribution. The discriminator

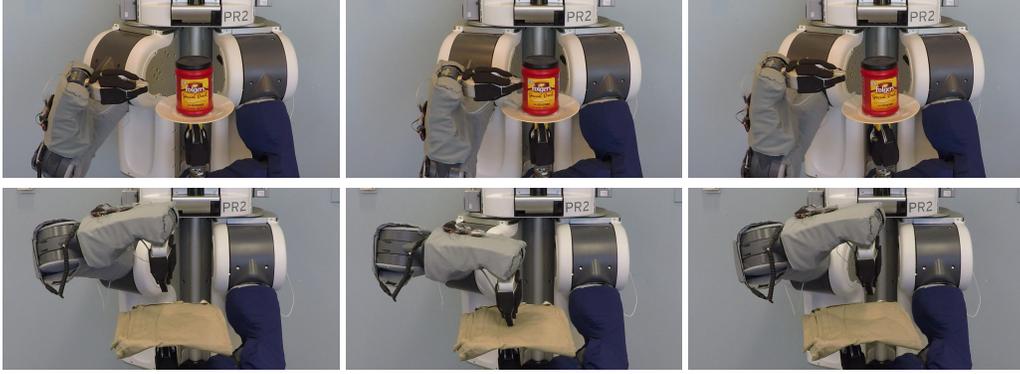


Figure 3: The PR2 interacting with a coffee container and pair of khaki shorts through either horizontal motion (top) or vertical motion (bottom).

is then optimized to differentiate between a real data sample, \mathbf{x} , and a sample from the generator, $\tilde{\mathbf{x}} = G(\mathbf{z})$. Recently, GANs have found success in a range of applications such as image generation, super-resolution imaging, video frame prediction, and image classification [17, 18, 19, 20]. Several works have also extended GANs into semi-supervised learning for image classification [21, 22].

A typical multiclass classifier will label a new data point as one of K discrete classes. For semi-supervised learning, we can add samples from the generator G into our dataset with a new label of $K + 1$, corresponding to a “fake” class. As described by Salimans et al. [20], we can learn from unlabeled data by defining a loss function for the discriminator as,

$$\begin{aligned}
 L &= L_{\text{supervised}} + L_{\text{unsupervised}} \tag{1} \\
 L_{\text{supervised}} &= -\mathbb{E}_{\mathbf{x}, y \sim p_{\text{data}}(\mathbf{x}, y)} \log p_{\text{model}}(y | \mathbf{x}, y < K + 1) \\
 L_{\text{unsupervised}} &= -\mathbb{E}_{\mathbf{x}_u \sim p_{\text{data}}(\mathbf{x}_u)} \log(1 - p_{\text{model}}(y = K + 1 | \mathbf{x}_u)) - \mathbb{E}_{\tilde{\mathbf{x}} \sim G(\mathbf{z})} \log p_{\text{model}}(y = K + 1 | \tilde{\mathbf{x}}).
 \end{aligned}$$

Here, $p_{\text{model}}(y = K + 1 | \mathbf{x})$ represents the probability that \mathbf{x} is fake and corresponds to $1 - D(\mathbf{x})$ from the original GAN architecture [16]. \mathbf{x}_u denotes an unlabeled data sample. We can train the classifier with gradient descent by minimizing both $L_{\text{supervised}}$ and $L_{\text{unsupervised}}$. Intuitively, $L_{\text{supervised}}$ emphasizes correctly classifying a material given labeled sensor measurements, x . The first term in $L_{\text{unsupervised}}$, focuses on classifying unlabeled sensor measurements as one of the K classes, whereas the second term aims to classify images from the generator as fake, i.e. $K + 1$.

At each training iteration, we stochastically update the discriminator by descending the gradient of equation 1:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m -\log \sigma(\mathbf{x}^{(i)})_{y^{(i)}} - \log D(\mathbf{x}_u^{(i)}) - \log(1 - D(G(\mathbf{z}^{(i)}))),$$

where $\sigma(\mathbf{x})_j = p_{\text{model}}(y = j | \mathbf{x})$ is the softmax function applied to the output of the discriminator for all m samples in a minibatch. Salimans et al. [20] proposed a technique, feature matching, to improve classification performance by defining a new cost function for the generator that discourages it from overtraining on the discriminator. We update the generator with feature matching by descending its gradient given by the activations, $\mathbf{f}(\mathbf{x})$, on an intermediate layer of the discriminator:

$$\nabla_{\theta_g} \left\| \frac{1}{m} \sum_{i=1}^m \mathbf{f}(\mathbf{x}_u^{(i)}) - \mathbf{f}(G(\mathbf{z}^{(i)})) \right\|_2^2$$

Our generator has three fully connected neural network layers with a softplus activation function applied to the output of the first two layers and batch normalization between the first and second layer. The generator takes as input a 100-dimensional noise vector sampled from a Gaussian distribution and outputs a vector corresponding to a generated time series of haptic signals.

Our discriminator has six fully connected layers with ReLU activations and Gaussian noise added to the output of each hidden layer for regularization. The vector output of the discriminator can be mapped to a single class estimate by applying an argmax function. We use the Adam optimizer [23] while training both the generator and discriminator with a learning rate of 0.0006 and $\beta_1 = 0.5$.

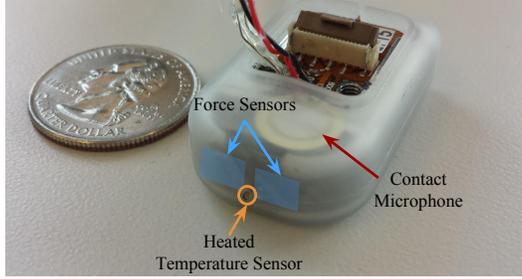


Figure 4: The different sensors used while interacting with objects. We use two force sensors, one actively heated temperature sensor, and one contact microphone all at the PR2’s fingertip. Instructions for replicating the sensors can be found at: <http://healthcare-robotics.com/mr-gan>.

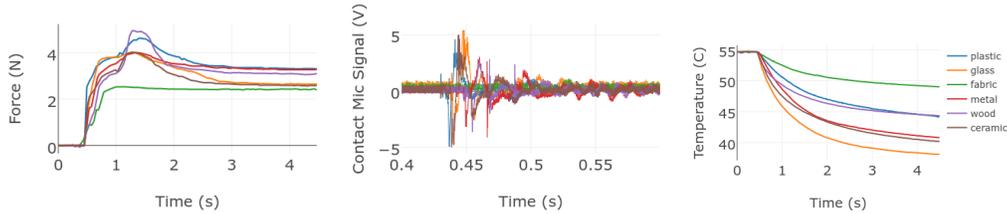


Figure 5: Example measurements from interactions with each of the six materials.

5 Evaluation

We conducted several semi-supervised experiments to evaluate the performance of our generative adversarial approach given unlabeled training samples and new objects. We measured sensor data as the robot interacted with each object, as discussed in Section 3. Prior to learning, we mapped these unprocessed data to vectors of fixed length via linear interpolation over each sensor signal with uniform resampling. We interpolated force and temperature signals to match a sampling frequency of 100 Hz. We interpolated the contact microphone signals to a high quality audio frequency of 48 kHz and then computed the input features as a Mel-scaled spectrogram. We used 128 bands for the Mel filter bank which separates the signal into multiple single-frequency components. Similar to Deng et al. [24], we found that Mel-scaled spectrogram features resulted in stronger classification performance than did raw vibration signals. One second of vibration data resulted in a vector of size 48,000, which we converted to a Mel-scale spectrogram of length 12,032 with 128 frequency bands.

5.1 Semi-Supervised Learning

For semi-supervised learning, we evaluate our GAN’s classification accuracy given varying amounts of labeled and unlabeled training data. Our dataset consists of 7,200 recorded time series, with 1,200 time series per material. We computed results via stratified 6-fold cross-validation. Each fold had 1,200 samples consisting of 200 randomly selected time series from each of the six material classes.

For temperature and force features, we considered only measurements between 0.1 seconds before the robot’s gripper contacts an object and 4 seconds after contact. Similarly, we only used measurements from the contact microphone between 0.1 seconds before contact and 0.1 seconds after contact. We trained each model over 100 epochs.

Table 1 summarizes our results. 1% of labeled training data amounts to providing 10 labeled time series from each of the six classes, with the other 990 provided without labels. We place the remaining 200 time series from each class into the testing fold. With all three modalities and all training data labeled, our classifier can recognize materials with 96.2% accuracy. Our classifier achieves over 95% accuracy when half of the training data were unlabeled and almost 90% accuracy even when 92% of the data were unlabeled. The classification performance tended to drop more rapidly when there were fewer than 80 labeled samples per class and this remains an area for future research. Surprisingly, a contact microphone actively hurt material recognition when less than 4% of the training data were labeled. We suspect this is due to variability in the high frequency vibration signals which makes these signals hard to disambiguate without enough labeled samples.

Table 1: Semi-supervised learning results. Results are presented as accuracies in percent and are averaged over six test sets via 6-fold cross validation. Each column represents the percentage of training data that was labeled during training, with the remaining data provided without labels.

Modalities	Percentage of training data labeled						
	1%	2%	4%	8%	16%	50%	100%
Force	62.1	70.4	72.2	77.7	79.8	85.8	87.9
Temperature	53.8	59.0	64.1	68.1	69.0	80.0	82.1
Contact mic	42.9	53.9	62.6	67.5	73.4	79.8	83.1
Force, Temperature	74.3	81.4	85.6	88.5	90.2	94.2	95.3
Force, Contact mic	58.2	67.5	73.8	80.2	84.7	89.7	91.8
Temperature, Contact mic	52.4	68.3	79.2	84.9	87.4	91.2	92.2
Force, Temperature, Contact mic	62.8	75.4	85.6	89.4	92.0	95.4	96.2

Table 2: Classification results with supervised learning using both an SVM and the discriminator network from the proposed GAN architecture. This compares directly with Table 1 by training the SVM and fully-connected network on only the labeled training data (disregarding unlabeled data).

Supervised SVM with RBF Kernel	Percentage of dataset used for training						
	1%	2%	4%	8%	16%	50%	100%
Force, Temperature	60.4	64.6	69.6	74.1	78.2	86.0	88.8
Force, Temperature, Contact mic	58.9	64.6	70.4	75.8	80.5	88.9	91.8
Supervised Discriminator	Percentage of dataset used for training						
	1%	2%	4%	8%	16%	50%	100%
Force, Temperature	62.9	71.3	78.6	83.3	87.0	91.9	93.6
Force, Temperature, Contact mic	59.1	69.8	78.6	85.8	89.7	92.6	94.8

We also compared these results to that of supervised learning by training on only labeled data and disregarding any unlabeled data. We compared against both an SVM with a radial basis function (RBF) kernel and a fully-connected neural network, which used the discriminator’s architecture from our GAN. Results are shown in Table 2. When less than half of the data are labeled, semi-supervised learning can improve performance by up to 10% over a supervised SVM or neural network.

5.2 Generalizing to New Objects

When robots begin entering real world environments, they will likely be faced with new objects that they have not interacted with before. As such, we also evaluated how well our generative model performs when classifying objects that are not found in the training set. We assessed generalization over a wide variety of objects via leave-one-object-out cross-validation. For this, we trained our model on 71 objects (7,100 interactions) and computed classification accuracy on the remaining 100 interactions from the left out object. We repeated this process for every object and report the resulting accuracy as the average over all 72 training and test sets. We followed a similar training structure as specified above, using the same duration of contact data and training for 100 epochs.

Table 3 shows our generalization results using leave-one-object-out cross-validation with semi-supervised training. Our approach achieved an accuracy of 75.1% with all three modalities and all training sets fully labeled. We note that generalization is challenging for our dataset since most objects are fairly distinct from one another, as seen in Figure 2. This variation is further exaggerated as many objects have paint or coatings, such as the metal frying pan which is the only object with a Teflon coating. We again compare results to supervised learning with an SVM and a fully-connected discriminator, shown in Table 4. For generalization, we observe that semi-supervised learning with a GAN provides little benefit over supervised learning. This remains an open area for further research.

Overall, we found that objects with odd or uneven surface contours posed the largest problem for generalization as these contours could lead to unbalanced force distributions among the two force sensors. Similarly, the temperature sensor could slip along a highly curved surface and give inconsis-

Table 3: Generalization results via leave-one-object-out cross-validation. Each result is presented as the average accuracy over 72 GAN models, where each model is trained on 71 objects and tested on the remaining object.

Leave-One-Object-Out Modalities	Percentage of training data labeled				
	1%	4%	16%	50%	100%
Force, Temperature	65.4	68.1	69.7	68.4	70.3
Force, Temperature, Contact mic	56.4	67.0	70.6	73.6	75.1

Table 4: Generalization results for supervised learning with an SVM and fully-connected discriminator model. Results are computed via leave-one-object-out cross-validation.

Leave-One-Object-Out Supervised SVM with RBF Kernel	Percentage of dataset used for training				
	1%	4%	16%	50%	100%
Force, Temperature	50.9	61.5	67.5	70.7	72.6
Force, Temperature, Contact mic	50.5	59.6	68.0	72.7	75.1

Supervised Discriminator	Percentage of dataset used for training				
	1%	4%	16%	50%	100%
Force, Temperature	58.3	66.1	69.3	72.1	74.9
Force, Temperature, Contact mic	54.8	68.5	74.7	74.6	76.6

tent readings between interactions. For instance, our models had difficulty generalizing classification to the glass and plastic bowls due to their highly curved surfaces. However, our models generalized well to objects such as clothing articles, large cylindrical glass jars, and objects with flat surfaces.

5.3 Varying Duration of Interaction

We also investigated how the duration of contact with an object impacts material recognition. Several past studies, such as work by Kerr et al. [4] and Takamuku et al. [25], have shown that strong material recognition results can be achieved when the robot remains in contact with an object for over 20 seconds. However, these long duration contact requirements reduce the rate at which a robot can interact with and sense its surroundings. To address this issue, we compared how GAN models performed using different sensing modalities and decreasingly short durations of contact.

The overall testing procedure, data processing, and network training remained similar to that of semi-supervised learning in Section 5.1. However, we provided labels for all training data and instead varied the duration of data that we considered once the robot made contact with an object. Table 5 summarizes our results when we trained our GAN architecture on various lengths of interaction.

With 0.5 seconds of temperature and force data, our method achieved a classification accuracy of 92.4% which is just 2.9% lower than when the robot interacted with an object for an entire 4 seconds. As seen in Figure 5, temperature and force signals began to stabilize after 2 seconds of contact, thus leading to no significant performance gains with longer durations of contact.

When contacting an object for 0.2 seconds, our approach was 8% more accurate with a contact microphone than either force or temperature sensing alone. Therefore, high frequency vibrations may be useful for rapid classification upon contact with an object. For example, these vibrations may give a robot insight into the material of an object that it briefly bumps into when manipulation in cluttered spaces. However, contact microphones also have limitations. Namely, these high frequency vibrations lead to large feature vectors which are more difficult to process and evaluate in real time.

5.4 Impact of Unlabeled Data

Once deployed in the real world, it may be valuable for a robot to continue training on unlabeled data that it observes when interacting with objects. Yet, it is unclear if providing more unlabeled data during training will lead to any improvements in performance. To address this, we trained models on increasing quantities of unlabeled data, while holding the amount of labeled training data constant.

Table 5: Classification accuracy as the length of interaction with objects varies. Results are presented in percent and are computed with 100% of the training data labeled. While one second of contact microphone data provided higher accuracy for supervised learning, we found that 0.2 seconds of data improved results during semi-supervised learning.

Modalities	Length of interaction in seconds						
	4s	3s	2s	1s	0.5s	0.2s	0.1s
Force	87.9	87.6	87.6	86.9	81.8	75.1	70.9
Temperature	82.1	80.3	77.5	73.9	70.4	64.4	58.9
Force, Temperature	95.3	94.8	95.0	94.4	92.4	88.6	84.4
	1s	0.7s	0.5s	0.3s	0.2s	0.1s	0.05s
Contact mic	84.6	84.0	83.8	82.4	83.1	77.0	63.3

Table 6: Material recognition accuracy as the amount of unlabeled training data increases. We provided only 40 labeled samples per class during training. The far right column with 960 unlabeled training samples corresponds to the 4% column in Table 1.

40 labeled samples per class Modalities	Number of unlabeled training samples per class						
	0	40	80	160	320	640	960
Force, Temperature	74.8	77.5	78.5	81.1	83.0	85.8	85.6
Force, Temperature, Contact mic	71.9	72.4	76.4	79.3	81.7	83.2	85.6

We followed a similar training setup as in Section 5.1 with 6-fold cross-validation and 100 epochs of training. However, for each training fold we provided 40 randomly selected labeled training examples for each class (4% of the entire training set) and then varied the amount of unlabeled data provided during training. Table 6 presents our results.

With just force and temperature features, our GAN achieved 74% accuracy when trained with no unlabeled data and only 40 labeled samples per material class. With 40 unlabeled samples per class we see a classification improvement of just 2.7%. However, with 960 unlabeled training examples per class—24 times more unlabeled data than labeled—material recognition accuracy increases by over 10%. Given these results, it seems reasonable that a robot could improve its performance by continuing to train on unlabeled data that it observes when operating in everyday environments.

6 Conclusion

We describe how a robot can use generative adversarial networks to learn about material recognition when most of the provided training data are unlabeled. Our robot uses force, temperature, and vibration features in order to differentiate objects given six material categories. Our approach achieves state-of-the-art results and enables a robot to estimate the material class of household objects at almost 90% accuracy, even when 92% of the training data are unlabeled.

We show how a robot can classify an object’s material with 90% accuracy given half a second of contact and all labeled training samples. We compared the semi-supervised GAN to SVM and neural network classifiers trained on only labeled data, and we observed that learning from unlabeled data can improve performance considerably when few labeled examples are available. We also explore how well this generative adversarial approach can recognize the materials of new objects via leave-one-object-out cross-validation. GANs achieved >90% classification accuracy when provided with a few labeled training samples for each object. However, a GAN’s performance dropped by ~20% when classifying new objects with no labeled samples in the training set. It remains an open question how these generative models can more accurately recognize an object’s material when only unlabeled samples from the object are available. Finally, to support further research, we have released the haptics dataset from this work, which was collected on a PR2 robot that performed 7,200 interactions with everyday household objects.

Acknowledgments

We would like to thank Joshua Wade and Tapomayukh Bhattacharjee for their assistance and feedback throughout this work. We thank the developers of Keras, which we used throughout our experiments [26].

References

- [1] S. Bell, P. Upchurch, N. Snavely, and K. Bala. Material recognition in the wild with the materials in context database. *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [2] T. Bhattacharjee, J. Wade, and C. C. Kemp. Material recognition from heat transfer given varying initial conditions and short-duration contact. In *Robotics: Science and Systems*, 2015.
- [3] T. Bhattacharjee, J. Wade, Y. Chitalia, and C. C. Kemp. Data-driven thermal recognition of contact with people and objects. In *Haptics Symposium (HAPTICS), 2016 IEEE*, pages 297–304. IEEE, 2016.
- [4] E. Kerr, T. M. McGinnity, and S. Coleman. Material classification based on thermal properties a robot and human evaluation. In *Robotics and Biomimetics (ROBIO), 2013 IEEE International Conference on*, pages 1048–1053. IEEE, 2013.
- [5] S. Decherchi, P. Gastaldo, R. S. Dahiya, M. Valle, and R. Zunino. Tactile-data classification of contact materials using computational intelligence. *IEEE Transactions on Robotics*, 27(3): 635–639, 2011.
- [6] V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. M. Perez-Tejada, M. Arrigo, N. T. Fitter, J. C. Nappo, T. Darrell, and K. J. Kuchenbecker. Using robotic exploratory procedures to learn the meaning of haptic adjectives. In *2013 IEEE International Conference on Robotics and Automation*, 2013.
- [7] J. Sinapov, C. Schenck, K. Staley, V. Sukhoy, and A. Stoytchev. Grounding semantic categories in behavioral interactions: Experiments with 100 objects. *Robotics and Autonomous Systems*, 62(5):632 – 645, 2014.
- [8] Y. Gao, L. A. Hendricks, K. J. Kuchenbecker, and T. Darrell. Deep learning for tactile understanding from visual and haptic data. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pages 536–543. IEEE, 2016.
- [9] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson. Shape-independent hardness estimation using deep learning and a gelsight tactile sensor. *arXiv preprint arXiv:1704.03955*, 2017.
- [10] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [11] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems 28*, pages 3546–3554. 2015.
- [12] D. P. Kingma, S. Mohamed, D. Jimenez Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems 27*, pages 3581–3589. 2014.
- [13] S. Luo, X. Liu, K. Althoefer, and H. Liu. *Tactile Object Recognition with Semi-Supervised Learning*, pages 15–26. 2015.
- [14] C. L. Chen, J. O. Snyder, and P. J. Ramadge. Learning to identify container contents through tactile vibration signatures. In *Simulation, Modeling, and Programming for Autonomous Robots (SIMPAN), IEEE International Conference on*, pages 43–48. IEEE, 2016.
- [15] R. Murray-Smith, J. Williamson, S. Hughes, and T. Quaade. Stane: synthesized surfaces for tactile input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1299–1302. ACM, 2008.

- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, pages 2672–2680. 2014.
- [17] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *International Conference on Learning Representations*, 2016.
- [18] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.
- [19] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *International Conference on Learning Representations*, 2016.
- [20] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2226–2234, 2016.
- [21] A. Makhzani, J. Shlens, N. Jaitly, and I. Goodfellow. Adversarial autoencoders. In *International Conference on Learning Representations*, 2016.
- [22] J. T. Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *International Conference on Learning Representations*, 2016.
- [23] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] L. Deng, J. Li, J.-T. Huang, K. Yao, D. Yu, F. Seide, M. Seltzer, G. Zweig, X. He, J. Williams, et al. Recent advances in deep learning for speech research at microsoft. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 8604–8608. IEEE, 2013.
- [25] S. Takamuku, T. Iwase, and K. Hosoda. Robust material discrimination by a soft anthropomorphic finger with tactile and thermal sense. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 3977–3982. IEEE, 2008.
- [26] F. Chollet. Keras. <https://github.com/fchollet/keras>, 2015.