Material Classification with Active Thermography on Multiple Household Objects

Haofeng Chen, Haoping Bai, Tapomayukh Bhattacharjee, and Charles C. Kemp

May 2, 2019

Abstract

Active thermography is a technique to inject heat into a target sample and observe the temperature change along time. Such a technique enables a robot to perform material classification with machine learning algorithms and infer material properties of its surroundings. We present a study of material classification on 20 household objects of 5 material classes using active thermography, and analyze factors that impact on material classifiers' performance on generalizing to heating distances and object instances not present during training. By performing a 20-way classification of the object instances, we show that there is potential for classifiers to generalize to unseen objects made from known material classes. The best-performing algorithm trained on 15 object instances at 5 heating distances (20cm, 25cm, 30cm, 35cm, 40cm) gives an accuracy of 71.7% when generalizing to 5 objects that are not in the training set.

1 Introduction

Material classification is a task that can help robots understand the material properties of the nearby environment during manipulation. One widely studied approach is to use contact-based sensors to collect data and use machine learning algorithms to infer material classes. Relevant works collect data by using temperature data given by contact-based active thermal sensor [1], friction and normal force data given by tactile sensor [2], and the combination of different sensing modalities [3].

However, in material classification tasks in a real-world setting, making contact with the target object is not often possible. Researchers develop visual databases of materials like Flickr Material Database [4] and the Materials in the Context Database (MINC) [5] to solve this task with computer vision techniques. Yet, computer vision is often limited in low-light conditions, and sometimes hard to distinguish objects with similar appearance but different materials.

Recently, our work [6] proposed a method to use active thermography, injecting heat into the target object by a heat lamp and record infrared videos, to achieve material classification in a data-driven approach. We treated the temperature data over time at each pixel location to be a single training example and classified 12 materials. However, in the experiment, we used one object sample per material class and did not show the ability of material classifiers to generalize to objects not seen during training. In this study, we classify 5 material classes using 20 object instances and analyze factors that have an impact on material classifiers' ability to generalize to unseen circumstances.

The main contribution of this work is as follows:

- We show that training material classifiers with more heating distances, the distance between the sensing module and the target object, can improve the performance to generalize to unseen distances at test time;
- We demonstrate the similarity between objects with the same material classes and the difference between those with different material classes by a classification experiment on all object instances; the experiment shows the potential for material classifiers to infer material classes for unseen objects at test time;
- We show that training material classifiers with more object instances per material class can improve the classifiers' ability to generalize to unseen objects at test time.

2 Related Work

2.1 Material classification with contact-based sensing

Many researchers study the task of material classification with contact-based sensors. Contact-based sensing requires the sensing module to be in contact with the target object, which can provide a wide variety of information including force, vibration, and thermal data. Because each material has distinct object properties such as elasticity and thermal effusivity, different materials can give different responses given the contact. Researchers then apply machine learning algorithms to the contact data to perform material classification.

2.1.1 Force sensing

Contact-based force sensing can infer many mechanical properties such as surface texture, elasticity, and stiffness. Several researchers [2, 7, 8] developed artificial fingers to recognize material surfaces by analyzing the vibration generated by sliding along the surface. [2] developed a robot finger that can measure the normal and friction force when it comes in contact. They let the finger slide along the 13 target objects each for ten trials, and the Naïve Bayes classifier gives a mean accuracy of 0.885. However, they did not specify the performance of identifying materials with coatings. [9] used the LUCS haptic-hand II, an eight degree-of-freedom (DOF) robot with 45 piezoelectric touch sensors, to infer object shape, hardness, and texture. They used self-organizing maps (SOM) to cluster materials with different properties, but they did not perform material classification.

2.1.2 Active thermal sensing

Active thermal sensing is a technique to observe the temperature of an object after providing it with heat from an external energy source. Many researchers [1, 3, 10, 11, 12, 13, 14, 15] perform material classification with contact-based active thermal sensing. The sensing modules of these studies often consist of a heating element and a temperature sensor. Researchers let the module to inject heat into target objects when the sensor comes into contact, and use the temperature response to classify material types. [10] developed a combined thermal and tactile sensing module and compared the temperature response of 3 different materials after 50 seconds of heat injection. However, they did not use machine learning algorithms to evaluate recognition performance. [1] developed a 1-DOF robot with active thermal sensing module. They performed material classification on 11 materials with 500 trials each. The support vector machine (SVM) algorithm with principal component analysis (PCA) achieved an accuracy of 0.84 with 0.5s of contact data, and 0.98 with 1.5s of contact data. [11] used the SynTouch BioTAC sensor, a multimodal tactile sensor, to measure the derivative of temperature data and other sensing modalities. They achieved an accuracy of 0.99 when recognizing ten different objects.

2.1.3 Other contact-based sensing modalities

[16] used the Inertia-Based Surface Identification System (ISIS), a knocking device with an accelerometer, for material identification. They made the system knock the material sample and observe the acceleration response to infer object hardness, elasticity, and stiffness. They evaluated five different materials with 20 trials each, and the k-Nearest Neighbor algorithm gives an average accuracy of 0.85.

Material classification studies with contact-based sensing modules often require controlled laboratory settings and flat material samples to guarantee perfect contacts. However, when performing material classification in a real-world setting, many objects have irregular geometries and sometimes contact is not even viable. Therefore, it is crucial to study non-contact approaches to this task in order to achieve robust material classification during robotic manipulation.

2.2 Material classification with non-contact-based sensing

2.2.1 Computer Vision

Many researchers take advantage of computer vision techniques for material classification. [17] used an augmented Latent Dirichlet Allocation (aLDA) model under a Bayesian generative framework to identify material types of the images in the Flickr Material Database [4], and achieved an

accuracy of 0.446 in material classification. They claimed that the result outperformed state-of-the-art systems like [18]. [19] researched material classification in an unstructured environment with the Materials in the Context Database (MINC) [5]. They used different algorithms, and the GoogLeNet architecture achieved the highest mean accuracy of 0.852 when classifying 23 different material types. However, using computer vision to classify materials has drawbacks in poor lighting conditions, and when two objects are visually similar but made of different materials. Active thermography is a non-contact-based approach that can potentially solve these problems in this task.

2.2.2 Active Thermography

Active thermography is a technique that involves a heating element and an infrared camera to inspect the material properties of a target object. The heating element injects heat into the object, and then the IR camera gets the temperature data of the heated region. An application is Infrared Non-destructive Testing (IRNDT) [20], a method to inspect defects in laminar materials. After heat injection, anomalies can be observed from the IR image in the defect region of the object sample, if there were any. [21] used non-stationary heat injection schemes when performing IRNDT on fiber-reinforced plastic materials. [22] used IRNDT to detect flaws in different laminate materials and electric circuits.

Recently, we employed active thermography for classification of 12 materials using active thermography in a data-driven approach [6]. In this work, we developed a sensing module with a heat lamp and a thermal camera. We treated each pixel in the thermal video over time as a time series instance and performed multi-class classification with machine learning algorithms. However, in their experiment there is only one object instance per material type; it is not clear whether the classifier is able to identify materials in the wild, generalizing to object instances not present in the training dataset.

In summary, related studies on contact based material classification have the limitation to identify objects with irregular geometries, and those with non-contact based sensing are relatively low in recognition performance. On the other hand, material classification using active thermography is relatively under-explored, especially the ability of classifiers to generalize to unseen objects of known material classes. In this study, we train classifiers to identify materials based on multiple objects per material class at multiple distances, in order to provide guidelines for potential application of material sensing in the wild.

3 Experiment Setup

Following [6], we developed an active thermography sensing module containing a heat lamp and an infrared camera. The two components are mounted so that their relative position is fixed. We attached the module to a linear actuator to make it possible to collect data at different positions of the module.

We placed a polished stainless steel platform to hold the material samples. Since polished metal surfaces are known to have low emissivity [23], it will hardly absorb the heat lamp's radiation and emit heat. Thus, using a stainless steel board as the background platform can guarantee a low interference of the platform to the material sample. [6] mentions that ensuring a small heating angle, the angle between the normal vector of the target object's surface and the line between the heat lamp and the target object, improves material classification performance. The height of the target object is adjusted so that it can be appropriately seen in the thermal video frames, and that the heating angle between the heat lamp and the target object is small. We covered the experiment setup with sheets to minimize external interference on experiment data.

4 Data Collection and Preprocessing

4.1 Data Collection

In the experiment setup, the heat lamp, the infrared camera, and the linear actuator are all controlled by micro-controllers. For each object, we collect the temperature response video mentioned above 10 trials at 5 heating distances (the distance between the heat lamp and the object): 20cm, 25cm, 30cm, 35cm, and 40cm. In each trial, the heat lamp is turned on for 1 second to heat the

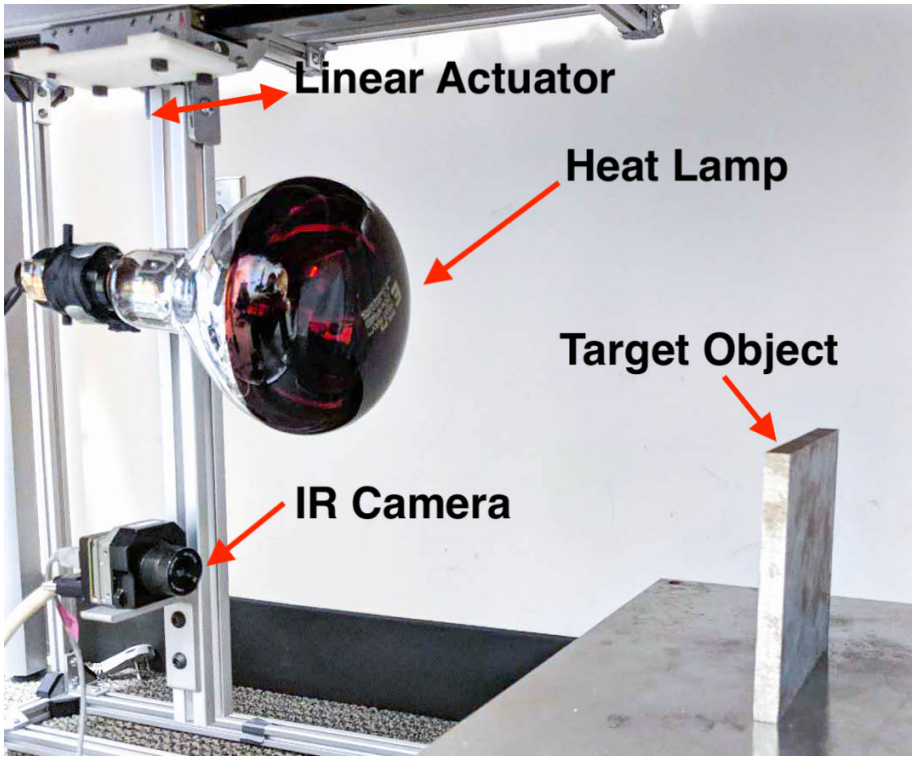

Figure 1: The experiment setup

object and switched off to wait for the material samples to cool down for 90 seconds. The infrared camera, in the meantime, records the temperature response video of the material sample during the 1-second heating phase and 4 seconds of the cooling phase, totaling to 5 seconds. We then use different models to classify materials based on the collected thermal video.

4.2 Object Instances for Material Classification

We choose 5 material types commonly seen in household: wood, plastic, fabric, paper, and metal. We select 4 object instances per material type, and labeled each object with an object index shown in Figure 2 used for training/testing split. All of the material samples are cooled to room temperature before each trial of the experiment.

4.3 Data Preprocessing

As described in previous sections, we collected 10 thermal videos per heating distance per object, totaling 1,000 videos. Each video is approximately 30 frames per second. For each video, we sampled 150 frames uniformly across the total duration of 5 seconds.

Because the target object in the foreground is much closer to the sensing module compared to the background, it is easy to identify these objects in the infrared images. After getting the thermal videos, we convert the frame of each video at 1s to an RGB heat map image and use the Scalabel annotation tool [24] to extract the pixels on the target materials (Figure 3). Following [6], we treat the temperature value of each pixel location over time as a time series instance of the corresponding material class. We only keep the data within the annotated masks in the time series dataset and label them with the corresponding material classes.

Because we employed the polygon-based annotation tool, we can get much more time series and therefore more variance in data compared with [6], who used fixed bounding boxes to extract the time series data. We used the data preprocessing procedure proposed by [6] to normalize each time series given by each pixel by subtracting its mean and dividing by its standard deviation.

For each heating position of each object, we randomly select 2,000 time series data in the first 8 thermal videos to form the training set and 500 time series data in the last 2 thermal videos to form the validation set. The numbers were selected considering the smallest number of pixels

Wood
① Ash
② White Oak
③ Red Oak
④ Walnut

Plastic
① PET
② Thermo Polypropylene
③ Thermo Teflon
④ Polypropylene sheet

Fabric
① T-shirt
② Hat
③ Linen
④ Towel

Paper
① Cardboard
② Newspaper
③ Colored paper
④ Notebook paper

Metal
① Copper
② Cast iron
③ Aluminum
④ Stainless Steel

Figure 2: Objects representing 5 different material classes. Each row consists of a single kind of material, with object indices labeled next to the individual objects.

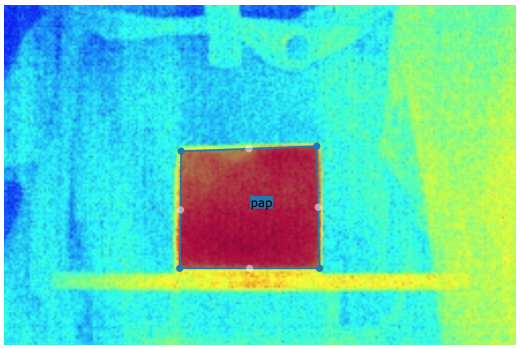

Figure 3: Annotation of the target object using Scalabel

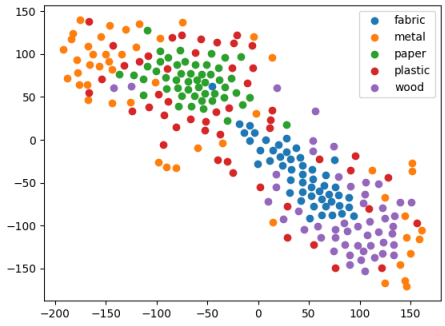

Figure 4: t-SNE plot of sampled time series data

among the annotated masks. We use the same number of time series for each heating distance and each object to avoid the class imbalance problem.

5 Data Visualization

Before performing material classification using deep models, we first make visualizations to have a better understanding of the data. [6] suggests that the normalized active thermography time series data is noisier, therefore hard to classify, when the magnitude of the time series is small, which can be caused by long heating distances and material properties. Here we visualize the time series data of the 5 materials we selected to see the potential effect of their properties on classification performance.

Figure 5 shows the randomly sampled time series data of the 5 material classes. Because each instance is normalized, the time series that has a lower magnitude is noisier, and vice versa [6]. Materials like metal, which often have a glossy surface and low thermal emissivity values, do not absorb heat as much as the others thus resulting in noisier temperature responses. The materials that have higher emissivity values like paper, on the other hand, give rise to time series that are more clearly defined and thus more similar to the thermal response signature of the material as mentioned in [6].

To further understand the time series data, we employ the t-Distributed Stochastic Neighbor Embedding method (t-SNE) [25], a dimensionality reduction algorithm commonly used for visualizing high dimensional data in two or three dimensions. The algorithm minimizes the Kullback-Leibler divergence between a conditional probability distribution of the data points and a conditional probability distribution of the mapped points. The mapping is modeled such that similar data points are closer to each other and vice versa. Figure 4 shows the t-SNE plot of randomly sampled time series data from different material classes. In the plot, one can observe that the data points of wood, fabric, and paper are more concentrated to form clusters, while those of plastic and metal classes scatter around the plot. Note that the time series data of these two classes are also noisier as shown in Figure 5.

6 Material Classification

In this part, we perform classification experiments of materials on the processed time series data. Our previous work [6] shows that there is certain heating distance, the distance between the heat lamp and the target object, that performs better than the others in material classification using active thermography. Such a distance balances between achieving a higher heating intensity and avoiding non-uniform heating patterns of the heat source. We first perform classification experiments to identify such heating distance and evaluate the classifiers' ability to generalize to unseen distances in the training dataset. We then study the intra- and inter-class similarity of the 20 objects and test the performance of classifiers to generalize to unseen objects of the known object classes. Finally, we study the potential use of depth information in material classification.

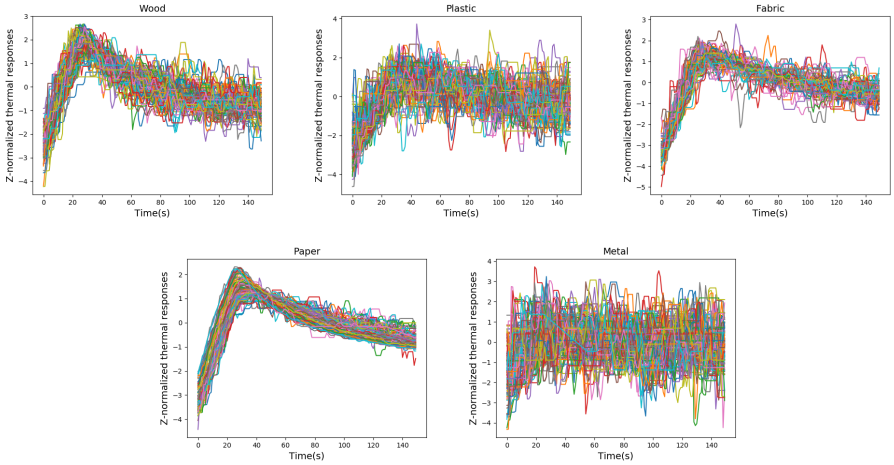

Figure 5: Time series data of the 5 material classes

6.1 Identifying the best-performing heating distance

We first run material classification experiments on different heating distances. For each experiment, we train classifiers on the training set of a given heating distance and test it on the validation set of the same heating distance. For simplicity, we only use object index 1 (see Figure 2) for each material class. Table 1 shows the classification accuracy of these experiments. We can observe that the heating distance of 25cm performs the best, and the best-performing heating distance of this study's experiment setup is between 20 and 30cm. For reference, a classifier that produces outputs by uniformly sampling across the 5 classes has an expected accuracy of 0.2.

Model \Heating Distance (cm)	20	25	30	35	40
VGG13	0.858	**0.917**	0.874	0.829	0.836
LSTM	0.829	0.904	0.861	0.804	**0.887**
ResNet34	**0.868**	0.912	**0.886**	**0.839**	0.884

Table 1: Material classification accuracy at different heating distances

6.2 Generalization to unseen heating distances

Next, we investigate the classifiers' ability to generalize to unseen heating distances. We first train classifiers on the heating distance of 30cm and test them on the validation set of 25 and 35cm. We then train another set of classifiers on the heating distances of 20, 30, and 40cm and test on the same validation set. We only use object index 1 for each material class. The testing accuracy is reported in Table 2.

We can observe that the classifiers that are only trained on the 30cm heating distance can generalize to unseen heating distances, producing comparable results with the one trained on more distances. Such performance is made possible by normalizing each time series, subtracting its mean from itself and dividing it by its standard deviation. This technique enables the classifiers to identify materials by their thermal response signatures [6]. Also, the time series dataset consists of pixel locations across the image frame, which includes a wide variance of data points with different heating distances and heating angles.

However, we can still see an increase in accuracy after including data collected at more heating distances to the training set. Therefore, to improve material classification performance, sampling multiple heating distances is a viable strategy which introduces more variance to the training set and thus improves the classifiers' ability to generalize to unseen distances.

6.3 Evaluation of intra- and inter-class similarity of target objects

We next extend to material classification with multiple objects in each material class. Before experimenting, we start with investigating the similarity of objects with the same material classes,

Model \Trained Heating Distance (cm)	30	20, 30, 40
VGG13	0.841	0.871
LSTM	0.798	**0.872**
ResNet34	**0.858**	0.871

Table 2: Generalization to the heating distances of 25 and 35cm

and that of objects with different material classes. By such an investigation, we evaluate the ability of classifiers to generalize to objects not seen in the training set.

We perform a 20-way classification on all of the 20 objects mentioned in Figure 2 using the ResNet34 model. We label each time series by its object index, instead of their material class as in the previous experiments. Figure 6 demonstrates the confusion matrix of the classification.

The rows of the confusion matrix, as well as the columns, are divided into 5 sections of materials, and each row or column represents a specific instance of the material class, with the object index shown next to it. The columns represent the ground-truth classes while the rows represent the predicted classes. We can observe that objects with the same material classes (the diagonal of the 5x5 "meta-grid") are often hard to distinguish between each other. Such a result is intuitive because objects with the same material classes tend to have similar properties compared with each other. The result also indicates that the classifiers are potentially able to identify objects that are not present in the training set with the material class that has similar properties.

We also see that some material classes are easier to distinguish from each other, such as paper and wood, fabric and plastic, and paper and metal. Such results can be due to the difference in the underlying thermal response signature as mentioned in [6], or that in the noise level of the time series data. This observation also shows that for an unseen object of a particular material class, the classifier is potentially able to tell the difference between the material properties of the target object and the other material classes.

Besides the above observations, we also see the confusion between data points of different material classes in Figure 6, for example, wood and fabric or plastic and paper. Also, the confusion matrix if not necessarily symmetric; given two classes A and B, class A is not often labeled as class B does not mean that class B is not often labeled as class A. For example, in the confusion matrix, we see that metal is more likely to be classified as wood than wood is to be classified as metal. The observation can be explained by the noisy nature of the time series data of the metal class; due to the high variance, some data point of the metal class might fall inside the decision boundary of the wood class. Such confusion between the time series data of different material classes poses challenges to the classification task and might need further investigation.

6.4 Generalization to unseen objects of known material classes

We next evaluate the ability of classifiers to generalize to unseen objects in the training set. We train classifiers on the data with different number of objects. Using the object indices mentioned in Figure 2, we train 3 classifiers on object index 1, indices 1 and 2, and indices 1, 2, and 3, respectively. Then we evaluate each classifier on the data with object index 4. For this experiment, we only use data with the heating distance of 30cm. In Table 3, we see that when more objects are included in the training set, the performance of the classifier improves from 0.584 with only 1 object in the training set, to 0.717 with 3 objects seen during training. Such a result is intuitive because training the classifier with more instances of the material class introduces more variance to the dataset and thus mitigates the bias imposed by specific target objects. When building a classifier to identify materials, it is essential to include a wide range of object instances to improve the performance to generalize to unseen objects.

Model \Trained Object Indices	1	1, 2	1, 2, 3
VGG13	0.555	**0.638**	**0.717**
LSTM	0.510	0.622	0.668
ResNet34	**0.584**	0.625	0.672

Table 3: Generalization to unseen objects of known material classes

We finally train classifiers on the data of multiple objects and multiple distances to see its performance to generalize to unseen objects under such setting. The training set consists of heating

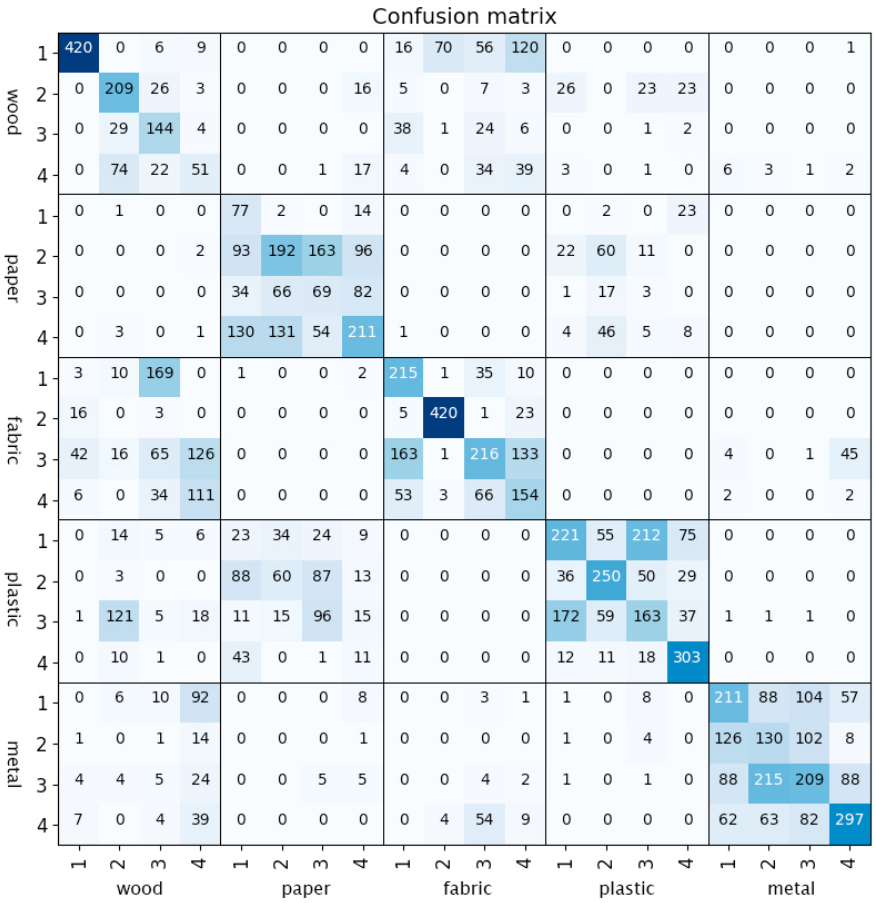

Figure 6: Classification of the 20 objects using ResNet34. Each column represents the ground-truth object index and material class, and each row represents the predicted ones. The diagonal represents the true positive examples. Note that a classifier whose output is uniformly sampled across all 20 classes has an expected number of examples of 25 for each cell in the confusion matrix.

distances of 20, 25, 30, 35, and 40cm, and object indices 1, 2, 3. We evaluate these classifiers on the data with the same heating distances and object index 4. Table 4 shows the classification results.

By comparing to Table 3, which classifies data with only the heating distance of 30cm, there isn't a significant decrease in accuracy of the classifiers identifying materials from different heating distances. An explanation could be training on multiple distances can introduce more variance to the data. Although the testing data now consist of multiple distances, the classifier can maintain classification performance by generalizing better thanks to the increased variance of training data.

Model	Accuracy
VGG13	0.686
LSTM	**0.693**
ResNet34	0.661

Table 4: Material classification trained on heating distances of 20, 25, 30, 35, 40cm, and object indices 1, 2, 3, and tested on the same heating distances and object index 4

7 Discussion

In Section 6 we perform material classification experiments with multiple heating distances and object instances. We first summarize a potential guideline to achieve better performance in material classification using active thermography.

In Section 6.1 we identified the heating distance with the best material classification accuracy, balancing the trade-off between higher heating intensity and a more uniform heating pattern as mentioned in [6]. In Section 6.2, we have shown that training classifiers with data sampled from multiple heating distances improve the accuracy when generalizing across different heating distances. Therefore, to achieve material classification through active thermography, one should first identify the best-performing heating distance of the experiment setup, and train classifiers with multiple sampled heating distances around the best-performing one. When making an inference of material type during manipulation, it is potentially useful to add a depth sensor to the sensing module. Following the guidance of the depth sensor, The robot can manipulate the sensing module so that the heating distance of the target object falls into the desired range.

In Section 6.3 we observe that object instances with the same material classes are often confused with each other, while the ones with different classes are often better distinguished. Such observation confirms the potential of classifiers to generalize to object instances not present in the training set. In Section 6.4 we have shown that training material classifiers with more object instances increases the accuracy when generalizing to unseen objects. Thus, to get better performance when classifying materials in the wild, it is crucial to have a wide range of object samples for each material class.

This work also has its limitations. First, we covered the experiment setup with sheets to lower external impact on the data. We posed the setup such that the sensing module faces perpendicularly to the surface of the target object. Also, all objects are placed on the stainless steel platform during the experiment. Such setup is not realistic during material classification in the wild. Therefore, research is needed to study material classification in real-world and less controlled settings. Also in this study, there are 4 object instances selected per material class. Future work can consider increasing the number of objects to mitigate the bias brought by small sample size.

In addition, when collecting data with various relative poses between the target object and the sensing module, we only used an 1 degree-of-freedom linear actuator to adjust the heating distance. Future work can use actuators with more degrees of freedom to get temperature data with a wider range of heating distances and heating angles, which might facilitate material classifiers in generalizing to different poses because of the increased variance. Future work can also pair calibrated depth sensors with the infrared camera to get depth and surface normal information for each pixel location in the infrared video. Such explicit information of relative pose can potentially enable classifiers to extract foreground pixels and improve classification accuracy by utilizing the relationship between temperature data and the relative pose.

Besides, in this work, temperature data over time at each pixel location is treated as an individual example; no spatial information is used. Future work can make use of the correlation between

pixel locations in the entire thermal video to perform full-frame material-based segmentation with convolutional neural networks.

8 Conclusion

In this work, we conducted experiments on different material classification tasks with active infrared sensing. We have shown that training classifiers with data collected at more heating distances with more object instances increase classification performance when generalizing to unseen circumstances in the training set. In future work, researchers can collect active thermography data with more object instances and a less controlled setting, and potentially incorporate other sensing modalities to achieve robust material classification during robotic manipulation.

References

[1] T. Bhattacharjee, J. Wade, and C. C. Kemp, "Material recognition from heat transfer given varying initial conditions and short-duration contact," Georgia Institute of Technology, 2015.

[2] H. Liu, X. Song, J. Bimbo, L. Seneviratne, and K. Althoefer, "Surface material recognition through haptic exploration using an intelligent contact sensing finger," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 52–57, IEEE, 2012.

[3] 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*, pp. 3977–3982, IEEE, 2008.

[4] L. Sharan, R. Rosenholtz, and E. Adelson, "Material perception: What can you see in a brief glance?," *Journal of Vision*, vol. 9, no. 8, pp. 784–784, 2009.

[5] 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.

[6] H. Bai, T. Bhattacharjee, H. Chen, A. Kapusta, and C. C. Kemp, "Towards material classification of scenes using active thermography," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4262–4269, IEEE, 2018.

[7] M. Tanaka, J. L. Lévêque, H. Tagami, K. Kikuchi, and S. Chonan, "The "haptic finger"–a new device for monitoring skin condition," *Skin Research and Technology*, vol. 9, no. 2, pp. 131–136, 2003.

[8] K. Hosoda, Y. Tada, and M. Asada, "Anthropomorphic robotic soft fingertip with randomly distributed receptors," *Robotics and Autonomous Systems*, vol. 54, no. 2, pp. 104–109, 2006.

[9] M. Johnsson and C. Balkenius, "Sense of touch in robots with self-organizing maps," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 498–507, 2011.

[10] D. Siegel, I. Garabieta, and J. Hollerbach, "An integrated tactile and thermal sensor," in *Robotics and Automation. Proceedings. 1986 IEEE International Conference on*, vol. 3, pp. 1286–1291, IEEE, 1986.

[11] D. Xu, G. E. Loeb, and J. A. Fishel, "Tactile identification of objects using bayesian exploration," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pp. 3056–3061, IEEE, 2013.

[12] J. Engel, J. Chen, Z. Fan, and C. Liu, "Polymer micromachined multimodal tactile sensors," *Sensors and Actuators A: physical*, vol. 117, no. 1, pp. 50–61, 2005.

[13] R. A. Russell, "A thermal sensor array to provide tactile feedback for robots," *The International journal of robotics research*, vol. 4, no. 3, pp. 35–39, 1985.

[14] J.-i. Yuji and K. Shida, "A new multifunctional tactile sensing technique by selective data processing," *IEEE Transactions on Instrumentation and Measurement*, vol. 49, no. 5, pp. 1091–1094, 2000.

[15] P. Dario, D. De Rossi, C. Domenici, and R. Francesconi, "Ferroelectric polymer tactile sensors with anthropomorphic features," in *Robotics and Automation. Proceedings. 1984 IEEE International Conference on*, vol. 1, pp. 332–340, IEEE, 1984.

[16] J. Windau and W.-M. Shen, "An inertia-based surface identification system," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 2330–2335, IEEE, 2010.

[17] C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz, "Exploring features in a bayesian framework for material recognition," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pp. 239–246, IEEE, 2010.

[18] M. Varma and A. Zisserman, "A statistical approach to material classification using image patch exemplars," *IEEE transactions on pattern analysis and machine intelligence*, vol. 31, no. 11, pp. 2032–2047, 2009.

[19] S. Bell, P. Upchurch, N. Snavely, and K. Bala, "Material recognition in the wild with the materials in context database," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3479–3487, 2015.

[20] H. Kaplan, *Practical applications of infrared thermal sensing and imaging equipment*, vol. 75. SPIE press, 2007.

[21] R. Mulaveesala, V. Ghali, and V. Arora, "Applications of non-stationary thermal wave imaging methods for characterisation of fibre-reinforced plastic materials," *Electronics Letters*, vol. 49, no. 2, pp. 118–119, 2013.

[22] G. E. VanDamme and J. W. McGarvey, "Infrared nondestructive testing of laminated structures and electrical circuits," tech. rep., ARMY WEAPONS COMMAND ROCK ISLAND IL GENERAL THOMAS J RODMAN LAB, 1972.

[23] X. Maldague, "Theory and practice of infrared technology for nondestructive testing," 2001.

[24] F. Yu, "Scalabel." https://github.com/ucbdrive/scalabel, commit = a14928b7bbbea2dbc6851b1a1a522ed30e1acdc3, 2019.

[25] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.