

Material Classification with Thermal Imagery

Philip Saponaro, Scott Sorensen, Abhishek Kolagunda, and Chandra Kambhamettu
Video/Image Modeling and Synthesis (VIMS) Lab,
Dept. of Computer and Information Science, University of Delaware
Newark, Delaware

{saponaro, sorensen, abhi, chandrak}@udel.edu

Abstract

Material classification is an important area of research in computer vision. Typical algorithms use color and texture information for classification, but there are problems due to varying lighting conditions and diversity of colors in a single material class. In this work we study the use of long wave infrared (i.e. thermal) imagery for material classification. Thermal imagery has the benefit of relative invariance to color changes, invariance to lighting conditions, and can even work in the dark. We collect a database of 21 different material classes with both color and thermal imagery. We develop a set of features that describe water permeation and heating/cooling properties, and test several variations on these methods to obtain our final classifier. The results show that the proposed method outperforms typical color and texture features, and when combined with color information, the results are improved further.

1. Introduction

Automatically classifying materials impacts real world applications such as recycling [9], mineralogy, and robotics [11]. Computer vision methods typically use standard RGB camera imagery and rely on texture and lighting cues to distinguish the different materials. Appearance based classification is challenging due to the variety of colors and illuminations. Objects of different materials that have the same color intensity response can be almost impossible to distinguish. In this work we consider the electromagnetic spectrum between 8-14 μ m with a long wave infrared (LWIR) camera, i.e. a thermal camera. We study thermal properties and how water permeates through different materials.

LWIR cameras detect infrared radiation, which is emitted by all objects above absolute zero according to Planck's black body radiation law [17]. The LWIR camera sees not only the emitted thermal radiation of a source object, but also reflected and transmitted thermal radiation. Emissivity is a property of a material which governs the effective-

ness of emitting thermal radiation compared to a black body source at the same temperature. Metals have very low emissivity and high reflectance, while woods have high emissivity. Thus many metals appear mirrorlike in the LWIR imagery.

When water comes into contact with an object, the permeation behavior changes based on the material. For example, in wood the water follows along the grain of the wood and is jagged in appearance, while in metals the water stays above the surface and hardly moves if the surface is flat, and in paper materials, the water spreads radially. We use the permeation behavior as a cue for classification. We construct a 3D model of the water permeation pattern for each class of material. Each frame of a video sequence of the water permeation is treated as a 2D slice of a 3D model. This model generalizes features of the permeation pattern such as rate of permeation and shape characteristics. Another cue we use is the thermal heating and cooling cycle. We show that different materials heat and cool at different rates when placed under a heating lamp. The materials we used for classification are shown in Figure 1.

The paper is organized as follows. Section 2 gives previous works for material classification. Section 3 details our method of feature extraction and model learning. Section 4 shows the data we collected and discusses the results of our classifier. Section 5 concludes the paper and discusses future work.

2. Background and Related Works

Many previous works consider material classification with a standard RGB camera. Color and texture information were extracted in [7]. [8] uses a Bidirectional Reflectance Distribution Function (BRDF) as a feature for per-pixel classification. [13] uses a Bidirectional Texture Function (BTF). Both use coded light illumination in an LED dome. Light polarization was used in [4, 21]. BRDF slices were used in [20]. Visible spectral reflectance (400-720nm) was used in [10]. In most of these works, the lighting was very controlled to give cues for the classification. In our

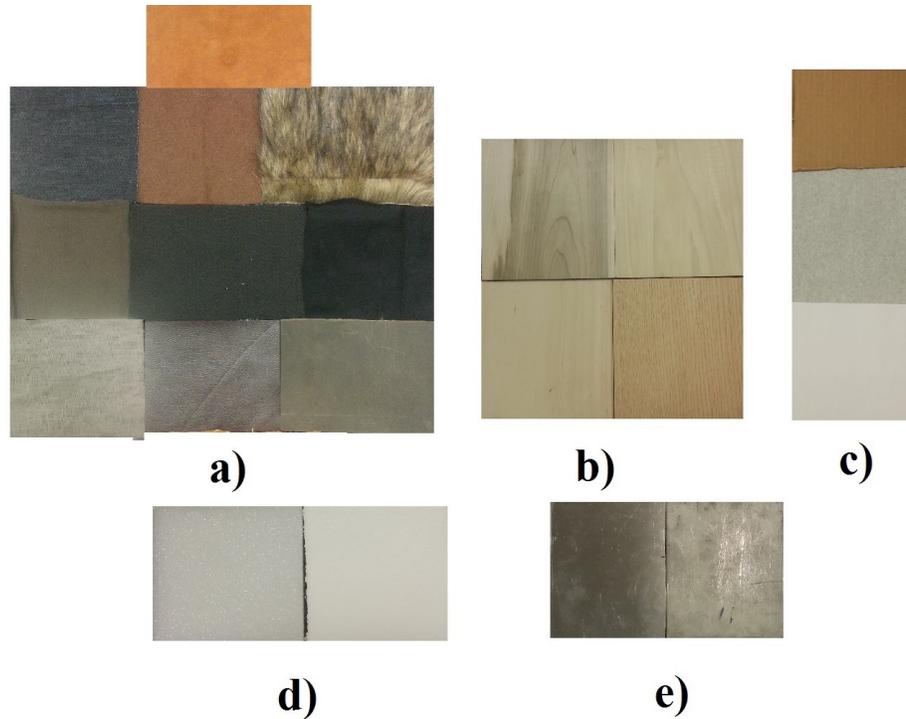


Figure 1: Materials used for classification. **a)** Cloth: suede, denim, wool, synthetic fur, cloth, felt, polyester, linen, synthetic leather, real leather. **b)** Wood: maple, poplar, birch, oak. **c)** Paper: corrugated cardboard, paper towel, printing paper. **d)** Plastic Foams: closed-cell expanded polystyrene, closed-cell extruded polystyrene. **e)** Metals: aluminum, steel.

work, the lighting is less important compared to thermal properties for classification.

Thermal imagery has not been heavily studied for material classification. [18] uses near infrared (NIR) to get a more intrinsic image of the material sample. [15] uses mid wave infrared for paper and board identification for food packaging. [12] measures thermal conductivity of materials using tactile feedback; a robot touches the material with a probe. LWIR can easily detect water damage in buildings and can detect permeation in materials [22]. To the best of our knowledge, this is the first work to use long wave infrared cameras to obtain image-based thermal properties for material classification.

3. Methods

An overview of our methods can be seen in Figure 2. Water permeation behavior is extracted by computing a characteristic model of permeation as discussed in Section 3.1, and is further transformed into the FFT of a binned spherical map for comparison. We developed two heating/cooling features. The first is extracted by sampling small patches over time to obtain a temperature curve as discussed in Section 3.2.1. The second is extracted by solving the heat equation for unknown constant parameters in Section 3.2.2.

3.1. CHAMP - CHAracteristic Model of Permeation

There are a few interesting features that can be obtained for water permeation such as the rate of permeation of the water into the material and the shape characteristics of the permeation. We generalize these features by creating a 3D model which we call the CHAMP (CHAracteristic Model of Permeation). To compute the CHAMP, we first define a 2D indicator function f for a material as

$$f_i(X, Y) = I_i(x, y) < \tau_f, \quad (1)$$

where τ_f is a material specific threshold value and i is the i th frame in the video sequence. The material threshold value can be calculated using Otsu's method [2], assuming a bimodal distribution of intensity values. The indicator function gives a value of 1 for points inside or on the CHAMP, and 0 for points outside the CHAMP. It is possible to use more complicated methods for calculating the indicator function, such as active contour models [3], but for most of our data the difference in temperature between the water and background is bimodal and significant enough that a simple threshold suffices. Moreover, the active contour model would lose some finer details depending on the snaxel resolution.

Once the indicator function is calculated, the boundary of the model can be quickly estimated using morphological

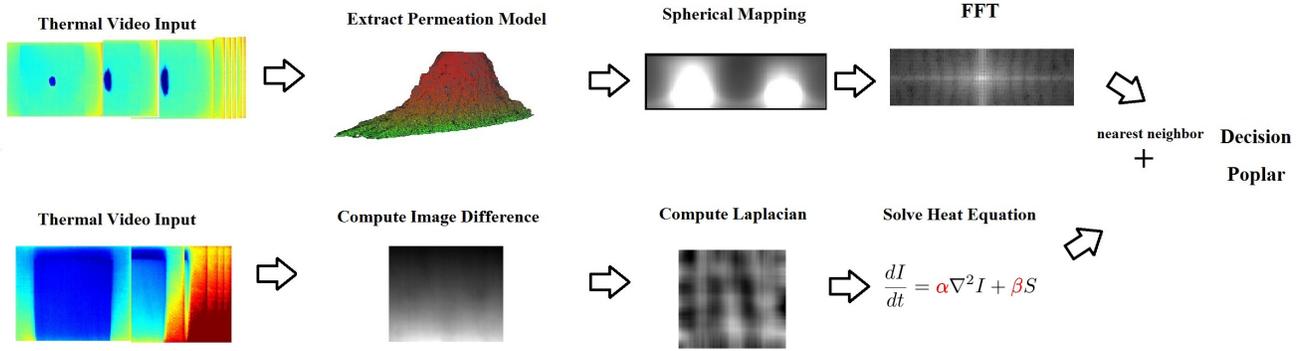


Figure 2: The proposed method consists of two types of features – water permeation and a heating/cooling cycle. For water permeation, we extract a 3D model called “CHAMP” which describes the water permeation rate and size. For heating and cooling, we solve a variation of the heat equation [5] for the constant parameters. Note that the permeation model is a mesh for display purposes, and the color refers to time (red = start). The FFT image is in a log-2 scale.

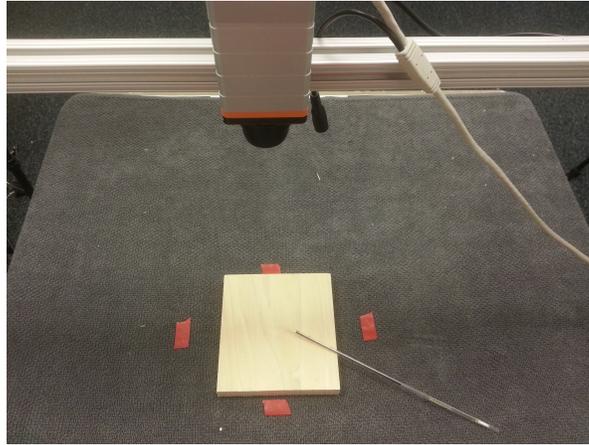


Figure 3: Camera setup for water permeation experiment. All materials were roughly centered in the image, and a pipette was used to control the amount of water.

operations [6, 14] as in

$$\Omega_i(X, Y) = f_i(X, Y) - f'_i(X, Y), \quad (2)$$

where $f'_i(X, Y)$ is obtained by eroding f using a small structuring element. To create the 3D model, simply concatenate the boundary $\Omega_i(X, Y)$ for each slice i along the Z dimension. The model can be “capped” by using the indicator function for $i = 1$ and $i = N$ for an N frame video sequence, instead of the boundary $\Omega_i(X, Y)$. That is

$$CHAMP(X, Y, i) = \begin{cases} f_i(X, Y), & \text{if } i = 1, N. \\ \Omega_i(X, Y), & \text{otherwise.} \end{cases}, \quad (3)$$

The CHAMP can be understood by visualizing each image in a video sequence as a 2D slice along the Z dimension. It implicitly contains the shape of the water permeation as well as the growth rate (i.e. curvature of the model).

3.1.1 Comparing CHAMPs with Binned Spherical Mapping

Once the models are created, a metric to compare the models is needed. We chose to compute an Fast Fourier Transform (FFT [19]) of a modified binned spherical mapping of the CHAMP. This is a robust method since it is invariant to rotations and translations between CHAMPs and is fast to compute.

A spherical coordinate system is a coordinate system where points are represented by 3 parameters: the radial distance from the origin, the polar angle measured in the zenith direction, and the azimuth angle orthogonal to the zenith direction. Before performing this mapping, we first center the CHAMP around its centroid. Next, we map (X, Y, Z) in the Cartesian coordinate system to (r, θ, ϕ) in the spherical coordinate system using simple trigonometric equations.

These spherical points are binned into a 2D histogram

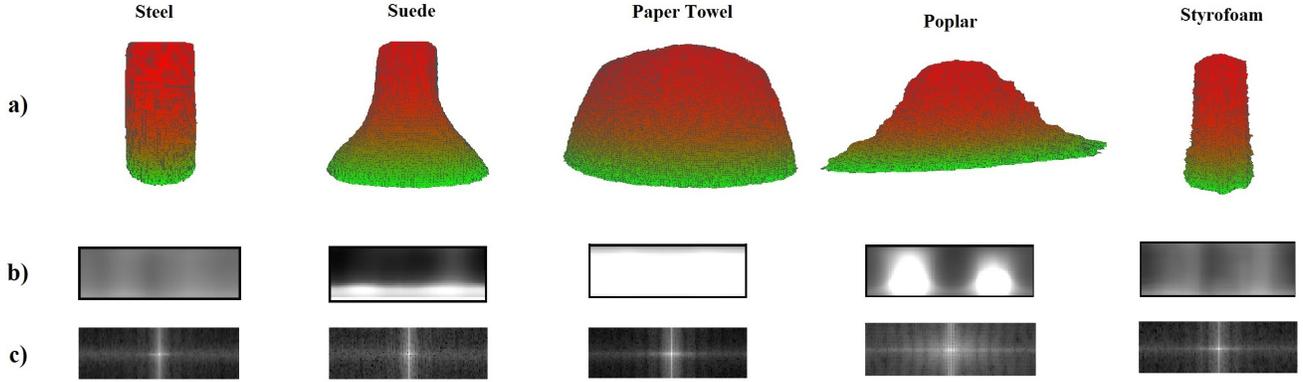


Figure 4: Comparison of select few materials from different coarse-grain classes. **a)** CHAMP. **b)** Binned spherical map. **c)** FFT of binned spherical map.

image. The intensity values of the histogram image are the r values multiplied by $\cos(\phi_k)$, the rows are varying θ , and the columns are varying ϕ . This is performed by

$$SPH(x, y) = \cos(\phi_k) * avg(r_k) \quad (4)$$

$$\{k \mid \theta_k \pm \epsilon = x * bin_x - \pi, \\ \phi_k \pm \epsilon = y * bin_y - \frac{\pi}{4}\},$$

where bin_x and bin_y are the desired bin size. Each row corresponds to a slice of the model, and the values are the distances from the centroid of that slice to the edge of the model.

Next, we compute the FFT of these 2D histogram images and shift the zero-frequency component to the center. Since the only misalignment of the spherical maps will be in the horizontal direction, we can ignore the phase and take the amplitude of the FFT image. This allows our FFT images to be aligned even if the CHAMPs are misaligned due to rotations and translations. For display purposes in Figure 4, we take the base-2 log of the FFT image. These FFT images are compared to each other using correlation, where a higher value corresponds to a better match. Figure 4 shows a comparison of the various models between a select few of the materials.

3.2. Material Heating and Cooling

In this section, we describe two features we extracted to represent the heating and cooling of the materials. The first feature we attempted was very simple, while the second one yielded higher accuracy and is more physically meaningful. These two features are described below.

3.2.1 Patched-Based Temperature Curves

Our first attempt of feature extraction of heating and cooling is quite simple. For each image in an infrared video

stream, we sample five patches as shown in Figure 5. For each patch, the mean temperature over the patch is plotted over time to give a temperature curve that should, ideally, be unique for each material. Patches were chosen over using all pixels to smooth over noise and to speed up the processing time.

To account for change in room temperature over the course of a day, we align the starting temperatures of each curve when comparing across materials, i.e. for materials x and y perform $T'_{xk}(t) = T_{xk}(t) + [T_{yk}(1) - T_{xk}(1)]$. Here $T_{xk}(t)$ is the temperature at time t for material x at patch k . Euclidean distance is used as a metric for comparison.

3.2.2 Solving the Heat Equation

The heat equation [5] is a parabolic partial differential equation that describes the distribution of heat over time. We augment the standard heat equation to more closely describe our physical setup by adding a second term as in

$$\frac{dI}{dt} = \alpha \nabla^2 I + \beta S(t), \quad (5)$$

where α, β are unknown constants, and S is a function which describes how heat is applied. In our setup, a heat lamp was the source of heat in the scene, and its temperature changed over time. To calculate S , we sampled the temperature of the heat lamp over time using an infrared thermometer, and fit a piecewise polynomial to the sample temperatures. Once S is known, we can calculate α, β by setting up an overconstrained linear system and applying a Moore-Penrose pseudoinverse [1]. The system is set up as

$$\left(\begin{array}{c} \nabla^2 I \\ S(t) \end{array} \right) \left(\begin{array}{c} \alpha \\ \beta \end{array} \right) = \left(\begin{array}{c} \frac{dI}{dt} \end{array} \right). \quad (6)$$

The resulting parameters are the feature vector for comparison and Euclidean distance was used as a difference metric.

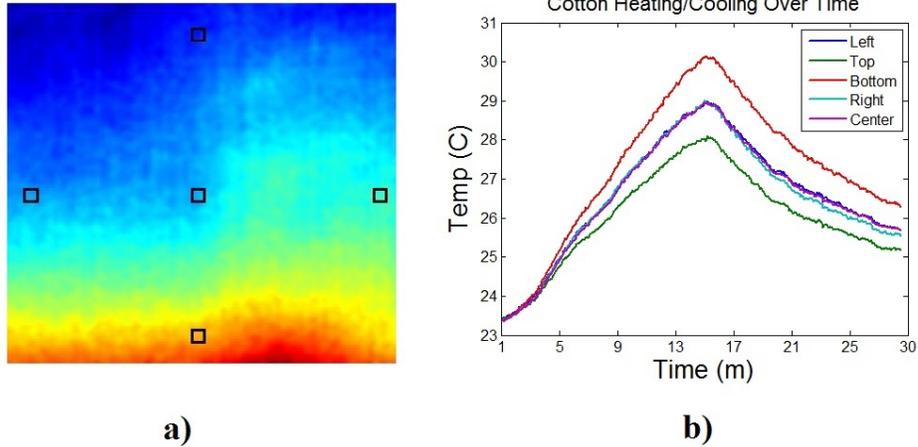


Figure 5: Simple heating/cooling feature. A heat lamp was placed at the bottom of each material and was turned on at $t=1m$ and turned off at $t=16m$. **a)** Location of sampled patches in infrared image for heating/cooling feature extraction. **b)** Corresponding graph of patch temperature over time.

Method	Rank 1 Coarse	Rank 2 Coarse	Rank 1 Fine	Rank 2 Fine	Rank 3 Fine
Capped, FFT	79.8	92.8	59.5	70.2	82.1
Uncapped, FFT	77.4	97.7	57.1	73.8	86.9
Capped, SPH	83.3	91.7	63.1	81.0	83.3
Uncapped, SPH	82.1	89.3	61.9	76.2	85.7
Capped, FFTCOS	83.3	92.9	64.3	73.8	82.1
Uncapped, FFTCOS	84.5	92.9	64.3	73.8	83.2
Capped, SPHCOS	82.1	89.3	61.9	79.8	85.7
Uncapped, SPHCOS	80.1	89.3	60.7	76.2	85.6
Heat Equation 100x100, 50	78.6	85.7	42.9	57.1	69.1
Heat Equation 100x100, 25	66.7	76.6	35.7	50.0	61.5
Heat Equation 100x100, 75	69.1	88.1	28.6	40.5	42.9
Heat Equation 50x50, 15	64.3	76.2	33.3	45.2	57.1
Heat Equation 150x150, 85	54.8	83.3	35.7	40.5	45.5
Heating/Cooling Graphs 3x3	57.1	85.7	35.7	52.4	63.1
Heating/Cooling Graphs 9x9	59.5	88.1	35.7	52.4	63.1
Heating/Cooling Graphs 15x15	59.5	83.1	28.6	52.4	59.2
Combination	95.3	100.0	71.4	85.7	92.9

Table 1: Results using variations on our proposed features given as accuracy over entire dataset. Rank n means the correct class was in the top n choices. Coarse refers to wood vs metal vs cloth vs paper vs plastic foams, whereas fine refers to a specific class (e.g. poplar) against all other 20 classes.

4. Experiments and Results

In our experiments we used a Xenics Gobi 640 GigE uncooled long wave infrared camera, which has a resolution of 640×480 and has a $50mC$ sensitivity at $30^\circ C$. The materials we used were broken up into 5 coarse classes: cloth, wood, paper, plastic foams, and metal. Each coarse class was further broken up into a total of 21 subclasses as shown in Figure 1. For each type of material, we imaged 4 samples; this

gives a total of 84 material samples. The physical setup of the camera and materials is shown in Figure 3. The LWIR camera was $0.4m$ above the materials looking downwards. The boundaries of the image were marked so that materials can be roughly aligned to the center of each image.

For the water permeation experiment, we used a pipette to drop $0.4mL$ of water onto the center of the material. We recorded a ten minute long video at 1fps for each sample. For each video, we extracted the CHAMPs and FFT

Method	Rank 1 Coarse	Rank 2 Coarse	Rank 1 Fine	Rank 2 Fine	Rank 3 Fine
Ours + DCT	98.8	100.0	81.0	91.7	94.0
Ours	95.2	100.0	71.4	85.7	92.9
NIR [18]	92.9	96.4	76.2	88.1	92.9
HSL [16]	83.3	95.3	59.5	85.7	92.9
DCT [7]	92.9	95.2	61.9	80.9	83.3
Gabor [7]	82.1	90.5	61.9	79.8	81.0
Co-occurrence [7]	81.0	81.0	60.0	64.8	67.1

Table 2: Results of comparison to other works using our dataset. Rank n means the correct class was in the top n choices. Coarse refers to wood vs metal vs cloth vs paper vs plastic foams, whereas fine refers to a specific class (e.g. poplar) against all other 20 classes.

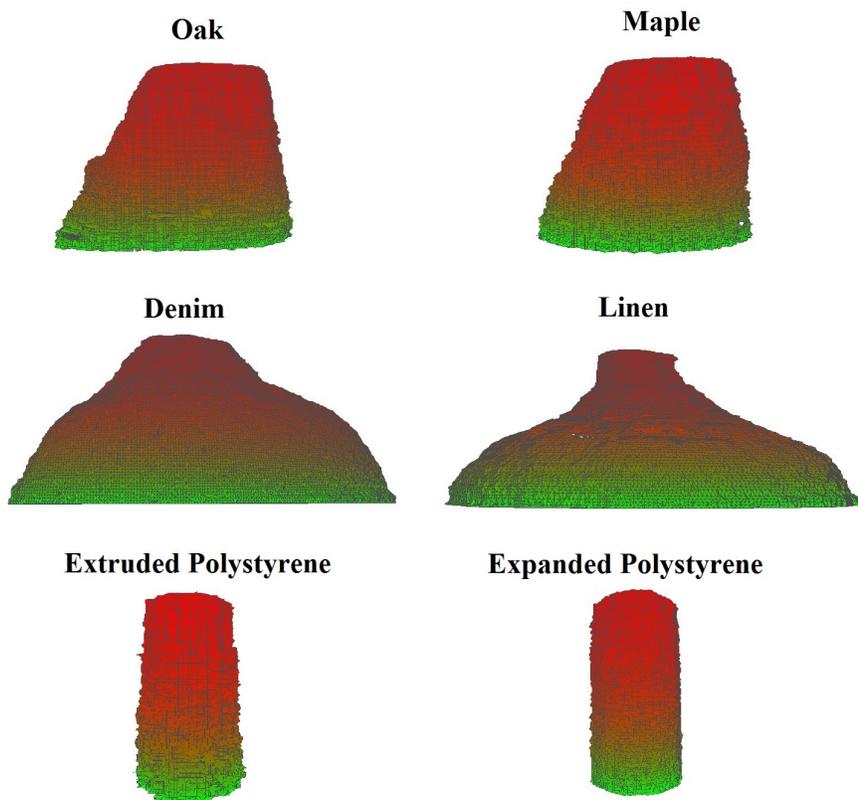


Figure 6: Misclassified CHAMPs of materials using only water permeation. When adding in heating/cooling information, these materials are distinguishable.

of binned spherical maps as described in Section 3.1. For each fine-grain class type, we created a mean model by averaging the FFT images for all samples. We used leave-one-out cross validation in this procedure by leaving one sample out for testing, and three samples to create the mean model. The mean model is compared against all other material samples. We tested a few variations including: “capping” the CHAMP, leaving off the caps, using aligned spherical maps without FFT, using the FFT, and using aligned spherical maps without the $\cos(\phi)$ term in Eq. 4. When uncapped,

ϕ is restricted to $-\frac{\pi}{4}$ to $\frac{\pi}{4}$ to avoid NaN results. These results are reported in Table 2. We give accuracy using rank1, rank 2, and rank 3 results, where rank n means if our choice was in the top n choices it is marked as correct. This is a useful metric to see how much each feature can narrow down the possible choices.

For the heating and cooling experiment, we recorded 30 minutes videos of the heating and cooling of each material sample. Each sample was placed in the center of the camera’s view and heated with a heat lamp placed in front and

above the material. The heat lamp was turned on for 15 minutes, and then switched off at the 15 minute mark. This process was automated using a programmable Arduino microcontroller to ensure precise timing. We tested a few different sized Laplacian of Gaussian (LoG) filter sizes and sigma values for the feature described in Section 3.2.2. We also tested a few different patch sizes for the feature described in Section 3.2.1. The results are included in Table 2.

To combine the features, we first took the top $n = 5$ choices of the best water permeation feature. Then, we used a simple linear combination of the normalized distances of the best water permeation variation and best heating/cooling patch size as in $\alpha_w * \text{UncappedFFT} + (1 - \alpha_w) * (\text{HeatEquation} [100 \ 100], 50)$. The weighting parameter α_w and n chosen were the ones with the highest average accuracy scores across all categories.

We compare to other material classification works that use color, texture, and near infrared. To the best of our knowledge there is no other work on image-based thermal properties for material classification. [7] gives a comparative study of approaches for classification of color texture images. We implemented the three features they recommended across all three color channels – Discrete Cosine Transform (DCT), Gabor Filters, and Co-occurrence. [18] uses NIR for classification due to its relative independence from color imagery. We recorded our materials with the Sony EVI-D70, which is a color video camera with a NIR mode and a resolution of 640x480.

We combined our infrared results with color results to further improve accuracy. The way we do this is similar as described above – a linear combination of the normalized distances of the best infrared features and best color features as in $1 - \beta_w * \text{Combination} + (\beta_w) * \text{DCT}$. In our experiments, $n = 5$, $\alpha_w = 0.14$, and $\beta_w = 0.25$. The smaller β_w is, the less color and texture information is relied on for classification.

4.1. Discussion

When water permeation is taken as a feature alone, the best rank 1 accuracy over the 5 coarse classes is 84.5% when using the FFT of the spherical map multiplied by $\cos(\phi)$. This is also the best version of the feature when looking at the best rank 1 accuracy over all 21 classes at 64.3%. Water permeation may be better utilized when combined with other features. This is because it is able to narrow down the possible classes very effectively. Using an uncapped CHAMP with the FFT of the spherical map we can achieve 97.7% rank 2 accuracy for coarse classes; this means the correct material class is almost always in the top 2 choices. Similarly for fine classes, this variation maximizes the accuracy where the correct class is in the top 3 choices about 87% of the time. That is why UncappedFFT

was chosen for the combination – it was able to narrow down the possible choices the best. Similarly, using the heat equation with a LoG filter of size 100x100 and $\sigma = 50$ led to the highest results across the categories. When combined together, the results are improved up to 16%, which implies these features are complimentary to each other.

When comparing to other works, our method gives the best result across all categories, although in some cases the gain is only slight. However, when color is added to our method using the linear combination described above, the results are significantly improved up to 10%. Moreover, weighting parameter β_w can be controlled to give more invariance to color and texture information, depending on the dataset being used.

One drawback of our method is the amount of time it takes to record the water permeation and heating/cooling videos, which were 10 minute and 30 minute respectively. The length of videos recorded for this project were conservative, and it may be possible to decrease the video length. Also, the heating/cooling can be sped up by placing the heating element closer to the material.

5. Conclusions and Future Work

In this paper, we described features that can be extracted from thermal imagery which give results that not only outperform other color and texture features, but also are complimentary to them, and can be combined to increase performance. We collected a dataset of 21 different classes with 84 total samples, and recorded thermal video of water permeation and heating/cooling, as well as color and NIR photographs. We present CHAMPs, which model the water permeation, and a method to extract the heat equation constants. Combining these features together results in higher accuracy than using either one individually. Our future work will focus on developing features that are faster to record, as well as testing for color invariance by collecting material classes that have the same color as each other.

6. Acknowledgements

This work is funded by NSF CDI Type I grant 1124664 and by Cooperative Agreement W911NF-11-2-0046 (ARO Proposal No. 59537-EL-PIR).

References

- [1] *Generalized Inverses of Linear Transformations: The MoorePenrose or Generalized Inverse*, chapter 2, pages 8–27. 4
- [2] A threshold selection method from gray-level histograms. *Systems, Man and Cybernetics, IEEE Transactions on*, 9(1):62–66, Jan 1979. 2

- [3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pages 694–699, Jun 1995. [2](#)
- [4] H. Chen and L. Wolff. Polarization phase-based method for material classification and object recognition in computer vision. In *Computer Vision and Pattern Recognition, 1996. Proceedings CVPR '96, 1996 IEEE Computer Society Conference on*, pages 128–135, Jun 1996. [1](#)
- [5] J. Crank and P. Nicolson. A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type. *Advances in Computational Mathematics*, 6(1):207–226, 1996. [3](#), [4](#)
- [6] E. Dougherty. *An introduction to morphological image processing*. Tutorial texts in optical engineering. SPIE Optical Engineering Press, 1992. [3](#)
- [7] A. Drimbarean and P. Whelan. Experiments in colour texture analysis. *Pattern Recognition Letters*, 22(10):1161 – 1167, 2001. [1](#), [6](#), [7](#)
- [8] J. Gu and C. Liu. Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral brdf. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 797–804, June 2012. [1](#)
- [9] K. A. Hendrik, R. Massey, B. B. Whitham, and M. D. Russell. Technologies for the identification, separation, and recycling of automotive plastics. In *International Journal of Environmentally Conscious Design and Manufacturing*, page 3750, 1997. [1](#)
- [10] A. Ibrahim, S. Tominaga, and T. Horiuchi. Spectral imaging method for material classification and inspection of printed circuit boards. *Optical Engineering*, 49(5):057201–057201–10, 2010. [1](#)
- [11] N. Jamali and C. Sammut. Material classification by tactile sensing using surface textures. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 2336–2341, May 2010. [1](#)
- [12] E. Kerr, T. McGinnity, and S. Coleman. Material classification based on thermal properties x2014; a robot and human evaluation. In *Robotics and Biomimetics (ROBIO), 2013 IEEE International Conference on*, pages 1048–1053, Dec 2013. [2](#)
- [13] C. Liu, G. Yang, and J. Gu. Learning discriminative illumination and filters for raw material classification with optimal projections of bidirectional texture functions. [1](#)
- [14] M. Mirmehdi and M. Petrou. Segmentation of color textures. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(2):142–159, Feb 2000. [3](#)
- [15] J.-R. Riba, T. Canals, and R. Cantero. Recovered paperboard samples identification by means of mid-infrared sensors. *Sensors Journal, IEEE*, 13(7):2763–2770, July 2013. [2](#)
- [16] Z. Ronghua, C. Hongwu, Z. Xiaoting, P. Ruru, and L. Jihong. Unsupervised color classification for yarn-dyed fabric based on fcm algorithm. In *Artificial Intelligence and Computational Intelligence (AICI), 2010 International Conference on*, volume 1, pages 497–501, Oct 2010. [6](#)
- [17] G. Rybicki and A. Lightman. *Radiative Processes in Astrophysics*. Physics textbook. Wiley, 2008. [1](#)
- [18] N. Salamati, C. Fredembach, and S. Ssstrunk. Material classification using color and nir images. [2](#), [6](#), [7](#)
- [19] C. Van Loan. *Computational Frameworks for the Fast Fourier Transform*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1992. [3](#)
- [20] O. Wang, P. Gunawardane, S. Scher, and J. Davis. Material classification using brdf slices. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2805–2811, June 2009. [1](#)
- [21] L. Wolff. Polarization-based material classification from specular reflection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(11):1059–1071, Nov 1990. [1](#)
- [22] L. Wu, S. Liu, and Y. Wu. Infrared imaging detection of water permeation on field large-scale rock relics. *Geoscience and Remote Sensing, IEEE Transactions on*, 49(1):581–590, Jan 2011. [2](#)