

Thermal Spread Functions (TSF): Physics-guided Material Classification

Aniket Dashpute^{1,3}, Vishwanath Saragadam³, Emma Alexander²,
 Florian Willomitzer⁴, Aggelos Katsaggelos¹, Ashok Veeraraghavan³, Oliver Cossairt^{1,2}

¹Electrical and Computer Engineering, ²Computer Science, Northwestern University,

³Electrical and Computer Engineering, Rice University,

⁴Wyant College of Optical Sciences, University of Arizona

Abstract

Robust and non-destructive material classification is a challenging but crucial first-step in numerous vision applications. We propose a physics-guided material classification framework that relies on thermal properties of the object. Our key observation is that the rate of heating and cooling of an object depends on the unique intrinsic properties of the material, namely the emissivity and diffusivity. We leverage this observation by gently heating the objects in the scene with a low-power laser for a fixed duration and then turning it off, while a thermal camera captures measurements during the heating and cooling process. We then take this spatial and temporal “thermal spread function” (TSF) to solve an inverse heat equation using the finite-differences approach, resulting in a spatially varying estimate of diffusivity and emissivity. These tuples are then used to train a classifier that produces a fine-grained material label at each spatial pixel. Our approach is extremely simple requiring only a small light source (low power laser) and a thermal camera, and produces robust classification results with 86% accuracy over 16 classes¹.

1. Introduction

Material classification is an important task pertinent to a diverse set of fields including but not limited to medicine and biology [1], chip manufacturing, recycling [2, 3], land and weather monitoring using satellites, and vision and robotics. Robust material classification is particularly critical in separating various parts of an object based on their constituent materials [2,3]. Common tools for material classification span a large spectrum including simple tools such as infrared spectroscopy, hyperspectral imaging to more exotic tools such as ultrasound, and x-ray fluorescent imagers.

Material classification primarily relies on various dimensions of light including bidirectional reflectance function

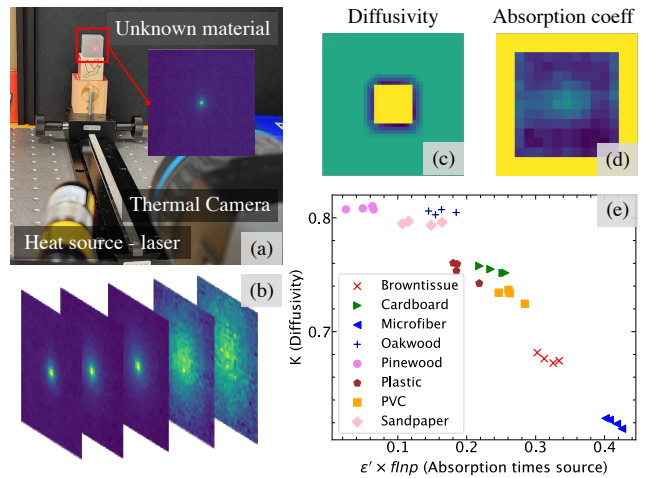


Figure 1. **Material classification with thermal properties.** Materials have unique thermodynamics that enables robust classification. We propose a simple setup composed of (a) a 60mW laser as a heat source and a thermal camera to capture the heat profile (inset). This results in a (b) stack of images we call Thermal Spread Function (TSF) that encodes heating and cooling effect around the laser dot. The TSF data is used to (c) estimate diffusivity map and (d) the external heat source term using inverse Finite Difference Method, which is then used to (e) classify the material robustly.

(BRDF) slices [4], color and NIR images [5], frequency and depth-dependent ToF distortion [6], spectral imaging methods [7,8], multi-modal methods [9], and thermal imaging [10]. Methods based on RGB images are popular due to availability of RGB cameras and large labeled datasets, but suffer from lack of robustness. In contrast, spectrum-based imaging based methods enable accurate classification but often require complex optical systems such as hyperspectral cameras, and are sensitive to external illumination conditions.

Human perception of materials is often multi-modal, such as relying on touch and vision to accurately classify material. This act of touching further involves a thermodynamic exchange that relies on the material composition of

¹Code: <https://github.com/aniketdashpute/TSF>

the object. A metallic object results in rapid conduction of heat whereas non-metallic objects like ones made of wood result in slower transfer rate. Thus the intrinsic thermal properties provides an insight into the material properties that is often missed or confused by vision alone. Previous work in contact ways of knowing conductivity - haptic sensing [11, 12], and haptic displays [13, 14] leveraged this idea by developing “an artificial fingertip”. The drawback of this method is that it is invasive - it requires to touch the scene and thus can lead to interfering with it. Thermal characterization for recycling has also been done using a spectrometer and a fluxmeter [15].

Thermal Imaging methods enable contact-free estimation of thermal properties, thus allowing us to classify materials rapidly and in a non-destructive manner. One of the most popular contact-less methods for determining thermal diffusivity is the laser flash method. A laser is flashed on a thin slice (microns thick) of a material and the temperature change is observed from the other side, providing a quantitative estimate of the thermal diffusivity or conductivity [16, 17]. This is restrictive due to the constrained lab setup and requirement of thin slices. Thermal imaging has also been used for non-destructive infrastructure inspection where the difference in thermal behaviour of unaltered and defected zones allow defect detection [18].

We take inspiration from the contact-less methods and develop a non-invasive thermal imaging system for material classification. As opposed to previous methods, our method is robust enough to be used in uncontrolled environments, and not limited to constrained lab setups. We use a visible laser beam as an external heat source that shines on a material, which absorbs a fraction of this beam corresponding to optical wavelength. The absorption of this energy leads to a rise in temperature that shows up in the long wave infrared domain (LWIR) and is captured by the thermal camera. The thermal camera is used to capture the heating process, and once the heat source is off, its cooling (refer Fig. 1). We define the temperature transients obtained from the heating-cooling as its **Thermal Spread Function (TSF)** and use it for robustly classifying materials.

A key challenge with using TSF for classifying materials is that a thermal camera requires a known emissivity (ϵ) (ratio of radiated energy of the object to that of a black-body) for accurately estimating the temperature. To overcome this ambiguity, we leverage a physically accurate heat diffusion equation (see Sec. 2) that carefully models the thermodynamic interactions between the ambient scene and the object. This estimated TSF is then used for training a material classifier which enables robust material classification.

Our approach and main contributions

When objects are heated through radiation on surface and allowed to cool down, they display characteristic tem-




Figure 2. **Heating and capturing process.** We use an external heat source to heat the object and a thermal camera to observe the heating and cooling effect. Refer Sec. 2 for detailed explanation.

perature changes. These changes are based on their initial temperature, surface absorption, and heat diffusivity. We inject heat through a small portion on the surface of a material which diffuses throughout the body over time. If we observe a small patch of material in the vicinity of injection, we observe the diffusion - both during the injection phase and during the cooling phase after no external heat is supplied. We call this varying 2D temperature profile as the Thermal Spread Function (TSF) of the material.

We measure the TSF of the material through a Long Wave Infrared (LWIR) thermal camera. We derive diffusivity and an absorption factor from the TSF to characterize the material as these properties are independent of the initial temperature of the object. Our main contributions are the following.

- We first derive a physically accurate model that characterizes the Thermal Spread Functions (TSFs) as a function of initial temperature of the object and an object’s thermodynamic properties.
- We then use a Finite Differences (FD) Method to solve the inverse heat problem for recovering parameters related to diffusion, absorption and emission
- Finally, we design and demonstrate a simple optical setup for non-invasively recovering the thermodynamic properties and using them to classify materials.

2. Background

We discuss the relevant heat equations and their physical significance below. For more detailed steps and explanation, please refer to [19, 20].

Heat Diffusion

Consider the scene shown in Fig. 2 where an object is illuminated by an external heat source. Let e denote the thermal energy density, ϕ represent the heat flux (energy per unit time per unit surface area) across the object boundary and F be the heat absorbed per unit time per unit volume.

The energy conservation equation can be written as,

$$\frac{\partial e}{\partial t} = -\frac{\partial \phi}{\partial x} + F. \quad (1)$$

The expression relating thermal energy density and temperature $u(x, t)$ is,

$$e(x, t) = c(x) \cdot \rho(x) \cdot u(x, t), \quad (2)$$

where $c(x)$ is the specific heat of the material at x , and $\rho(x)$ is its density. Fourier's law of conduction gives a relation between temperature $u(x, t)$ and flux $\phi(x, t)$:

$$\phi = -\sigma_0 \cdot \frac{\partial u}{\partial x}, \quad (3)$$

where σ_0 represents conductivity of the material. Substituting Equations (2) and (3) into Equation (1), we get,

$$c \cdot \rho \cdot \frac{\partial u}{\partial t} = \frac{\partial}{\partial x} \left(\sigma_0 \frac{\partial u}{\partial x} \right) + F, \quad (4)$$

where F is the heat received by the material per unit time per unit volume. Substituting $k = \frac{\sigma_0}{c\rho}$ and $\beta = \frac{1}{c\rho}$ we get,

$$\frac{\partial u}{\partial t} = \frac{\partial}{\partial x} \left(k \frac{\partial u}{\partial x} \right) + \beta F, \quad (5)$$

If we inject heat energy $f = f_{external}$ per unit time per unit volume, and $\epsilon_{hs} = \epsilon_{HeatSource}$ is the material absorptivity for the wavelength corresponding to heat source, then $F = \epsilon_{hs} \cdot f$. This denotes that the heat F absorbed by the material is a fraction of the heat f that is supplied.

Our goal is to estimate k, ϵ from the observations $u(x, t)$ measured with a thermal camera. However, a thermal camera can only measure radiance and not the temperature directly. To solve this, we start by assuming that the absorptivity of the object is $\epsilon = 1$ which gives us the radiation of an ideal black-body, $u_C = u_{black-body}$. The relationship between the real and ideal black body temperature is $\sigma u_{black-body}^4 = \sigma \epsilon u^4$, where σ is the Boltzmann constant. This gives us,

$$\epsilon^{1/4} \frac{\partial u_C}{\partial t} = \frac{\partial}{\partial x} \left(k \frac{\partial}{\partial x} \left(\epsilon^{1/4} u_C \right) \right) + \beta \cdot \epsilon_{hs} \cdot f. \quad (6a)$$

We now have an expression relating the measurements of the thermal camera u_C , the input light source f , and the properties of the material at each spatial location ϵ and k . Our goal is to use the measurements to estimate a spatial distribution of the thermal properties, $\epsilon(x, y)$ and $k(x, y)$.

Equation (6a) is an underconstrained problem as it has two unknowns for every equation. To regularize the problem, we assume that the diffusivity and absorptivity are constant over a small neighborhood giving us,

$$\frac{\partial u_C}{\partial t} = k \frac{\partial^2 u_C}{\partial x^2} + \frac{\beta \cdot \epsilon_{hs}}{\epsilon^{1/4}} f = k \frac{\partial^2 u_C}{\partial x^2} + \epsilon' \cdot f, \quad (6b)$$

where $\epsilon' = \frac{\beta \cdot \epsilon_{hs}}{\epsilon^{1/4}}$. Using the following shorthand notations:

$$u_t = \frac{\partial u_C}{\partial t}, \quad u_{xx} = \frac{\partial^2 u_C}{\partial x^2}, \quad (7a)$$

Equation (6b) can be written as:

$$u_t = k \cdot u_{xx} + \epsilon' \cdot f \quad (7b)$$

We derived all the equations above for the one-dimensional case for simplicity, but the analysis can be easily extended to 3 dimensions. By replacing the Laplacian term u_{xx} from 1-D Equation (7b) by the 3-D Laplacian $\Delta u_{xyz} = u_{xx} + u_{yy} + u_{zz}$, we obtain:

$$u_t = k \cdot \Delta u_{xyz} + \epsilon' \cdot f \quad (8)$$

Equation (8) is central to material classification proposed in this paper. The term k denotes thermal diffusivity of a material and ϵ' is a factor that depends on absorption (in heat source domain) and emission (in LWIR domain). The tuple of parameters forms a unique signature for materials that we use for the downstream task of classification.

2.1. Modeling the source

We use an external source that is switched on for a known t_{ON} duration and switched off. Let $\mathbf{x} = (x, y, z)$ be the combined spatial variable. We model $f(\mathbf{x}, t)$ as:

$$f(\mathbf{x}, t) = \begin{cases} f_s(\mathbf{x}), & \text{for } t \leq t_{ON} \\ 0, & \text{for } t > t_{ON} \end{cases} \quad (9)$$

where $f_s(\mathbf{x})$ is the spatial profile of the external heat source.

For simplicity, we make the following assumptions:

- Our object shape is known and has a flat surface. This constraint can be relaxed by using a method like Structured Light to get the shape of object
- The flat surface of the object faces the camera such that (x, y) coordinates lie in the plane of the surface and z lies perpendicular to it with $z = 0$ being at the surface and increasing as we go inside
- For a given (x, y) the initial temperature is constant for all z , that is,

$$u(x, y, z, t = 0) = u(x, y, t = 0), \quad \forall z \geq 0 \quad (10)$$

2.2. Choosing spatial profile of source $f_s(\mathbf{x})$

A spatially-uniform, known-intensity heat source allows us to estimate ϵ' from the temporal profile of heating and cooling. Estimation of k , however, requires spatial variation in temperature so that diffusion can occur. This can be seen from the Laplacian term, which will be zero across a patch of uniform temperature. This problem, in the absence of source, is the diffusion equation, and hence diffusivity recovery from a uniform-temperature patch is mathematically

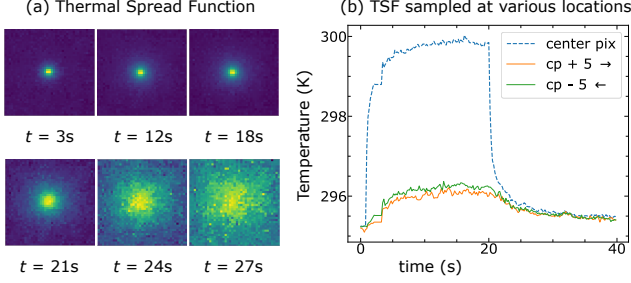


Figure 3. **Thermal Spread Function.** (a) The images captured by the thermal camera as the laser was switched on (top row), and then off (second row) are shown. We call this stack of images the thermal spread function as it encodes the spatial and temporal spread of thermal radiation in the object. In (b), **dashed blue** line graph on top shows the measured temperature at the point where the laser was projected and the bottom **solid green** and **solid orange** line plots show temperature 5 pixels away from center. TSF carries a unique signature about the material which is evident from the temperature plots.

identical to depth from Gaussian defocus on a texture-free patch [21,22]. Therefore, a useful illumination pattern must contain spatial gradients to enable robust estimation of diffusivity parameter. So we decided to use focused points on the scene that generate the required heat profile and hence the texture to get gradients.

2.3. Inverse Heat Problem

Given k , ϵ' , $f_s(x)$ and initial conditions, we can leverage the forward model to obtain temperature profile throughout the volume. However, we can only capture the surface measurements at various time steps. This requires us to solve the inverse heat problem of recovering the thermal parameters k and ϵ' . This is a challenging underconstrained problem. We relax the problem to make it invertible by leveraging assumptions made in Sec. 2.1. Given the assumptions, we can use Finite Difference method [23], Green's function-based approach [20], or a simple curve fitting-based approach to solve for the spatial distribution of k , ϵ' . We discuss their details in Sec. 3.2.

3. Proposed approach

3.1. Heating & Cooling - Thermal Spread Functions

A known external heat source $f_s(x)$ that injects heat into the materials for time t_{ON} increases the temperature of the object. The injected heat will increase temperature at the points of injection and this rise in temperature leads to diffusion of heat across the rest of the material.

This combined phenomenon of absorption and diffusion creates a distinct temperature profile over the surface that is observed by the thermal camera. This temperature profile varies over time as more heat is absorbed, leading to in-

crease in the temperature variation over the material, which in turn leads to more diffusion. Once the heat source is switched off, heat diffusion occurs until all points in that object are at the same temperature. A thermal camera measures this time-varying temperature on the surface of the object which we define as the Thermal Spread Function (TSF) of the material. For a given material, the TSF depends on initial temperature, and the spatial profile and intensity of external heat source $f_s(x)$. Fig. 3 demonstrates the TSF for some common material in a scene.

We parameterize the TSF as a function of the parameters k and ϵ' from Eq. (8). This gives us two-fold advantage, (1) it spares us from processing large data for each TSF directly for the classification problem, and (2) TSFs are dependent on initial temperature, so TSF by themselves cannot be used without accounting for the change in initial temperature.

3.2. Finite Difference Method for Inverse Heat problems

Given initial conditions, k , ϵ' , and external heat source $f_s(x)$, we can use many heat propagation methods to obtain the temperature distribution over time. Some of the methods include Finite Element Analysis (including FEA Softwares like Ansys, Abaqus), convolving with heat kernels (Green's functions) or Finite Difference (FD) Numerical Methods.

We choose FD method for our forward propagation model as they enable a computationally tractable way to obtain k and ϵ' via gradient descent-based optimization. We present below a short analysis of the steps for forward propagation using FD (please refer [20,23] for detailed analysis). Using the explicit Forward Time, Centered Space method, we obtain time derivatives with a forward difference as:

$$\frac{\partial u}{\partial t} = u_t = \frac{u_{x,y,z}^{t+1} - u_{x,y,z}^t}{\Delta t}, \quad (11)$$

Similarly, central difference approximation to the Laplacian term can be written as follows (writing only the x term, other terms can be written similarly):

$$\frac{\partial^2 u}{\partial x^2} = u_{xx} = \frac{u_{x+1,y,z}^t - 2u_{x,y,z}^t + u_{x-1,y,z}^t}{(\Delta x)^2}, \quad (12)$$

The complete Laplacian term can then be written as a sum of u_{xx} , u_{yy} and u_{zz} . Substituting Eqs. (11) and (12) into Eq. (8), we obtain:

$$u_{x,y,z}^{t+1} = u_{x,y,z}^t + \Delta t [k(x,y,z) \cdot \Delta u_{x,y,z}^t + \epsilon'(x,y,z) \cdot f^t(x)] \quad (13)$$

Eq. (13) is the time stepping equation where we obtain the temperature values for time-step $t + 1$ given the values at time-step t . We start these calculations by applying the initial conditions which give us values at $t = 0$. Other

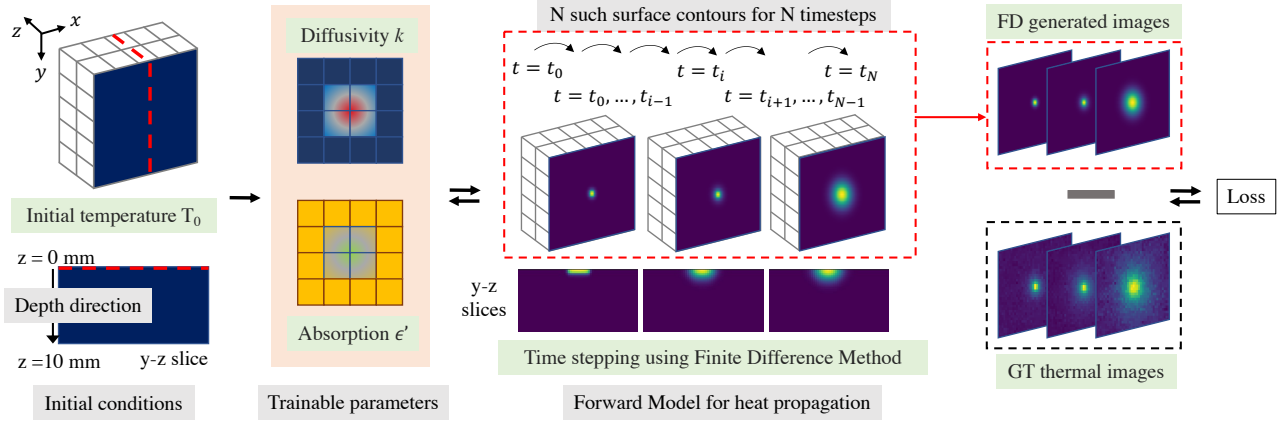


Figure 4. **Differentiable finite difference (FD) method for estimating thermal parameters.** Computing the diffusivity and emissivity of an object requires solving the heat equation. We propose a novel iterative approach using FD method. We implement a forward operator that takes diffusivity and absorptivity as inputs and uses FD method to generate the TSF for a given set of time stamps. The generated TSF is compared against the measured TSF from the camera with an MSE loss function. We then perform multiple iterations to refine the estimates of diffusivity and absorptivity. Since we leverage physics of TSF, the results are highly accurate, and robust to noise in measurements.

relevant boundary conditions are applied during each step as required. Using this method, we can obtain temperature values at any spatial and temporal point by time-stepping. Refer to Fig. 4 for a visualization of the process.

Thus, given the initial conditions, boundary conditions, and all thermal properties, we can define the forward model using FD as follows:

$$\mathcal{F} : u_{\mathbf{x}}^{t=0}, u_{\mathbf{x}=0}^t, k_{\mathbf{x}}, \epsilon'_{\mathbf{x}} \longrightarrow u_{\mathbf{x}}^t \quad (14)$$

implying the temperature at (x, y, z, t) can be obtained as $u(x, y, z, t) = \mathcal{F}(k, \epsilon', x, y, z, t)$.

This function \mathcal{F} forms the basis of our optimization process. This constitutes our forward model where we optimize for k and ϵ' given all the boundary and initial conditions, and external heat source. We constrain this optimization process using the ground truth 2D images we have of the surface temperature of the object over time.

3.3. Drawbacks of Image Based Analysis

As discussed in Sec. 2.3, one of the other methods to solve for k and ϵ' is to do data-based curve fitting. So we need to solve for k and ϵ' given a set of linear equations, where we would know the values of u_t and the Laplacian term $\Delta u_{x,y,z}$. Ideally once we have $u_C(x, y, z, t)$, the time derivative and Laplacian can be calculated using finite difference approximation. The drawback here is that we do not have $u_C(x, y, z, t)$ for all values of z but only for $z = 0$. This means from the image based values $u_C(x, y, z = 0, t)$, we can get u_t and only the x and y terms of the Laplacian $\Delta u_{x,y,z}$ but not the z term u_{zz} . So if we do the data-based curve fitting on this surface data, we will be ignoring the u_{zz} term which leads to inaccurate results. Fig. 5(a, b) demon-




Figure 5. **Drawbacks of using u_t and u_{xx} directly from images.** Image based method (a) gives a 2D Laplacian term (Eq. (8)) assuming 2D diffusion. This leads to inaccurate donut-shaped results. Our method - inverse FD (b) recovers accurate diffusivity and absorption in the region of heat diffusion by also accounting for diffusion in depth.

strates the inaccuracies observed in the recovered values of diffusivity k over the surface of the material.

Notice the doughnut shaped diffusivity result obtained for k in Fig. 5(a). Intuitively, this is because we ignore the diffusion taking place in the depth dimension. This is one of the main reasons we shift from the data-based solution to Finite Difference (FD) Method described in Sec. 3.2.

3.4. Classification of Materials based on k, ϵ'

Once we convert our TSF into the thermal parameters, we use those to characterize the materials and classify them into different categories and also identify them. We make 4 categories of materials - fabric, wood, paper, and metals. We have a dictionary of 22 different materials, each material falling into one of the four categories.

For known materials, we create a scatter plot using the recovered values of diffusivity k and absorption parameter ϵ' as shown in Fig. 6(b). For each material, we take multiple

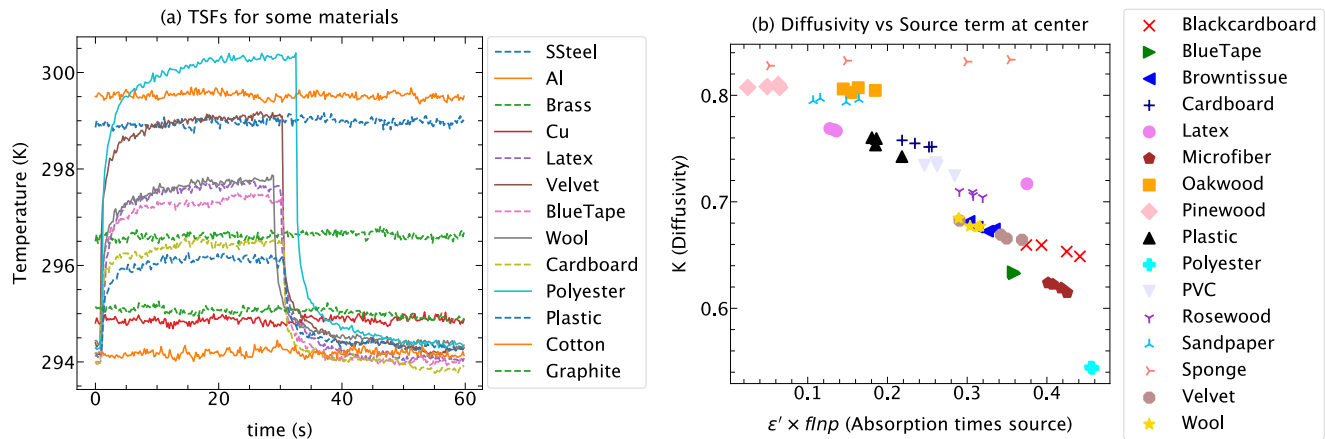


Figure 6. **Material classification with thermal properties** (k , ϵ'). (a) TSFs at center pixel for 13 materials. The curves for heat conductors - Stainless Steel, Aluminium, Brass, Copper and Graphite are almost flat, owing to their high diffusivity and low absorptivity; TSF for cotton is flat due to its high transmissivity. We use the other TSFs to obtain the (k, ϵ') tuples as shown in (b). The plot (b) shows clustering of diffusivity and absorptivity for various materials. Expectedly, the tuple of thermal properties form separate natural clusters corresponding to each material which enables us to use simple classification approaches. Here, we use a random forest classifier that takes a few values of (k, ϵ') around the center as input and outputs the material class.

readings by changing its orientation and initial temperature. Our method is robust to these variations and gives clusters of data points for each of the materials. Once we have this dictionary, we can use the recovered k and ϵ' obtained from an unknown material and classify it into one of the material classes. We use a multi-layer perceptron (MLP, one hidden layer with 90 neurons) as our classifier for its simplicity and speed, and its suitability for small dataset. We train the MLP classifier on the dictionary we create and use that to classify a new material based on k and ϵ' . We also experiment and compare performance with other suitable classifiers - Support Vector Classifiers (SVC, nuSVC), and Random Forest classifier (RF) (refer Tab. 1, Sec. 4.3).

4. Experimental Setup & Results

4.1. Simulations

We use ANSYS Fluent FEM Analysis Software for simulating the heat diffusion process. We use a custom function to define the heat injection which we allow to penetrate to a certain depth. We model this injection to be equivalent of using a relatively high power laser to heat the objects. Modeling the laser means that we also need to simulate the Gaussian beam profile of the heating source.

We can choose one of the existing materials or can create a new one based on the properties from Eq. (5) and Eq. (6b). We model the custom heat injection based on our laser source and choose its intensity accordingly. We assume same surface penetration depth for all materials (the penetration is so low that we can assume it only hits the surface layer of our discretized voxel grid). We run the simulations for 40 seconds, with the heat source being on for

$t_{ON} = 20$ seconds. We use a timestep Δt of 0.25s. The total size of our object is $50\text{mm} \times 50\text{mm} \times 30\text{mm}$ where each discretized voxel size is $\Delta x = \Delta y = \Delta z = 0.5\text{mm}$.

For the simulation process, we define a mesh on a cuboid and only take the surface readings from the simulation. A thermal camera would also provide us only surface readings, so we use these simulated surface contours as ground truth to train our FD model and test its performance. We used a material from ANSYS library (wood) with thermal properties as - density $\rho = 700 \text{ kg/m}^3$, specific heat $C_p = 2310 \text{ J/(kg K)}$ and thermal conductivity $\sigma = 0.173 \text{ W/(m K)}$. From Eq. (5), the thermal diffusivity of this material is $1.069 \times 10^{-7} \text{ m}^2/\text{s}$. When we solve the inverse problem stated above using our FD optimization framework, we obtain a diffusivity value of $1.066 \times 10^{-7} \text{ m}^2/\text{s}$, which is very close to that of the ground truth value.

4.2. Experimental Setup

Fig. 1(a) shows our experimental setup. We use a FLIR A655sc thermal camera, which has a resolution of 640×480 , and works in the LWIR range from $7\text{-}15 \mu\text{m}$, and a 633 nm , 60 mW laser with an Arduino controlled relay switch. This power rating ensures that the laser is powerful enough to heat while safe enough to use. We control the system remotely as any warm body present near the setup can hamper with the readings because of thermal radiation.

We use a spatially varying source, a gaussian profiled laser beam. We use a total capture time of 60 seconds, in which the laser is on for $t_{ON} = 20$ seconds. In the streaming mode of the camera, we capture 300 frames for 60 seconds, implying each time step measures $\Delta t = 0.20$ seconds. We found that our approach resulted in similar results for

Features	SVC	nuSVC	RF	MLP
2	68.8%	76.6%	82.8%	76.6%
18	65.6%	82.8%	84.4%	81.2%
50	64.0%	85.9%	84.4%	85.9%

Table 1. **Comparing accuracies of classifiers.** We experimented with various classifiers and number of features for our dataset. We observed that an MLP and nuSVC perform best with 25 features taken each from the recovered k and ϵ' images (window size of 5). The accuracy did not improve upon adding further features.

a total capture duration of 20s or higher, detailed analysis on this is presented in the supplementary. Our algorithm can also handle multiple simultaneous scans as shown in Fig. 8(a) We placed a graded ruler in the scene next to the target objects to calibrate the physical size of each pixel. We perform this calibration every time we move the setup, but on average our value of spatial Δx was 0.5 mm (same for all dimensions). This process could be automated using structured light methods (refer Sec. 6).

We code the differentiable FD algorithm in python using PyTorch [24]. We use an ADAM classifier with a decreasing learning rate for 400 epochs. Our system consisting of an RTX 3060, takes 5 minutes for the process to complete and recover the k and ϵ' maps. We captured measurements for each material in different initial temperatures and orientations. The code along with the dataset has been released.

4.3. Results

In this section, we analyze the experimental observations made in lab and look at the performance of our parameter recovery and classification algorithms.

Similar to [10], we chose a few materials from the coarse categories of fabric, wood, paper, and metals. We made a dataset consisting of 22 materials or subclasses, applied our recovery algorithm on the data and used the generated features to train the classifier. Since metals have very low absorption and high diffusion, their TSF variation tends to have a very low magnitude. Since the ϵ' values are almost zero for metals, the FD algorithm did not converge for them. Although we cannot differentiate between different metals, this can be used as an indicator for metals. The data in Fig. 6(a) for metals validates the theory for metals. Barring the conducting materials, we performed FD analysis on the other materials and present our analysis on them.

Fig. 6 shows a plot of k vs ϵ' sampled at the center pixels of their recovered maps. For each measurement, we generate a feature vector by sampling k vs ϵ' in a window of specified size around the center. This feature vector along with the material label forms our training set for the classification algorithm. We perform a leave-one-out cross validation approach to obtain the classification accuracy of our method. The classification results can be found in the con-

fusion matrix in Fig. 7(b). We obtain an overall accuracy of 85.9% using a multi-layer perceptron (MLP) classifier with a feature size of 50 pixels. With more features, the subclasses become more separable compared to what we see in Fig. 6(b). We also compare performance with varying feature sizes and classifiers suitable for dataset Tab. 1.

Materials that have similar diffusivity, and similar absorptivity for red wavelength (used in our experiments), produce similar TSFs. This in turn makes the recovered properties of these materials similar to each other. Such an issue occurs for brown tissue and wool as seen in Fig. 6(b) and Fig. 7(b) where the algorithm gets confused between the two. This can be mitigated by using more lasers like red, green and blue which gives us a set of absorption coefficients enabling more efficient segregation.

5. Limitations


Our current modeling of thermodynamics requires certain simplifying assumptions for accurate results, such as planar geometry and constant temperature on the surface and bulk of the objects. These assumptions have enabled robust classification results in our preliminary experiments, but can be relaxed to extend to complex shaped objects with varying temperature profiles. Our approach can be combined with approaches for estimating scene geometry such as structured light [25–27] to estimate material properties of complex scenes and will be pursued as promising future direction.

6. Future Work

Our experiments covered a limited set of materials to show that their TSFs can be used for material classification. Collecting a larger dataset that covers varying surface color, roughness, and geometry will enable a more robust material classification and will be pursued as future direction. We used Finite Differences in the forward model of our analysis. This part can also be potentially replaced by other approaches such as Physics Informed Neural Networks (PINNs) [28–30]. We can also potentially perform joint estimation of shape and the material properties using thermal structured Light. A known pattern of laser dots on the scene can be leveraged to estimate the thermal properties using the heat diffusion as well as the shape using the texture that is generated.

6.1. Going Below the Surface

An advantage of thermal material classification over other methods like spectroscopy is that analyzing the heat diffusion can give us information that is hidden to the eye (even the thermal one). Consider a system as shown in Fig. 8(b), where the top layer is a different material from the one below it. When we inject heat, the upper layer ab-



(a) Some of the materials used for experiments

BlueTape	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Cardboard	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PVC	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Pinewood	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Plastic	0	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Polyester	0	0	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Wool	0	0	0	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0
microfiber	0	0	0	0	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0
Blackcard	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	1	0
Browntiss	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0
Latex	0	0	0	0	0	0	0	0	0	0	0	3	0	1	0	0	0	0	0	0
Oakwood	0	0	0	0	0	0	0	0	0	0	0	0	3	0	1	0	0	0	0	0
Rosewood	0	0	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0
Sandpaper	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	1	0
Velvet	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0
Sponge	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	2
BlueTape																				
Cardboard																				
PVC																				
Pinewood																				
Plastic																				
Polyester																				
Wool																				
microfiber																				
Blackcard																				
Browntiss																				
Latex																				
Oakwood																				
Rosewood																				
Sandpaper																				
Velvet																				
Sponge																				

(b) Confusion Matrix for classification on non-metals

Figure 7. **Material library and classification results:** We applied our recovery and classification algorithm on our dictionary of materials (a), and show the classification results using confusion matrix in (b). We use random forest as our classifier and use leave-one-out cross validation technique for testing our algorithm. We get an overall accuracy of 81.3%.

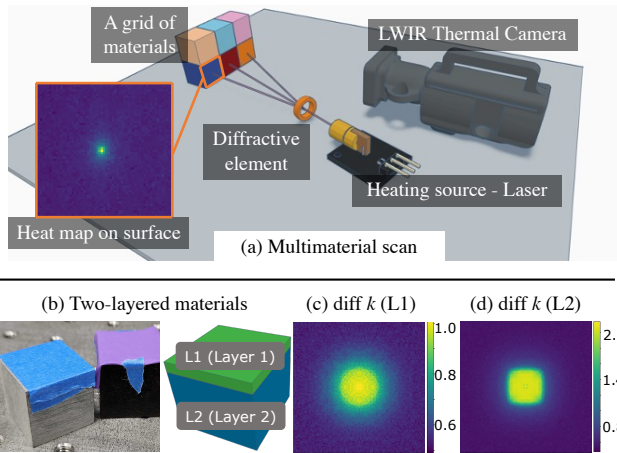


Figure 8. **Extending our approach.** (a) Our setup can be extended to accommodate simultaneous scans. (b) We ran simulations to understand the performance of our approach on multi-layered materials (b) and to analyse the possibility of recovering thermal properties below the surface. We recovered diffusivity of layers 1 and 2 (c, d) with 0.01% error, showing promise in "seeing the hidden"

sorbs it based on its surface and absorption properties, while the diffusion is affected by the diffusivity of both the upper layer and the layer underneath. We analyzed this problem in simulation where we assume a known thickness of upper layer but unknown thermal parameters. We used the FD forward model and optimized for diffusivity values - for the surface as well as the underneath layer, and the absorption coefficient of the surface layer. Fig. 8(c,d) show the recovered thermal properties for our simulations. We assumed a thickness of 0.5 mm and obtained the diffusivity values of

the two layers with 0.01% error. This shows promise in the approach and is a good direction for future work.

7. Conclusion

Tapping into the intrinsic properties of objects such as a material's thermodynamic constants (k, ϵ) enables robust classification, and alleviates some of the challenges faced by the more common visible camera-based approaches. We proposed a simple optical setup consisting of a light source and thermal camera that captures the TSF of the scene. We then used a novel FD-based iterative solver that robustly estimates k, ϵ which was used for the downstream task of material classification. Thermal material classification can be used as a complementary modality along with existing techniques including visible sensing. In biometrics, spoofing-based attacks such as using a mannequin for a real human become easier to detect due to the complementary strengths of thermal and visible cameras. Hence our work also has the potential to impact numerous areas in vision and security surveillance.

Acknowledgements: This work was supported by NSF awards IIS-2107313 and IIS-1730574, and ONR award N00014-21-1-2035. We thank Manuel Ballester for his help in formulating numerical methods for the heat equation, Rishubh Parihar for discussion on optimization, and Lionel Fiske for discussion on PDE solutions. We also thank the reviewers for their valuable suggestions to improve the quality of the manuscript.

References

- [1] Patitapabana Parida, Ajit Behera, and S Chandra Mishra. Classification of biomaterials used in medicine. *International Journal of Advances in Applied Sciences (IJAAS)*, 2012. 1
- [2] Joerg Hendrix. Technologies for the identification, separation and recycling of automotive plastics. *Int. J. Environmentally Conscious Design & Manufacturing*, 5(1):35–47, 1996. 1
- [3] Alfons Buekens and Xujian Zhou. Recycling plastics from automotive shredder residues: A review. *Journal of Material Cycles and Waste Management*, 16(3):398–414, 2014. 1
- [4] Oliver Wang, Prabath Gunawardane, Steve Scher, and James Davis. Material classification using brdf slices. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2805–2811. IEEE, 2009. 1
- [5] Neda Salamati, Clément Fredembach, and Sabine Süssstrunk. Material classification using color and nir images. In *Color and Imaging Conference*, volume 2009, pages 216–222. Society for Imaging Science and Technology, 2009. 1
- [6] Kenichiro Tanaka, Yasuhiro Mukaigawa, Takuya Funatomi, Hiroyuki Kubo, Yasuyuki Matsushita, and Yasushi Yagi. Material classification using frequency-and depth-dependent time-of-flight distortion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 79–88, 2017. 1
- [7] Abdelhameed F Ibrahim, Shoji Tominaga, and Takahiko Horiuchi. Spectral imaging method for material classification and inspection of printed circuit boards. *Optical Engineering*, 49(5):057201, 2010. 1
- [8] Vishwanath Saragadam and Aswin Sankaranarayanan. Programmable spectrometry: Per-pixel material classification using learned spectral filters. *IEEE Intl. Conf. Computational Photography (ICCP)*, 2020. 1
- [9] Zackory Erickson, Eliot Xing, Bharat Srirangam, Sonia Chernova, and Charles C Kemp. Multimodal material classification for robots using spectroscopy and high resolution texture imaging. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10452–10459. IEEE, 2020. 1
- [10] Philip Saponaro, Scott Sorensen, Abhishek Kolagunda, and Chandra Kambhamettu. Material classification with thermal imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4649–4656, 2015. 1, 7
- [11] Jakub Bednarek, Michal Bednarek, Piotr Kicki, and Krzysztof Walas. Robotic touch: Classification of materials for manipulation and walking. In *2019 2nd IEEE international conference on Soft Robotics (RoboSoft)*, pages 527–533. IEEE, 2019. 2
- [12] Jun Zhang, Han Li, Aiguo Song, Yizhuang Ding, and Juan Wu. Thermal perception for information transmission: Theoretical analysis, device design, and experimental verification. *IEEE Transactions on Haptics*, 2022. 2
- [13] Mohamed Guiatni, Abdelaziz Benallegue, and Abderrahmane Kheddar. Thermal display for telepresence based on neural identification and heat flux control. *Presence: Teleoperators and Virtual Environments*, 18(2):156–169, 2009. 2
- [14] Anshul Singhal. *Development of thermal displays for haptic interfaces*. PhD thesis, Massachusetts Institute of Technology, 2016. 2
- [15] Amine Tilioua, Laurent Libessart, and Stéphane Lassue. Characterization of the thermal properties of fibrous insulation materials made from recycled textile fibers for building applications: Theoretical and experimental analyses. *Applied Thermal Engineering*, 142:56–67, 2018. 2
- [16] WJ Parker, RJ Jenkins, CP Butler, and GL Abbott. Flash method of determining thermal diffusivity, heat capacity, and thermal conductivity. *Journal of applied physics*, 32(9):1679–1684, 1961. 2
- [17] Robert D Cowan. Pulse method of measuring thermal diffusivity at high temperatures. *Journal of Applied Physics*, 34(4):926–927, 1963. 2
- [18] I Garrido, S Lagüela, R Otero, and P Arias. Thermographic methodologies used in infrastructure inspection: A review—post-processing procedures. *Applied Energy*, 266:114857, 2020. 2
- [19] Theodore L Bergman, Adrienne S Lavine, Frank P Incropera, and David P DeWitt. *Introduction to heat transfer*. John Wiley & Sons, 2011. 2
- [20] Richard Haberman. *Applied partial differential equations*, volume 4. Prentice Hall Upper Saddle River, NJ., 2003. 2, 4
- [21] Qi Guo, Emma Alexander, and Todd Zickler. Focal track: Depth and accommodation with oscillating lens deformation. In *Proceedings of the IEEE international conference on computer vision*, pages 966–974, 2017. 4
- [22] Emma Alexander, Qi Guo, Sanjeev Koppal, Steven Gortler, and Todd Zickler. Focal flow: Measuring distance and velocity with defocus and differential motion. In *European conference on computer vision*, pages 667–682. Springer, 2016. 4
- [23] Keith W Morton and David Francis Mayers. *Numerical solution of partial differential equations: an introduction*. Cambridge university press, 2005. 4
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochele, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 7
- [25] Martin Landmann, Stefan Heist, Anika Brahm, Simon Schindwolf, Peter Kühmstedt, and Gunther Notni. 3d shape measurement by thermal fringe projection: optimization of infrared (ir) projection parameters. In *Dimensional Optical Metrology and Inspection for Practical Applications VII*, volume 10667, pages 9–18. SPIE, 2018. 7
- [26] Martin Landmann, Stefan Heist, Patrick Dietrich, Henri Speck, Peter Kühmstedt, Andreas Tünnermann, and Gunther

Notni. 3d shape measurement of objects with uncooperative surface by projection of aperiodic thermal patterns in simulation and experiment. *Optical Engineering*, 59(9):094107, 2020. 7

- [27] Jack Erdozain, Kazuto Ichimaru, Tomohiro Maeda, Hiroshi Kawasaki, Ramesh Raskar, and Achuta Kadambi. 3d imaging for thermal cameras using structured light. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 2795–2799. IEEE, 2020. 7
- [28] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019. 7
- [29] Shengze Cai, Zhicheng Wang, Sifan Wang, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks for heat transfer problems. *Journal of Heat Transfer*, 143(6), 2021. 7
- [30] Zhili He, Futao Ni, Weiguo Wang, and Jian Zhang. A physics-informed deep learning method for solving direct and inverse heat conduction problems of materials. *Materials Today Communications*, 28:102719, 2021. 7