

Towards Material Classification of Scenes Using Active Thermography

Haoping Bai^{*1}, Tapomayukh Bhattacharjee², Haofeng Chen¹, Ariel Kapusta¹, and Charles C. Kemp¹

Abstract—By briefly heating the local environment with a heat lamp and observing what happens with a thermal camera, robots could potentially infer properties of their surroundings. However, this form of active thermography introduces large signal variations compared to traditional active thermography, which has typically been used to characterize small regions of materials in carefully controlled settings. We demonstrate that a data-driven approach with modern machine learning methods can be used to classify material samples over relatively large surface areas and variable distances. We also introduce the use of z-normalization to improve material classification and reduce variation due to distance and heating intensity. Our best performing algorithm achieved an overall accuracy of 77.7% for multi-class classification among 12 materials placed at varying distances (20 cm, 30 cm, and 40 cm). The observations were made for 5 seconds with 1s of heating and 4s of cooling. We also provide a demonstration of performance with a multi-material scene.

I. INTRODUCTION

Active thermography involves heating a target region of interest and observing its response over time with a thermal camera. Many robotics applications could potentially benefit from active thermography, since it allows heat-transfer based sensing without the necessity of contact. A robot can provide thermal excitation to a region of interest and make inferences based on the captured thermographic video. In this work, we specifically apply active thermography to material classification.

We treat material classification as a scene labeling problem. We refer to the values of a pixel in the thermographic video of a static scene over time as the *thermal response* corresponding to that pixel location. Our goal is to assign a material label to each pixel location in the thermographic video based on the recorded thermal response.

In comparison to the laser heat excitation used recently in [1] for material classification, we use a heat lamp as the heat source so that we can provide heat excitation to an entire region at once. However, a regional heat source also introduces challenges: 1) some heat lamps have nonuniform heating patterns that can introduce variation across a surface of the same material; 2) unlike a laser heat source which concentrates energy along a single direction, a heat lamp radiates its heat in a cone. As a result, the relative position between the heat source and target surface can impact heating intensity and alter material thermal responses. We propose to apply z-normalization [2] (normalization to zero mean and

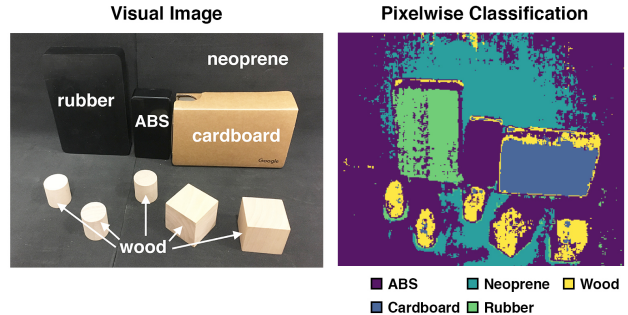


Fig. 1: Demonstration of pixel-wise material classification performance with a multi-material scene. Left: visual scene. Right: predicted class labels of each pixels of the scene.

unit of energy) to each thermal response to mitigate data variability due to such factors. We observed considerable improvements in material classification performance with z-normalized thermal responses. We then provide a demonstration of material-based scene labeling by performing material classification on a pixel-by-pixel basis over an entire scene (see Fig 1).

The key contributions of our work follow:

- We present a system that actively provides heat excitation to a larger region of a visible scene and assigns material labels to each pixel location in the recorded thermographic video.
- We introduce z-normalized thermal responses as an effective representation that is diagnostic to materials and invariant to heat intensity, which allows us to identify materials with non-uniform heating from varied distances.
- We provide a classification baseline for active thermography data by evaluating the performance of selected classification algorithms.
- We provide an empirical evaluation of strengths, weaknesses, and interesting phenomena of material classification using active thermography.

The problem of material-aware scene understanding is a common problem for robotic perception. We expect our work to lay out a baseline for further exploration in non-contact, heat-transfer based material classification.

II. RELATED WORK

A. Active Thermography

Active thermography is widely applied in infrared non-destructive testing (IRNDT) to detect internal defects in test samples [3]–[8]. In this section, we will discuss both model-based applications and data-driven applications of active thermography.

¹H. Bai, H. Chen, A. Kapusta, and C. C. Kemp are with the Healthcare Robotics Lab, Institute for Robotics and Intelligent Machines, Georgia Institute of Technology

²T. Bhattacharjee is with the Paul G. Allen School of Computer Science and Engineering, University of Washington

Work was performed at the Healthcare Robotics Lab, Institute for Robotics and Intelligent Machines, Georgia Institute of Technology

^{*}H. Bai is the corresponding author {fbai@gatech.edu}




Fig. 2: Active thermography sensing module consisting of a heat lamp and a thermal camera

1) *Model-based Applications*: According to [3], conventional active thermography applications usually rely on analytical models or finite difference simulations to reverse engineer observed thermal responses for quantitative material characterization [9]–[12]. For example, Parker [9] introduces the flash method, which measures thermal diffusivity, thermal conductivity, and heat capacity of a thin, camphor-coated test sample. In [12], the authors propose an in-situ method to measure the thermal diffusivity with angle-adjustable Gaussian beam and an IR camera.

2) *Data-Driven Applications*: Data-driven algorithms provide an alternative to the complex inverse modeling by learning to make predictions directly from data. Maldague [13] provides an example of using trained neural networks to detect corrosion in aluminum specimens. Halloua [14] proposes a method to create 3D reconstruction of the subsurface depth geometry by training a neural network to estimate the depths of subsurface defects. In [15], the authors train supervised learning algorithms to identify bruised apples based on the thermal responses.

In the context of active thermography for material classification, deriving an analytical model can be infeasible due to real-world complexities. In our work, we employed a data-driven approach with the z-normalization technique backed by a simple model.

B. Material Classification

In the field of robotics, there are a variety of sensing modalities that can identify materials based on their properties. Some approaches rely on contact with the target material, such as force and vibration sensing that identify materials based on elasticity and texture [16]–[18]. Another example is heat-transfer-based material classification based on tactile sensors [19]–[26], which often seeks to identify materials by their thermal responses over time.

When contact is not viable, non-contact sensing modalities can be useful. Many researchers take advantage of computer vision to perform material classification [27]–[30]. However, such approaches are not robust to illumination conditions.

There are few methods that employ active thermography for material classification. In [1], the authors demonstrate a system that uses a single laser to induce thermal responses in each of the five selected high-emissivity materials and classifies materials based on the thermal responses. An advantage




Fig. 3: The view factor of two differential areas is calculated based on the relative position (d, θ, θ_0) of the areas, as defined in Eq. (1).

of a laser heat source is that it can provide consistent and concentrated energy excitation over a small area. However, looking at a surface point-by-point can be inefficient for gathering information from a scene. In our work, we use a heat lamp as the excitation source, which makes it possible to collect a region of data using a thermal camera for a single heat excitation. The z-normalization technique further allow us to train classify algorithms that are more robust to data variations introduced by varying heating positions. We perform material classification on 12 materials of distinct properties with different classification algorithms.

III. THEORETICAL INSIGHTS

A. Heat Flux into the Material Surface

To derive a strategy for reducing unwanted variance in data, we look into the heat transfer process of active thermography. First, we derive an expression for the rate of heat absorption at the surface of the material sample. We consider the heat lamp and each pixel region to be two differential areas. The proportion of heat flux (flow of energy per unit area per unit time) leaving the heat lamp that strikes the surface of the material sample is given by the view factor [31] F , with

$$F = \frac{\cos \theta_0 \cos \theta}{\pi d^2} \quad (1)$$

where θ_0 and θ are the angles between the corresponding surface normal (\mathbf{n}_1 for dA_1 and \mathbf{n}_2 for dA_2) and the line between the two differential areas, and d is the length of the line, as shown in Fig. 3. Let the heat flux leaving the heat lamp be \dot{Q} , we can then express the heat flux into the material sample, \dot{q}_s , as

$$\dot{q}_s = \alpha F \dot{Q} = \alpha \frac{\cos \theta_0 \cos \theta}{\pi d^2} \dot{Q} \quad (2)$$

where F is the view factor. Absorptivity α is the surface's efficiency in absorbing incident thermal radiation [31]. The value of absorptivity ranges from 0 to 1.

Assuming no heat source beneath the material surface, heat transfer inside a material is usually modeled using the heat equation [32], which, in this case, is a linear homogeneous second-order partial differential equation with spatial and temporal dimensions and constant coefficients. Intuitively, in our scenario, assuming homogeneity implies that applying twice the heat flux at the surface will result in twice the magnitude of the observed thermal response.

Considering the complications of material properties and the heating process in a real-world scenario, we provide a simplified model that decouples the observed thermal response into two factors: the heat absorbed by the material and the material properties that induce the heat response given




Fig. 4: 12 selected materials for classification: 1) ABS, 2) Aluminum, 3) Rubber, 4) Concrete, 5) Polyester Foam, 6) Cardboard, 7) Neoprene, 8) Pine Wood, 9) Clay, 10) EPS Foam, 11) Graphite, 12) Stainless Steel

the absorbed heat. Assuming the homogeneous property of the heat equation, we can approximate the relation between the thermal response and the absorbed heat flux with the following equation:

$$T(t) - T(0) = \dot{q}_s s(t) \quad (3)$$

where $T(t)$ is the temperature of the material surface as a function of time, and $T(0)$ is the initial surface temperature just before heating, which we assume to be the ambient temperature. \dot{q}_s is the absorbed heat flux as quantified by eq. (2). We define $s(t)$ as the *thermal response signature* of the material, which we can treat as a black box. Ideally, we can classify materials based on $s(t)$ because it is assumed to be invariant to the heat flux \dot{q}_s and diagnostic of the properties that affect the heat transfer and heat propagation *inside* the material, such as internal structure, thermal effusivity, and thermal conductivity.

Next, we derive an expression for the thermal response that the thermal camera observes. All objects emit a certain amount of infrared radiation as a function of their temperature [33]. A thermal camera can detect this radiation in a way similar to how an optical camera detects visible light [3], [34], [35]. The Stefan-Boltzmann Law quantifies this amount of radiation (radiance) of a given surface:

$$j = \epsilon \beta T^4 \quad (4)$$

where the radiance j has dimensions of energy flux (energy per unit time per unit area). T is the surface temperature of the object, β is the Stefan-Boltzmann constant, and ϵ is the emissivity, which is the effectiveness of the surface in emitting thermal radiation [31]. The value of emissivity ranges from 0 to 1, with 1 being the emissivity of an ideal black body. As a result, low emissivity surfaces such as shiny or polished metals [36] can be difficult to observe with thermography because there is not much thermal emission to observe.

Usually, IR cameras require manual inputs of target surface emissivity to correct the raw readings. However, since the emissivity of a given material sample is often unknown, we configure the thermal camera to report data assuming the emissivity of the target surface is 1 [37]. If the actual




Fig. 5: Left: Diagram of the heat transfer process. Right: Diagram of the heat excitation and observed thermal response.

emissivity of the surface is ϵ , then the relation between the actual temperature T of the surface and the reported temperature by the thermal camera y can be derived as

$$j = \epsilon \beta T^4 = \beta y^4 \implies y = \sqrt[4]{\epsilon} T \quad (5)$$

Based on Eq. (2), Eq. (3), and Eq. (5), we have

$$y(t) - y(0) = \sqrt[4]{\epsilon} [T(t) - T(0)] = \gamma s(t) \quad (6)$$

where $\gamma = \sqrt[4]{\epsilon} \dot{q}_s = \sqrt[4]{\epsilon} \alpha \frac{\cos \theta_0 \cos \theta}{\pi d^2} \dot{Q}$

We thus express the thermal response $[y(t) - y(0)]$ as $s(t)$ scaled by the factor γ . We refer to γ as the scaling factor for the thermal response signature. We define the *magnitude* of the thermal response be the standard deviation of $[y(t) - y(0)]$ over time. This magnitude is proportional to γ , which is determined by the heat flux from the heat lamp \dot{Q} , the heating position (d, θ) , and material surface properties α and ϵ .

In real-world scenarios, it is difficult to directly measure the value of the factor γ to extract $s(t)$ from $y(t)$. Therefore, we eliminate the effect of γ on $y(t)$ by performing z-normalization along the time dimension, which gives us

$$\hat{y}(t) = \frac{[y(t) - y(0)] - \mu}{\sigma} \quad (7)$$

where $\hat{y}(t)$ is the z-normalized thermal response, μ is the mean of the thermal response $[y(t) - y(0)]$ over time and σ is the standard deviation, or the magnitude of the thermal response. We will next illustrate the effect of z-normalization with data observations and material classification experiments.

IV. DATA COLLECTION AND PROCESSING

A. Material Selection

To include a variety of materials samples with different properties, we selected 12 materials to represent the four material categories (metals/alloys, ceramics/glasses, polymers/elastomers, and composites) found in the CES Edupack database [38]. These material samples are shown in Fig. 4.

B. Experiment Setup

Figure 2 shows the active thermography sensing module used in our experiments, consisting of a FLIR Tau 2 324 7.5mm Thermal Imaging Camera Core (46324007H-FRNLX) equipped with Pleora's iPORTTMCL-U3 External Frame Grabber. We fixed the relative position between the thermal camera and the heat lamp. We also adjusted the




Fig. 6: Formation of training/validation datasets

thermal camera to an angled position to avoid the reflection of the heat lamp.

C. Data Collection

For the experiments, we centered material samples directly beneath the filament of the heat lamp as shown in Fig. 5 (left). The length of the line between the center of the sample and the filament corresponds to the heating distance d in eq. (1), and the angle between the line and the surface normal corresponds to the view angle θ in eq. (1). We call the configuration (d, θ) the heating position of the sensing module.

We performed four sets of experiments with varied heating distances: $d = 20, 30, 40$, and 50 cm, $\theta = 0^\circ$. Each experiment used 100 trials of data collection and yielded 100 thermographic videos. While it is usually suggested to introduce heat flux normal to the material surface (zero view angle) to maximize absorption and emission [39], we collected 100 trials of data for the heating position $d = 30$ cm, $\theta = 30^\circ$ to observe the effect of a non-zero view angle θ on classification.

During each trial of data collection, the heat lamp automatically turns on and the IR camera starts recording. The heat lamp turns off after 1 second and the thermal camera stops recording after 5 seconds. After the recording ends, the material sample is left to cool down for 90 seconds, which we observed to be sufficient for each material to return to the ambient temperature.

D. Data Processing

The thermal camera produces video frames of size 324×256 at 60 Hz. For the video of each heating position (d, θ) , we manually selected a 32×32 window in the thermal video that contained no background pixels to extract material thermal responses. Doing so ensures all the materials have the same amount of data. Each of the $32 \times 32 = 1024$ pixels within the window corresponds to a time series of pixel values. All pixels have a known ground truth material class. For simplicity, we assume that each pixel in the window has the same d and θ values.

We synchronized all the time series based on their corresponding timestamps to a uniform 5 second time vector with size $N = 200$ by linear interpolation. We then independently normalized each time series using z-normalization [2] for relevant experiments.

For classification, we assigned training and validation videos with stratified 10-fold cross-validation. To make sure no pixel location was observed during both training and




Fig. 7: Magnitude comparison of the average thermal response of ABS with different heating positions




Fig. 8: Magnitude comparison of the average thermal responses across materials at heating distance $d = 30$ cm with $\theta = 0^\circ$

validation, we split the 32×32 window of each video into two regions by specifying a "training window" in the center. We only kept the pixels *within* the training window of the training videos to form the training set, and the pixels *outside* the training window of the validation videos to form the validation set, as shown in Fig. 6. Unless mentioned, we use a training window of size 24×24 and report the average of cross-validation accuracies.

V. OBSERVATIONS FROM DATA

Based on the collected data, we validated the scaling effect of the factor γ over material thermal responses. Fig. 7 shows that the magnitude of thermal responses of a material increases as the heating distance d decreases or when the heat flux is normal to the surface (view angle $\theta = 0^\circ$). This result corresponds with eq. (7) for which the magnitude of the thermal response is proportional to γ . Thus, a small heating distance d or a small view angle θ may result in a higher magnitude of the thermal response.

We also observed that different materials have different magnitudes of thermal response with the same heating position, as shown in Fig. 8. Materials with high response magnitudes tended to have rough surfaces and dark color, and those with low response magnitudes were metals or had smooth/polished surfaces. Such an observation could be due to the difference of materials in absorptivity α and emissivity ϵ in Eq. (7). With greater α and ϵ , we would expect the magnitude of the thermal responses to be higher due to the increased absorption and emission of energy. Note that the magnitude of the thermal response signature $s(t)$ can also contribute to the overall magnitude of the thermal response of a material.




Fig. 9: Z-normalized thermal responses (over 5 seconds) of materials from varying heating distance and view angle, averaged over 100 trials. **Green:** 20 cm, 0°, **Red:** 30 cm, 0°, **Blue:** 40 cm, 0°, **Black:** 50 cm, 0°, **Yellow:** 30 cm, 30°. We can observe that the z-normalized thermal responses of most materials are invariant to heating distance and view angle.




Fig. 10: Noise level comparison of z-normalized thermal responses of ABS obtained at $d = 30$ cm (Top) and 50 cm (Bottom) with $\theta = 0^\circ$.

We then illustrate the thermal response signature by comparing z-normalized data collected at different heating positions. Fig. 9 shows the average z-normalized thermal responses at $d = 20, 30, 40, 50$ cm with $\theta = 0^\circ$ and $d = 30$ cm with $\theta = 30^\circ$ for each material. We observed that for most materials, thermal responses from varying distances and view angle align well after z-normalization. Since z-normalization can eliminate the scaling effect of γ , it can potentially make material classification more robust against environmental variations such as the heating intensity of the heat source \dot{Q} and heating geometry (d, θ) .

Another observation is that a higher magnitude of thermal response corresponds with a higher signal to noise ratio after z-normalization. With a higher magnitude, the z-normalized time series data becomes more representative of the thermal response signature of the material sample and less dominated by noise. By Eq. (7), a lower heating distance d , a lower view angle θ and materials with high absorptivity α and emissivity ε tend to result in a lower noise level in the z-normalized thermal response. Fig. 10 illustrates this effect with z-normalized time series data collected at different heating distances.

VI. CLASSIFICATION RESULTS

We performed two sets of classification experiments. We first evaluated the effect of the scaling factor γ on classification performance by analyzing the components of γ .

Then, we performed an ablation study of z-normalization to examine its role on classification performance.

A. Effect of γ on Classification Performance

We performed a multi-class classification of the 12 selected materials. We analyzed the effect of γ , a function of the heating position (d, θ) and material properties, on classification accuracy. As a reminder, lower d , lower θ and higher α and ε can lead to higher γ , as indicated by Eq. (7).

In this section, we use the Gaussian Naive Bayes classifier (GNB) for its simplicity. With the GNB classifier, each class corresponds to a data-generating normal distribution centered around a mean feature vector that is analogous to the thermal response signature. The covariance matrix of the distribution models the effect of noise on the time series data.

1) *Effect of Heating Distance d :* Table I (left) shows the classification accuracy on the data with different heating distances. When the heating distance d decreases from 50 cm to 30 cm, the accuracy increases. As discussed in Section V, with lower d , γ becomes higher. The z-normalized time series data then become less dominated by noise, and it is easier to classify the thermal response according to the underlying thermal response signature with a lower heating distance.

However, as d continues to decrease, the accuracy of the $d = 20$ cm data drops. This may be due to the effect of the non-uniform heating pattern of the heat lamp. As the sensing module moves closer to the material sample, the non-uniform heating effect of the heat lamp becomes more significant with a clearly visible pattern, which can lead to more variations in the collected data.

2) *Effect of View Angle θ :* Table I (right) shows the classification accuracy on the data with zero and non-zero view angles. We observed that a zero view angle yields better classification accuracy. A zero view angle leads to higher γ in Eq. (7) and thus the resulting time series are less dominated by noise, which can improve classification performance.

TABLE I: Effect of heating distance (left) and view angle (right) on the multi-class classification cross-validation accuracy

Experiment	Accuracy (%) mean (SD)	Experiment	Accuracy (%) mean (SD)
20 cm, 0°	72.6 (3.35)	30 cm, 0°	81.4 (3.25)
30 cm, 0°	81.4 (3.25)	30 cm, 30°	71.7 (3.02)
40 cm, 0°	73.9 (3.09)		
50 cm, 0°	58.1 (3.44)		




Fig. 11: Confusion matrix of GNB multi-class classification with the $d = 30$ cm, $\theta = 0^\circ$ data. Materials are sorted based on the magnitude of their thermal responses. Materials with low thermal responses tend to have low cross-validation accuracies.

3) *Effect of Surface Properties α and ϵ* : Fig. 11 shows the confusion matrix of classification on the $d = 30$ cm, $\theta = 0^\circ$ data. The material classes, from top to bottom, are sorted by the average magnitude of their thermal responses. With the same heating position, a higher thermal response magnitude can be attributed to higher absorptivity and emissivity. We observe that the materials with high thermal response magnitudes like neoprene and wood have higher cross-validation accuracy. The classifier encounters the most confusion among materials with low thermal response magnitudes, possibly due to lower energy absorption/emission resulting in z-normalized thermal responses with more noise, as suggested by Section V.

The classification results suggest that our system has higher classification performance over a range of heating distances and material types. In general, performance is better with a zero view angle on materials with high absorptivity and emissivity. When adopting active thermography for material classification, it is important to consider the range of heating distances over which the sensing is effective.

We next investigate the effect of z-normalization on material classification with an ablation study, to see whether this technique can improve classification performance on data with more variation.




Fig. 12: Learning Curve of GNB classifier with respect to different training window sizes

B. Ablation Study of Z-Normalization

To demonstrate the effect of z-normalization on classification performance, we trained a GNB classifier with different training window sizes. The smaller the training window size, the less variation of data the classifier can encounter during the training period. Fig. 12 shows the cross-validation accuracy of the GNB classifier with respect to training window size. We can observe that the classification accuracy on the z-normalized thermal responses is consistently higher than without z-normalization. With z-normalization, the classifier is able to learn a more robust representation of a material class across thermal responses with varying γ , making classification more accurate - especially with combined data of varying distances.

We next used the combined data of $d = 20, 30, 40$ cm with a fixed training window size of 24 to train and evaluate multi-class classifiers using 10-fold cross-validation. We did not include $d = 50$ cm data because of the significantly lower classification performance. At 50 cm, we observed that our heat lamp does not produce thermal responses strong enough for our selected algorithms to reach good material classification accuracy. We performed classification on the dataset with and without z-normalization to evaluate the effect of this technique on classification performance. We evaluated the performance of the following four algorithms:

- Gaussian Naive Bayes (GNB) classifier [40]. We used the Scikit-Learn implementation [41] of GNB with default parameters.
- Very Deep Convolutional Network (VGG) [42]. We used the VGG13 (with batch normalization) variation of the network [42] as provided by the PyTorch Model Zoo [43], and replaced 2D operations on an image with corresponding 1D operations on a time series.
- Residual Network (ResNet) [44]. We used the ResNet34 implementation provided in the PyTorch Model Zoo [43], and replaced 2D operations on an image with corresponding 1D operations on a time series.
- Long Short Term Memory (LSTM) [45]. We constructed a 3-layer, bi-directional LSTM with 120 hidden units and 0.2 dropout rate using the nn.LSTM module in

TABLE II: Cross-validation accuracies (%) of classifier on each material class. Materials are sorted from left to right by their thermal response magnitude in decreasing order. Algorithms are tested with/without z-normalization as indicated by the z-norm column. The All column denotes the overall cross-validation accuracy across all material classes, with SD indicating the standard deviation of the cross-validation accuracies. The best accuracy achieved for each material are marked in bold.

	z-norm	Neoprene	Polyester Foam	ABS	Cardboard	Rubber	Wood	Clay	EPS Foam	Concrete	Stainless Steel	Aluminum	Graphite	All mean (SD)
GNB	\times	77.3	76.6	25.1	17.7	26.2	8.8	29.7	44.1	78.4	58.4	38.9	42.0	43.6 (0.52)
	\checkmark	93.2	89.1	78.6	63.7	85.4	81.6	40.7	58.4	57.1	33.6	65.6	54.0	67.1 (0.55)
VGG	\times	97.7	95.2	86.3	76.1	92.5	92.7	64.0	78.2	73.7	47.5	59.9	58.5	76.9 (0.88)
	\checkmark	97.5	96.3	86.4	76.2	94.0	93.0	66.1	78.5	73.9	49.0	61.1	60.0	77.7 (0.93)
LSTM	\times	85.9	93.2	59.2	55.5	59.9	73.4	45.6	65.4	63.3	48.5	36.6	52.6	61.6 (14.1)
	\checkmark	98.8	95.5	85.6	73.6	94.5	93.1	63.7	80.7	66.6	47.2	53.3	64.1	76.4 (0.57)
ResNet	\times	97.6	94.3	82.8	72.5	93.5	90.8	63.7	80.9	69.8	48.6	55.4	57.6	75.6 (0.70)
	\checkmark	98.0	95.7	85.3	73.2	93.7	92.1	61.9	76.3	69.6	47.1	54.0	53.9	75.1 (0.99)

PyTorch [43]. The output of the LSTM is processed by a multilayer perceptron with one hidden layer consisting of 120 ReLU units with a dropout rate of 0.2 to produce class predictions.

We trained VGG, ResNet, and LSTM for 10 epochs with early stopping using the Adam optimizer with default settings.

Table II shows the multi-class classification performance of the selected models. We observed that z-normalization improves the accuracy in most circumstances, especially in the case of GNB. The improvement in GNB classification accuracy suggests that z-normalization moves thermal responses of a material towards a distribution centered around an underlying thermal response signature, as discussed in Section V. LSTM also receives considerable performance gain from z-normalization. Note the standard deviation of cross-validation accuracies of LSTM is much smaller than that without z-normalization, as shown in Table II, which means that z-normalization helps to stabilize the training of LSTM to achieve consistent performance. Judging from classification performance across materials, VGG and ResNet benefit less from z-normalization than does LSTM. One possible explanation is that the batch normalization layers [46] in the networks mitigate the effect of data variation during training. However, we did observe faster convergence for both algorithms during training when z-normalization was applied to the data.

VII. DISCUSSION

Our results suggest that active thermography can be used by robots to efficiently acquire information about the materials in their surroundings. This approach works better for materials with high absorptivity and high emissivity, and has difficulty with low-absorptivity and low-emissivity materials. In addition, a small view angle θ and a carefully chosen heating distance d can improve performance. The potential performance over larger volumes, including farther away from the center of heating and observation, remains unclear.

Z-normalization can help classification algorithms handle signal variations caused by varying heating intensity. This is relevant to the problem of performing material classification in scenes with real-world geometries. As an initial assessment of the feasibility of this application, we set up a qualitative experiment of pixel-wise classification of a

scene. We selected 5 materials with high thermal response magnitudes: ABS, cardboard, neoprene, rubber and wood, and trained a VGG13 classifier with the data collected from corresponding material blocks. Then, we set up a scene with objects in the real world: an ABS box, a cardboard box, a rubber block, several wood blocks, and a neoprene mat as background. Fig. 1 shows the scene and the classification result. We can observe that with z-normalization, the classifier is able to extend predictions to real-world geometries with non-zero view angles and curvatures near the center of the heated region. This demonstrates the potential of active thermography for robotic perception in the wild, where a robot could heat up a given scene and use a pre-trained classifier to label the material associated with each pixel’s time series. The quantitative performance of this form of scene perception under real-world conditions remains an open question.

VIII. CONCLUSION

We investigated region-based material classification enabled by active thermography. Using a heat lamp as the heat source, we applied heat excitation to 12 materials from varying heating locations and recorded corresponding thermal responses with a thermal camera. We also demonstrated the consistency between z-normalized thermal responses from a material despite the variations in heating distances and angles. We ran specifically designed classification experiments with GNB to demonstrate how factors including heating distance, angle and material surface property impact material classification performance. We benchmarked different classification algorithms to provide a baseline for material classification. VGG achieved the best performance with an overall accuracy of 77.7% for multi-class classification among 12 materials placed at varying distances (20 cm, 30 cm, and 40 cm). We also performed a pixel-wise classification of a scene using a pre-trained classifier trained on material blocks, which shows the potential of using active thermography for material-aware scene understanding.

ACKNOWLEDGMENT

This work was supported in part by NSF Awards EFRI-1137229 and IIS-1150157, the National Institute on Disability, Independent Living, and Rehabilitation Research (NI-DILRR) grant 90RE5016-01-00 via RERC TechSAGE, and a Google Faculty Research Award.

REFERENCES

- [1] T. Aujeszkzy, G. Korres, and M. Eid, "Thermography-based material classification using machine learning," in *2017 IEEE International Symposium on Haptic, Audio and Visual Environments and Games (HAVE)*, Oct 2017, pp. 1–6.
- [2] D. Q. Goldin and P. C. Kanellakis, "On similarity queries for time-series data: constraint specification and implementation," in *International Conference on Principles and Practice of Constraint Programming*. Springer, 1995, pp. 137–153.
- [3] X. P. Maldague, "Introduction to ndt by active infrared thermography," *Materials Evaluation*, vol. 60, no. 9, pp. 1060–1073, 2002.
- [4] H. Kaplan, *Practical applications of infrared thermal sensing and imaging equipment*. SPIE press, 2007, vol. 75.
- [5] T. Durrani, A. Rauf, K. Boyle, F. Lotti, and S. Baronti, "Thermal imaging techniques for the non destructive inspection of composite materials in real-time," in *ICASSP '87. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 12, Apr 1987, pp. 598–601.
- [6] N. Pech-May, A. Oleaga, A. Mendioroz, and A. Salazar, "Fast characterization of the width of vertical cracks using pulsed laser spot infrared thermography," vol. 35, 06 2016.
- [7] S. Burrows, S. Dixon, S. Pickering, T. Li, and D. P. Almond, "Thermographic detection of surface breaking defects using a scanning laser source," vol. 44, pp. 589–596, 11 2011.
- [8] 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.
- [9] W. Parker, R. Jenkins, C. Butler, and G. Abbott, "Flash method of determining thermal diffusivity, heat capacity, and thermal conductivity," *Journal of applied physics*, vol. 32, no. 9, pp. 1679–1684, 1961.
- [10] J.-C. Krapez, L. Spagnolo, M. Frie, H.-P. Maier, and G. Neuer, "Measurement of in-plane diffusivity in non-homogeneous slabs by applying flash thermography," *International Journal of Thermal Sciences*, vol. 43, no. 10, pp. 967 – 977, 2004. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S129007290400047X>
- [11] F. Lakestani, A. Salerno, and A. Volcan, "Modulated spot heating for the measurement of thermal diffusivity," *Journal of Applied Physics*, vol. 97, no. 1, p. 013704, 2005. [Online]. Available: <https://doi.org/10.1063/1.1828217>
- [12] H. Dong, B. Zheng, and F. Chen, "Infrared sequence transformation technique for in situ measurement of thermal diffusivity and monitoring of thermal diffusion," *Infrared Physics & Technology*, vol. 73, pp. 130 – 140, 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1350449515002443>
- [13] X. Maldague, "Theory and practice of infrared technology for nondestructive testing," 2001.
- [14] H. Halloua, A. Obadi, Y. Errami, S. Sahnoun, and A. Elhassnaoui, "Neural networks for back wall geometry reconstruction using the active thermography," in *Electrical Sciences and Technologies in Maghreb (CISTEM), 2016 International Conference on*. IEEE, 2016, pp. 1–7.
- [15] P. Baranowski, W. Mazurek, J. Wozniak, and U. Majewska, "Detection of early bruises in apples using hyperspectral data and thermal imaging," *Journal of Food Engineering*, vol. 110, no. 3, pp. 345–355, 2012.
- [16] 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*. IEEE, 2012, pp. 52–57.
- [17] M. Tanaka, J. L. L  v  que, 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.
- [18] 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.
- [19] 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 (RSS), Rome, Italy, 2015*.
- [20] 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*. IEEE, 2016, pp. 297–304.
- [21] 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*. IEEE, 2008, pp. 3977–3982.
- [22] 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*. IEEE, 2013, pp. 3056–3061.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] 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. IEEE, 1984, pp. 332–340.
- [27] 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*. IEEE, 2010, pp. 239–246.
- [28] 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.
- [29] 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.
- [30] S. Bell, P. Upchurch, N. Snaveley, 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*, 2015, pp. 3479–3487.
- [31] Y. A. Cengel and A. Ghajar, "Heat and mass transfer (a practical approach, si version)," 2011.
- [32] W. A. Strauss, *Partial differential equations*. Wiley New York, 1992, vol. 92.
- [33] C. Meola, *Infrared thermography recent advances and future trends*. Oak Park, Ill.: Bentham Science, 2012.
- [34] E. G. Henneke, K. L. Reifsnider, and W. W. Stinchcomb, "Thermographic ndi method for damage detection," *JOM*, vol. 31, no. 9, pp. 11–15, 1979.
- [35] M. Planck, "On the law of distribution of energy in the normal spectrum," *Annalen der Physik*, vol. 4, no. 553, p. 1, 1901.
- [36] M. Q. Brewster, *Thermal radiative transfer and properties*. John Wiley & Sons, 1992.
- [37] I. OMEGA Engineering, *FLIR Thermal Imager Sensor Manual*, 2012. [Online]. Available: <https://www.omega.com/manuals/manualpdf/M5230.pdf>
- [38] CES EduPack 2017, Granta Design Limited, UK (www.grantadesign.com).
- [39] C. Ibarra-Castaneda, F. Galmiche, A. Darabi, M. Pilla, M. Klein, A. Ziadi, S. Vallerand, J.-F. Pelletier, and X. P. Maldague, "Thermographic nondestructive evaluation: overview of recent progress," in *Thermosense XXV*, vol. 5073. International Society for Optics and Photonics, 2003, pp. 450–460.
- [40] J. A. J. A. Richards, *Remote sensing digital image analysis : an introduction*, 3rd ed. Berlin ; New York: Springer, 1999.
- [41] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [42] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [43] A. Paszke, S. Gross, S. Chintala, and G. Chanan, "Pytorch," 2017.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [45] S. Hochreiter and J. Schmidhuber, "Lstm can solve hard long time lag problems," in *Advances in neural information processing systems*, 1997, pp. 473–479.
- [46] S. Ioffe and C. Szegedy, "Batch normalization: accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*. JMLR. org, 2015, pp. 448–456.