

# Differential Angular Imaging for Material Recognition

Jia Xue<sup>1</sup> Hang Zhang<sup>1</sup> Kristin Dana<sup>1</sup> Ko Nishino<sup>2</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, Rutgers University, Piscataway, NJ 08854

<sup>2</sup>Department of Computer Science, Drexel University, Philadelphia, PA 19104

{jia.xue, zhang.hang}@rutgers.edu, kdana@ece.rutgers.edu, kon@drexel.edu

## Abstract

Material recognition for real-world outdoor surfaces has become increasingly important for computer vision to support its operation “in the wild.” Computational surface modeling that underlies material recognition has transitioned from reflectance modeling using in-lab controlled radiometric measurements to image-based representations based on internet-mined images of materials captured in the scene. We propose a middle-ground approach that takes advantage of both rich radiometric cues and flexible image capture. We develop a framework for differential angular imaging, where small angular variations in image capture provide an enhanced appearance representation and significant recognition improvement. We build a large-scale material database, Ground Terrain in Outdoor Scenes (GTOS), geared towards real use for autonomous agents. This publicly available database<sup>1</sup> consists of over 30,000 images covering 40 classes of outdoor ground terrain under varying weather and lighting conditions. We develop a novel approach for material recognition called a Differential Angular Imaging Network (DAIN) to fully leverage this large dataset. With this network architecture, we extract characteristics of materials encoded in the angular and spatial gradients of their appearance. Our results show that DAIN achieves recognition performance that surpasses single view or coarsely quantized multiview images. These results demonstrate the effectiveness of differential angular imaging as a means for flexible, in-place material recognition.

## 1. Introduction

Real world scenes consist of surfaces made of numerous materials, such as wood, marble, dirt, metal, ceramic and fabric, which contribute to the rich visual variation we find in images. Material recognition has become an active

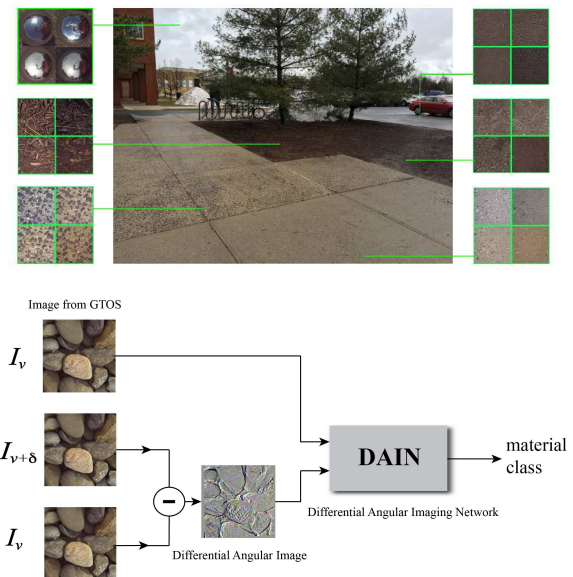


Figure 1: (Top) Example from GTOS dataset comprising outdoor measurements with multiple viewpoints, illumination conditions and angular differential imaging. The example shows scene-surfaces imaged at different illumination/weather conditions. (Bottom) Differential Angular Imaging Network (DAIN) for material recognition.

area of research in recent years with the goal of providing detailed material information for applications such as autonomous agents and human-machine systems.

Modeling the apparent or latent characteristic appearance of different materials is essential to robustly recognize them in images. Early studies of material appearance modeling largely concentrated on comprehensive lab-based measurements using dome systems, robots, or gonireflectometers collecting measurements that are dense in angular space (such as BRDF, BTF) [10]. These reflectance-based studies have the advantage of capturing intrinsic in-

<sup>1</sup><http://ece.rutgers.edu/vision/>

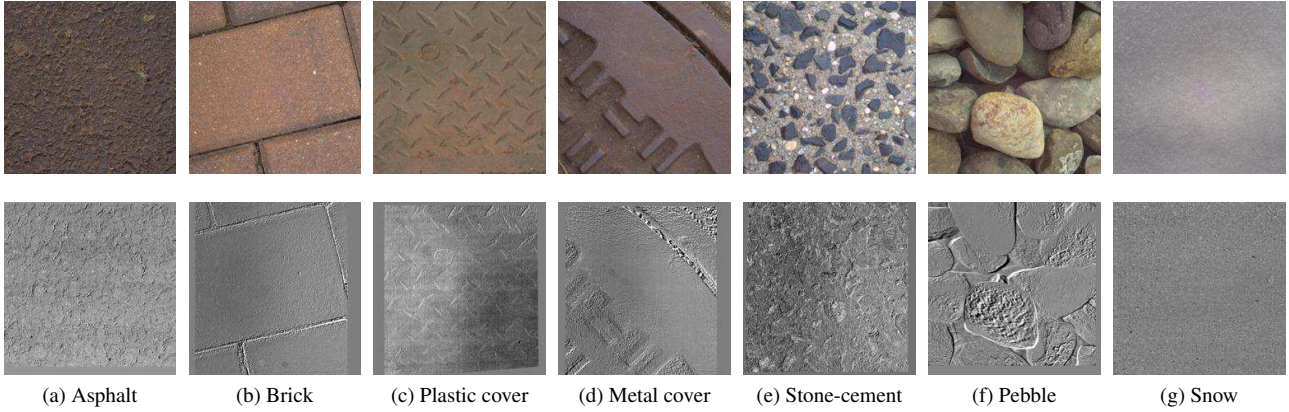


Figure 2: Differential Angular Imaging. (Top) Examples of material surface images  $I_v$ . (Bottom) Corresponding differential images  $I_\delta = I_v - I_{v+\delta}$  in our GTOS dataset. These sparse images encode angular gradients of reflection and 3D relief texture.

variant properties of the surface, which enables fine-grained material recognition [27, 33, 42, 47]. The inflexibility of lab-based image capture, however, prevents widespread use in real world scenes, especially in the important class of outdoor scenes. A fundamentally different approach to reflectance modeling is image-based modeling where surfaces are captured with a single-view image in-scene or “in-the-wild.” Recent studies of image-based material recognition use single-view internet-mined images to train classifiers [1, 7, 20, 28] and can be applied to arbitrary images casually taken without the need of multiview reflectance information. In these methods, however, recognition is typically based more on context than intrinsic material appearance properties except for a few purely local methods [34, 35].

Between comprehensive in-lab imaging and internet-mined images, we take an advantageous middle-ground. We capture in-scene appearance but use controlled viewpoint angles. These measurements provide a sampling of the full reflectance function. This leads to a very basic question: how do multiple viewing angles help in material recognition? Prior work used differential camera motion or object motion for shape reconstruction [2, 3, 43], here we consider a novel question: Do small changes in viewing angles, *differential changes*, result in significant increases in recognition performance? Prior work has shown the power of angular filtering to complement spatial filtering in material recognition. These methods, however, rely on a mirror-based camera to capture a slice of the BRDF [48] or a light-field camera to achieve multiple differential viewpoint variations [44] which limits their application due to the need for specialized imaging equipment. We instead propose to capture surfaces with differential changes in viewing angles with an ordinary camera and compute discrete approximations of *angular gradients* from them. We present an approach called *angular differential imaging* that augments

image capture for a particular viewing angle  $v$  a differential viewpoint  $v + \delta$ . Contrast this method with lab-based reflectance measurements that often quantize the angular space measuring with domes or positioning devices with large angular spacing such as  $22.5^\circ$ . These coarse-quantized measurements have limited use in approximating angular gradients. Angular differential imaging can be implemented with a small-baseline stereo camera or a moving camera (e.g. handheld). We demonstrate that differential angular imaging provides key information about material reflectance properties while maintaining the flexibility of convenient in-scene appearance capture.

To capture material appearance in a manner that preserves the convenience of image-based methods and the important angular information of reflectance-based methods, we assemble a comprehensive, first-of-its-kind, *outdoor* material database that includes multiple viewpoints and multiple illumination directions (partial BRDF sampling), multiple weather conditions, a large set of surface material classes surpassing existing comparable datasets, multiple physical instances per surface class (to capture intra-class variability) and differential viewpoints to support the framework of differential angular imaging. We concentrate on outdoor scenes because of the limited availability of reflectance databases for outdoor surfaces. We also concentrate on materials from ground terrain in outdoor scenes (GTOS) for applicability in numerous application such as automated driving, robot navigation, photometric stereo and shape reconstruction. The 40 surface classes include ground terrain such as grass, gravel, asphalt, concrete, black ice, snow, moss, mud and sand (see Figure 2).

We build a recognition algorithm that leverages the strength of deep learning and differential angular imaging. The resulting method takes two image streams as input, the original image and a differential image as illustrated in Fig-

Datasets	samples	classes	views	illumination	in scene	scene image	camera parameters	year
CUReT [11]	61	61	205		N	N	N	1999
KTH-TIPS [18]	11	11	27	3	N	N	N	2004
UBO2014 [46]	84	7	151	151	N	N	N	2014
Reflectance disk [48]	190	19	3	3	N	N	Y	2015
4D Light-field [44]	1200	12	1	1	Y	N	N	2016
NISAR [5]	100	100	9	12	N	N	N	2016
<b>GTOS(ours)</b>	<b>606</b>	<b>40</b>	<b>19</b>	<b>4</b>	<b>Y</b>	<b>Y</b>	<b>Y</b>	<b>2016</b>

Table 1: Comparison between GTOS dataset and some publicly available BRDF material datasets. Note that the 4D Light-field dataset [44] is captured by the Lytro Illum light field camera.

ure 1. We optimize the two-stream configuration for material recognition performance and call the resulting network DAIN—differential angular imaging network.

We make three significant contributions in this paper: 1) Introduction of differential angular imaging as a middle-ground between reflectance-based and image-based material recognition; 2) Collection of the GTOS database made publicly available with over 30000 in-scene outdoor images capturing angular reflectance samples with scene context over a large set of material classes; 3) The development of DAIN, a material recognition network with state-of-the-art performance in comprehensive comparative validation.

## 2. Related Work

Texture recognition, the classification of 3D texture images and bidirectional texture functions, traditionally relied on hand-designed 3D image features and multiple views [8, 24]. More recently, features learned with deep neural networks have outperformed these methods for texture recognition. Cimpoi *et al.* [7] achieves state-of-art results on FMD [36] and KTH-TIPS2 [18] using a Fisher vector representation computed on image features extracted with a CNN.

The success of deep learning methods in object recognition has also translated to the problem of material recognition, the classification and segmentation of material categories in arbitrary images. Bell *et al.*, achieve per-pixel material category labeling by retraining the then state-of-the-art object recognition network [37] on a large dataset of material appearance [1]. This method relies on large image patches that include object and scene context to recognize materials. In contrast, Schwartz and Nishino [34, 35] learn material appearance models from small image patches extracted inside object boundaries to decouple contextual information from material appearance. To achieve accurate local material recognition, they introduced intermediate material appearance representations based on their intrinsic properties (e.g., “smooth” and “metallic”).

In addition to the apparent appearance, materials can be discerned by their radiometric properties, namely the bidi-


rectional reflectance distribution function (BRDF) [30] and the bidirectional texture function (BTF) [11], which essentially encode the spatial and angular appearance variations of surfaces. Materials often exhibit unique characteristics in their reflectance offering detailed cues to recognize the difference of subtle variations in them (e.g., different types of metal [27] and paint [42]). Reflectance measurements, however, necessitate elaborate image capture systems, such as a gonireflectometer [30, 45], robotic arm [25], or a dome with cameras and light sources [12, 27, 42]. Recently, Zhang *et al.* introduced the use of a one-shot reflectance field capture for material recognition [48]. They adapt the parabolic mirror-based camera developed by Dana and Wang [9] to capture the reflected radiance for a given light source direction in a single shot, which they refer to as a reflectance disk. More recently, Zhang *et al.* showed that the reflectance disks contain sufficient information to accurately predict the kinetic friction coefficient of surfaces [49]. These results demonstrate that the angular appearance variation of materials and their gradients encode rich cues for their recognition. Similarly, Wang *et al.* [44] uses a light field camera and combines angular and spatial filtering for material recognition. In strong alignment with these recent advances in material recognition, we build a framework of spatial and angular appearance filtering. In sharp contrast to past methods, however, we use image information from standard cameras instead of a multilens array as in Lytro. We explore the difference of using a large viewing angle range (with samples coarsely quantized in angle space) by using differential changes in angles which can easily be captured by a two-camera system or small motions of a single ordinary camera.

Deep learning has achieved major success in object classification [4, 19, 23], segmentation [17, 22, 32], and material recognition [7, 26, 49, 50]. In our goal of combining spatial and angular image information to account for texture and reflectance, we are particularly motivated by the two-stream fusion framework [15, 37] which achieves state-of-art results in UCF101 [38] action recognition dataset.





(a) material classes



(b) one sample at multiple viewing directions

Figure 3: (a) The 40 material categories in the GTOS dataset introduced in this paper. (Right) The material surface observation points. Nine viewpoint angles separated along an arc spanning  $80^\circ$  are measured. For each viewpoint, a differential view is captured  $\pm 5^\circ$  in azimuth from the original orientation (the sign is chosen based on robotic arm kinematics.)

**Datasets:** Datasets to measure reflectance of real world surfaces have a long history of lab-based measurements including: CURET database [11], KTH-TIPS database by Hayman *et al.* [18], MERL Reflectance Database [29], UBO2014 BTF Database [46], UTIA BRDF Database [16], Drexel Texture Database [31] and IC-CERTH Fabric Database [21]. In many of these datasets, dense reflectance angles are captured with special image capture equipment. Some of these datasets have limited instances/samples per surface category (different physical samples representing the same class for intraclass variability) or have few surface categories, and all are obtained from indoor measurements where the sample is removed from the scene. More recent datasets capture materials and texture in-scene, (a.k.a. in-situ, or in-the-wild). A motivation of moving to in-scene capture is to build algorithms and methods that are more relevant to real-world applications. These recent databases are from internet-mined databases and contain a single view of the scene under a single illumination direction. Examples include the the Flickr Materials Database by Sharan *et al.* [36] and the Material in Context Database by Bell *et al.* [1]. Recently, DeGol *et al.* released GeoMat Database [13] with 19 material categories from outdoor sites and each category has between 3 and 26 physical surface instances, with 8 to 12 viewpoints per surface. The viewpoints in this dataset are irregularly sampled in angle space.

### 3. Differential Angular Imaging

We present a new measurement method called differential angular imaging where a surface is imaged from a particular viewing angle  $v$  and then from an additional viewpoint  $v + \delta$ . The motivation for this differential change in viewpoint is improved computation of the angular gradient of intensity  $\partial I_v / \partial v$ . Intensity gradients are the basic building block of image features and it is well known that discrete approximations to derivatives have limitations. In particular, spatial gradients of intensities for an image  $I$  are approximated by  $I(x + \Delta) - I(x)$  and this approximation is most reasonable at low spatial frequencies and when  $\Delta$  is small. For angular gradients of reflectance, the discrete approximation to the derivative is a subtraction with respect to the viewing angle. Angular gradients are approximated by  $I(v + \delta) - I(v)$  and this approximation requires a small  $\delta$ . Consequently, differential angular imaging provides more accurate angular gradients.

The differential images as shown in Figures 1 and 2 have several characteristics. First, the differential image reveals the gradients in BRDF/BTF at the particular viewpoint. Second, relief texture is also observable in the differential image due to non-planar surface structure. Finally, the differential images are sparse. This sparsity can provide a computational advantage within the network. (Note that  $I(v + \delta)$  and  $I(v)$  are aligned with a global affine transformation before subtraction.)

## 4. GTOS Dataset

**Ground Terrain in Outdoor Scenes Dataset** We collect the GTOS database, a first-of-its-kind in-scene material reflectance database, to investigate the use of spatial and angular reflectance information of outdoor ground terrain for material recognition. We capture reflectance systematically by imaging a set of viewing angles comprising a partial BRDF with a mobile exploration robot. Differential angular images are obtained by also measuring each of  $N_v = 9$  base angles  $v = (\theta_v, \phi_v)$ ,  $\theta_v \in [-40^\circ, -30^\circ, \dots, 40^\circ]$ , and a differential angle variation of  $\delta = (0, 5^\circ)$  resulting in 18 viewing directions per sample as shown in Figure 3 (b). Example surface classes are depicted in Figure 3 (a). The class names are (in order of top-left to bottom-right): cement, asphalt, painted asphalt, brick, soil, muddy stone, mud, mud-puddle, grass, dry leaves, leaves, asphalt-puddle, mulch, metal grating, plastic, sand, stone, artificial turf, aluminum, limestone, painted turf, pebbles, roots, moss, loose asphalt-stone, asphalt-stone, cloth, paper, plastic cover, shale, painted cover, stone-brick, sandpaper, steel, dry grass, rusty cover, glass, stone-cement, icy mud, and snow. The  $N_c = 40$  surface classes mostly have between 4 and 14 instances (samples of intra-class variability) and each instance is imaged not only under  $N_v$  viewing directions but also under multiple natural light illumination conditions. As illustrated in Figure 1, sample appearance depends on the weather condition and the time of day. To capture this variation, we image the same region with  $N_i = 4$  different weather conditions (cloudy dry, cloudy wet, sunny morning, and sunny afternoon). We capture the samples with 3 different exposure times to enable high dynamic range imaging. Additionally, we image a mirrored sphere to capture the environment lighting of the natural sky. In addition to surface images, we capture a scene image to show the global context. The robot measurement device is depicted in Figure 4. Although, the database measurements were obtained with robotic positioning for precise angular measurements, our recognition results are based on subsets of these measurements so that an articulated arm would not be required for an in-field system. The total number of surface images in the database is 34,243. As shown in Table 1, this is the most extensive outdoor in-scene multiview material database to date.

## 5. DAIN for Material Recognition

**Differential Angular Imaging Network (DAIN)** Consider the problem of in-scene material recognition with images from multiple viewing directions (multiview). We develop a two-stream convolutional neural network to fully leverage differential angular imaging for material recognition. The differential image  $I_\delta$  sparsely encodes reflectance angular gradients as well as surface relief texture. The spa-




Figure 4: The measurement equipment for the GTOS database: Mobile Robots P3-AT robot, Cyton gamma 300 robot arm, Basler aca2040-90uc camera with Edmund Optics 25mm/F1.8 lens, DGK 18% white balance and color reference card, and Macmaster-Carr 440C Stainless Steel Sphere.

tial variation of image intensity remains an important recognition cue and so our method integrates these two streams of information. A CNN is used on both streams of the network and then combined for the final prediction result. The combination method and the layer at which the combination takes place leads to variations of the architecture.

We employ the ImageNet [14] pre-trained VGG-M model [4] as the prediction unit (labeled CNN in Figure 5). The first input branch is the image  $I_v$  at a specific viewing direction  $v$ . The second input branch is the differential image  $I_\delta$ . The first method of combination shown in Figure 5 (a) is a simple averaging of the output prediction vectors obtained by the two branches. The second method combines the two branches at the intermediate layers of the CNN, i.e. the feature maps output at layer  $M$  are combined and passed forward to the higher layers of the CNN, as shown Figure 5 (b). We empirically find that combining feature maps generated by Conv5 layer after ReLU performs best. A third method (see Figure 5 (c)) is a hybrid of the two architectures that preserves the original CNN path for the original image  $I_v$  by combining the layer  $M$  feature maps for both streams *and* by combining the prediction outputs for both streams as shown in Figure 5 (c). This approach is the best performing architecture of the three methods and we call it the differential angular imaging network (DAIN).

For combining feature maps at layer  $M$ , consider features maps  $x_a$  and  $x_b$  from the two branches that have width  $W$ , height  $H$ , and feature channel depth  $D$ . The output feature map  $y$  will be the same dimensions  $W \times H \times D$ . We can combine feature maps by: (1) *Sum*: pointwise sum of  $x_a$  and  $x_b$ , and (2) *Max*: pointwise maximum of  $x_a$  and  $x_b$ . In Section 6 we evaluate the performance of these methods of combining lower layer feature maps.

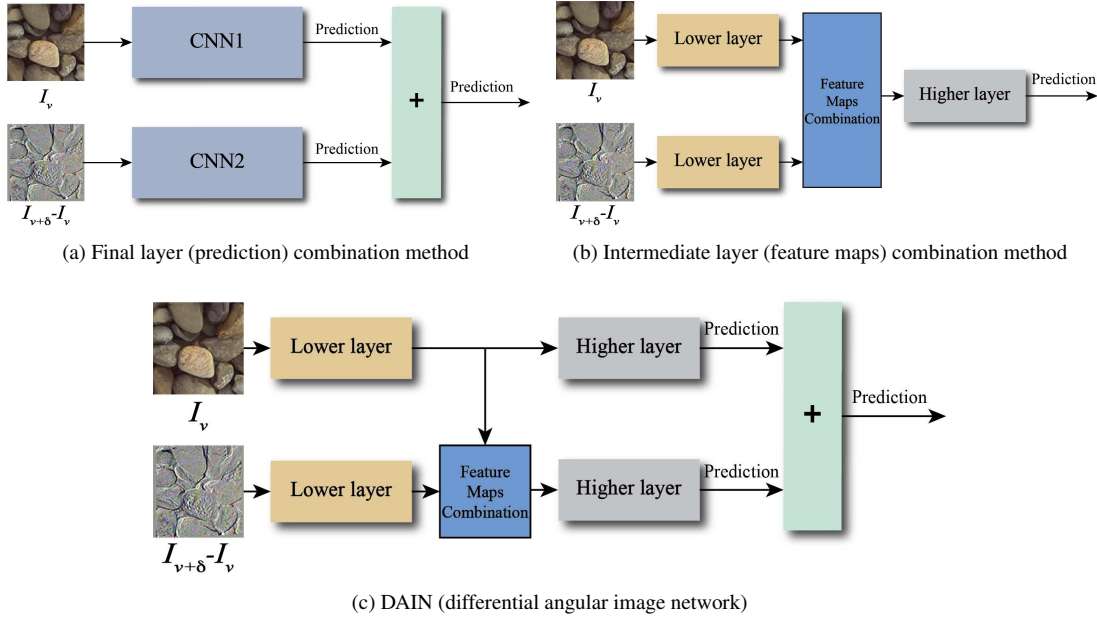


Figure 5: Methods to combine two image streams, the original image  $I_v$  and the differential image  $I_\delta = I_{v+\delta} - I_v$ . The best performing configuration is the architecture in (c), which we refer to as differential angular imaging network (DAIN).

### 5.1. Multiple Views

Our GTOS database has multiple viewing directions on an arc (a partial BRDF sampling) as well as differential images for each viewing direction. We evaluate our recognition network in two modes: (1) **Single view DAIN**, with inputs from  $I_v$  and  $I_\delta$ , with  $v$  representing a single viewing angle; (2) **Multi view DAIN**, with inputs  $I_v$  and  $I_\delta$ , with  $v \in [v_1, v_2, \dots, v_N]$ . For our GTOS database,  $v_1, v_2, \dots, v_N$  are viewing angles separated by  $10^\circ$  representing a  $N \times 10^\circ$  range of viewing angles. We empirically determine that  $N = 4$  viewpoints are sufficient for recognition. For a baseline comparison we also consider non-differential versions: **Single View** with only  $I_v$  for a single viewing direction and **Multi View** with inputs  $I_v, v \in [v_1, v_2, \dots, v_N]$ .

To incorporate multi view information in DAIN we use three methods: (1) voting (use the predictions from each view to vote), (2) pooling (pointwise maximum of the combined feature maps across viewpoints), (3) 3D filter + pooling (follow [40] to use a  $3 \times 3 \times 3$  learned filter bank to convolve the multi view feature maps). See Figure 6. After 3D filtering, pooling is used (pointwise maximum across viewpoints). The computational expense of this third method due to learning the filter weights is significantly higher.

## 6. Experiments

In this section, we evaluate the DAIN framework for material recognition and compare the results on GTOS with several state-of-the-art algorithms. The first evaluation de-

termines which structure of the two stream networks from Figure 5 works best on the GTOS dataset, leading to the choice in (c) as the DAIN architecture. The second evaluation considers recognition performance with different variations of DAIN recognition. The third experimental evaluation compares three other state-of-the-art approaches on our GTOS-dataset, concluding that multiview DAIN works best. Finally, we apply DAIN to a lightfield dataset to show performance in another multiview material dataset.

**Training procedure** We design 5 training and testing splits by assigning about 70% of ground terrain surfaces of each class to training and the rest 30% to testing. Note that, to ensure that there is no overlap between training and testing sets, if one sample is in the training set, all views and illumination conditions for that sample is in the training set.

Each input image from our GTOS database is resized into  $240 \times 240$ . Before training a two branch network, we first fine-tune the VGG-M model separately with original and differential images with batch size 196, dropout rate 0.5, momentum 0.9. We employ the augmentation method that horizontally and vertically stretch training images within  $\pm 10\%$ , with an optional 50% horizontal mirror flips. The images are randomly cropped into  $224 \times 224$  material patches. All images are pre-processed by subtracting a per color channel mean and normalizing for unit variance. The learning rate for the last fully connected layer is set to 10 times of other layers. We first fine-tune only the last fully connected layer with learning rate  $5 \times 10^{-2}$  for 5 epochs; then, fine-tune all the fully connected layers with learning






Figure 6: Multiview DAIN. The 3D filter + pooling method to combine two streams (original and differential image) from multiple viewing angles.  $W$ ,  $H$ , and  $D$  are the width, height, and depth of corresponding feature maps,  $N$  is the number of view points.

rate  $10^{-2}$  for 5 epochs. Finally we fine-tune all the layers with learning rate starting at  $10^{-3}$ , and decrease by a factor of 0.1 when the training accuracy saturates. Since the snow class only has 2 samples, we omit them from experiments.

For the two branch network, we employ the fine-tuned two-branch VGG-M model with batch size 64 and learning rate starting from  $10^{-3}$  which is reduced by a factor of 0.1 when the training accuracy saturates. We augment training data with randomly stretch training images by  $\pm 25\%$  horizontally and vertically, and also horizontal mirror flips. The images are randomly cropped to  $224 \times 224$  material patches. We first backpropagate only to feature maps combination layer for 3 epochs, then fine tunes all layers. We employ the same augmentation method for the multiview images of each material surface. We randomly select the first viewpoint image, then subsequent  $N = 4$  view point images are selected for experiments.

**Evaluation for DAIN Architecture** Table 2 shows the mean classification accuracy of the different three branch combination methods depicted in Figure 5. Inputs are single view images ( $I_v$ ) and single view differential images ( $I_\delta$ ). Combining the two streams at the final prediction layer (77% accuracy) is compared with the intermediate layer combination (74.8%) or the hybrid approach in Figure 5(c) (79.4%) which we choose as the differential angular imaging network. The combination method used is Sum and the feature maps are obtained from Conv5 layers after ReLU.

**DAIN Recognition Performance** We evaluate DAIN recognition performance for single view input (and differential image) and for multiview input from the GTOS database. Additionally, we compare the results to recognition using a standard CNN without a differential image stream. For all multiview experimental results we choose the number of viewpoints  $N = 4$ , separated by  $10^\circ$  with the starting viewpoint chosen at random (and the corresponding differential input). Table 3 shows the resulting recog-

nition rates (with standard deviation over 5 splits shown as a subscript). The first three rows shows the accuracy *without* differential angular imaging, using both single view and multiview input. Notice the recognition performance for these non-DAIN results are generally lower than the DAIN recognition rates in the rest of the table. The middle three rows show the recognition results for single view DAIN. For combining feature maps we evaluate both Sum and Max which have comparable results. Notice that single view DAIN achieves better recognition accuracy than multiview CNN with voting (79.4% vs. 78.1%). This is an important result indicating the power of using the differential image. Instead of four viewpoints separated by  $10^\circ$  a single viewpoint and its differential image achieves a better recognition. These results provide design cues for building imaging systems tailored to material recognition. We also evaluate weather using inputs from the two viewpoints directly (i.e.  $I_v$  and  $I_{v+\delta}$ ) is comparable to using  $I_v$  and the differential image  $I_\delta$ . Interestingly, the differential image as input has an advantage (79.4% over 77.5%). The last three rows of Table 3 show that recognition performance using multiview DAIN beats the performance of both single view DAIN and CNN methods with no differential image stream. We evaluate different ways to combine the multiview image set including voting, pooling, and the 3D filter+pooling illustrated in Figure 6.

The CNN module of our DAIN network can be replaced by other state-of-the-art deep learning methods to further improve results. To demonstrate this, we change the CNN module in a single view DAIN (Sum) (with inputs  $I_v$ ,  $I_\delta$ ) to ImageNet pre-trained ResNet-50 model [19] on split1. Combining feature maps generated from the Res4 layer (the fourth residual unit) after ReLU with training batch size 196, recognition rate improves from 77.5% to 83.0%.

Table 4 shows the recognition rates for multiview DAIN that outperforms three other multi-view classification method: FV+CNN [6], FV-N+CNN+N<sub>3D</sub> [13], and MVCNN

Method	Final Layer Combination	Intermediate Layer Combination	DAIN
Accuracy	77.0 $\pm$ 2.5	74.8 $\pm$ 3.4	79.4 $\pm$ 3.4

Table 2: Comparison of accuracy from different two stream methods as shown in Figure 5. The feature-map combination method for (b) and (c) is Sum at Conv5 layers after ReLU. The reported result is the mean accuracy and the subscript shows the standard deviation over 5 splits of the data. Notice that the architecture in (c) gives the best performance and is chosen for the differential angular imaging network (DAIN).

[39]. The table shows recognition rates for a single split of the GTOS database with images resized to  $240 \times 240$ . All experiments are based on the same pre-trained VGG-M model. We use the same fine-tuning and training procedure as in the MVCNN [39] experiment. For FV-N+CNN+N<sub>3D</sub> applied to GTOS, 10 samples (out of 606) failed to get geometry information by the method provided in [13] and we removed these samples from the experiment. The patch size in [13] is  $100 \times 100$ , but the accuracy for this patch size for GTOS was only 43%, so we use  $240 \times 240$ . We implement FV-N+CNN+N<sub>3D</sub> with linear mapping instead of homogeneous kernel map [41] for SVM training to save memory with this larger patch size.

**DAIN on 4D Light Field Dataset** We tested our multi-view DAIN (Sum + pooling) method on a recent 4D light field (Lytro) dataset [44]. ResNet-50 is used as the CNN module. The recognition accuracy with full images on 5 splits is 83.0 $\pm$ 2.1. Note that a subset of the lightfield data is used to mimic the differential imaging process, so these results should not be interpreted as a comparison of our algorithm to [44].

The Lytro dataset has  $N = 49$  views, from the  $7 \times 7$  lenslet array, where each lenslet corresponds to a different viewing direction. Using  $(i, j)$  as an index into this array, we employ the viewpoints indexed by (4, 1), (4, 3), (4, 5), (4, 7) as the 4 views in multiview DAIN. We use the viewpoint indexed by (3, 1), (5, 3), (3, 5), (5, 7) as the corresponding differential views. This is an approximation of multiview DAIN; the lightfield dataset does not capture the range of viewing angles to exactly emulate multiple viewpoints and small angle variations of these viewpoints. Instead of using all  $N = 49$  viewpoints as in [44], we generate comparable recognition accuracy by only 8 viewpoints.

## 7. Conclusion

In summary, there are three main contributions of this work: 1) Differential Angular Imaging for a sparse spatial

Method	First input	Second input	Accuracy
single view CNN	$I_v$	-	74.3 $\pm$ 2.8
multiview CNN, voting	$I_v$	-	78.1 $\pm$ 2.4
multiview CNN, 3D filter	$I_v$	-	74.8 $\pm$ 3.2
single view DAIN (Sum)	$I_v$	$I_{v+\delta}$	77.5 $\pm$ 2.7
single view DAIN (Sum)	$I_v$	$I_\delta$	79.4 $\pm$ 3.4
single view DAIN (Max)	$I_v$	$I_\delta$	79.0 $\pm$ 1.8
multiview DAIN (Sum/voting)	$I_v$	$I_\delta$	80.0 $\pm$ 2.1
multiview DAIN (Sum/pooling)	$I_v$	$I_\delta$	81.2 $\pm$ 1.7
multiview DAIN (3D filter/pooling)	$I_v$	$I_\delta$	81.1 $\pm$ 1.5

Table 3: Results comparing performance of standard CNN recognition without angular differential imaging (first three rows) to our single-view DAIN (middle three rows) and our multi-view DAIN (bottom three rows).  $I_v$  denotes the image from viewpoint  $v$ ,  $I_{v+\delta}$  is the image obtained from viewpoint  $v+\delta$ , and  $I_\delta = I_v - I_{v+\delta}$  is the differential image. The differential angular imaging network (DAIN) has superior performance over CNN even when comparing single view DAIN to multiview CNN. Multiview DAIN provides the best recognition rates.

Architecture	Accuracy
FV+CNN [6]	75.4%
FV-N+CNN+N <sub>3D</sub> [13]	58.3%
MVCNN [39]	78.1%
<b>multiview DAIN (3D filter), pooling</b>	<b>81.4%</b>

Table 4: Comparison with the state of art algorithms on GTOS dataset. Notice that our method, multiview DAIN, achieves the best recognition accuracy.

distribution of angular gradients that provides key cues for material recognition; 2) The GTOS Dataset with ground terrain imaged by systematic in-scene measurement of partial reflectance instead of in-lab reflectance measurements. The database contains 34,243 images with 40 surface classes, 18 viewing directions, 4 illumination conditions, 3 exposure settings per sample and several instances/samples per class. 3) We develop and evaluate an architecture for using differential angular imaging, showing superior results for differential inputs as compared to original images. Our work in measuring and modeling outdoor surfaces has important implications for applications such as robot navigation (determining control parameters based on current ground terrain) and automatic driving (determining road conditions by partial real time reflectance measurements). We believe our database and methods will provide a sound foundation for in-depth studies on material recognition in the wild.

## Acknowledgment

This work was supported by National Science Foundation award IIS-1421134. A GPU used for this research was donated by the NVIDIA Corporation. Thanks to Di Zhu, Hansi Liu, Lingyi Xu, and Yueyang Chen for help with data collection.



## References

- [1] S. Bell, P. Upchurch, N. Snavely, and K. Bala. Material recognition in the wild with the materials in context database. *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 3, 4
- [2] M. Chandraker. The information available to a moving observer on shape with unknown, isotropic brdfs. *IEEE transactions on pattern analysis and machine intelligence*, 38(7):1283–1297, 2016. 2
- [3] M. Chandraker, J. Bai, and R. Ramamoorthi. On differential photometric reconstruction for unknown, isotropic brdfs. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2941–2955, 2013. 2
- [4] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014. 3, 5
- [5] G. Choe, S. G. Narasimhan, and I. S. Kweon. Simultaneous estimation of near ir brdf and fine-scale surface geometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 3
- [6] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014. 7, 8
- [7] M. Cimpoi, S. Maji, and A. Vedaldi. Deep filter banks for texture recognition and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3828–3836, 2015. 2, 3
- [8] O. G. Cula and K. J. Dana. Recognition methods for 3d textured surfaces. In *Proceedings of SPIE conference on human vision and electronic imaging VI*, number 209-220, page 3, 2001. 3
- [9] K. Dana and J. Wang. Device for convenient measurement of spatially varying bidirectional reflectance. *Journal of the Optical Society of America A*, 21:pp. 1–12, January 2004. 3
- [10] K. J. Dana. Capturing computational appearance: More than meets the eye. *IEEE Signal Processing Magazine*, 33(5):70–80, 2016. 1
- [11] K. J. Dana, B. Van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics (TOG)*, 18(1):1–34, 1999. 3, 4
- [12] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '00*, pages 145–156, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co. 3
- [13] J. DeGol, M. Golparvar-Fard, and D. Hoiem. Geometry-informed material recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1554–1562, 2016. 4, 7, 8
- [14] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009. 5
- [15] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. *arXiv preprint arXiv:1604.06573*, 2016. 3
- [16] J. Filip and R. Vávra. Template-based sampling of anisotropic brdfs. *Comput. Graph. Forum*, 33(7):91–99, Oct. 2014. 4
- [17] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 3
- [18] E. Hayman, B. Caputo, M. Fritz, and J.-O. Eklundh. On the significance of real-world conditions for material classification. In *European conference on computer vision*, pages 253–266. Springer, 2004. 3, 4
- [19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. 3, 7
- [20] D. Hu, L. Bo, and X. Ren. Toward Robust Material Recognition for Everyday Objects. In *BMVC*, pages 48.1–48.11, 2011. 2
- [21] C. Kampouris, S. Zafeiriou, A. Ghosh, and S. Malassiotis. *Fine-Grained Material Classification Using Micro-geometry and Reflectance*, pages 778–792. Springer International Publishing, Cham, 2016. 4
- [22] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128–3137, 2015. 3
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 3
- [24] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textures. *International journal of computer vision*, 43(1):29–44, 2001. 3
- [25] M. Levoy and P. Hanrahan. Light Field Rendering. In *Computer Graphics Proceedings, ACM SIGGRAPH 96*, pages 31–42, Aug. 1996. 3
- [26] T.-Y. Lin, A. RoyChowdhury, and S. Maji. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1449–1457, 2015. 3
- [27] C. Liu and J. Gu. Discriminative Illumination: Per-Pixel Classification of Raw Materials Based on Optimal Projections of Spectral BRDF. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(1):86–98, January 2014. 2, 3
- [28] C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz. Exploring Features in a Bayesian Framework for Material Recognition. In *CVPR*, pages 239–246, 2010. 2
- [29] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003. 4
- [30] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis. Geometric Considerations and Nomenclature for Reflectance. National Bureau of Standards (US), 1977. 3

- [31] G. Oxholm, P. Bariya, and K. Nishino. The Scale of Geometric Texture. In *European Conference on Computer Vision*, volume I, pages 58–71, 2012. 4
- [32] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 3
- [33] N. Salamati, C. Fredembach, and S. Süssstrunk. Material Classification using Color and NIR Images. In *IS&T/SID Color Imaging Conference*, 2009. 2
- [34] G. Schwartz and K. Nishino. Visual Material Traits: Recognizing Per-Pixel Material Context. In *IEEE Color and Photometry in Computer Vision Workshop*, 2013. 2, 3
- [35] G. Schwartz and K. Nishino. Automatically discovering local visual material attributes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2, 3
- [36] L. Sharan, R. Rosenholtz, and E. Adelson. Material perception: What can you see in a brief glance? *Journal of Vision*, 9(8):784–784, 2009. 3, 4
- [37] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, pages 568–576, 2014. 3
- [38] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 3
- [39] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 945–953, 2015. 8
- [40] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4489–4497. IEEE, 2015. 6
- [41] A. Vedaldi and A. Zisserman. Efficient additive kernels via explicit feature maps. *IEEE transactions on pattern analysis and machine intelligence*, 34(3):480–492, 2012. 8
- [42] O. Wang, P. Gunawardane, S. Scher, and J. Davis. Material classification using brdf slices. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2805–2811, 2009. 2, 3
- [43] T.-C. Wang, M. Chandraker, A. A. Efros, and R. Ramamoorthi. Svbrdf-invariant shape and reflectance estimation from light-field cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5451–5459, 2016. 2
- [44] T.-C. Wang, J.-Y. Zhu, E. Hiroaki, M. Chandraker, A. A. Efros, and R. Ramamoorthi. A 4d light-field dataset and cnn architectures for material recognition. In *European Conference on Computer Vision*, pages 121–138. Springer, 2016. 2, 3, 8
- [45] G. Ward. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH 92*, pages 265–272, 1992. 3
- [46] M. Weinmann, J. Gall, and R. Klein. Material classification based on training data synthesized using a btf database. In *European Conference on Computer Vision*, pages 156–171. Springer, 2014. 3, 4
- [47] H. Zhang, K. Dana, and K. Nishino. Reflectance hashing for material recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [48] H. Zhang, K. Dana, and K. Nishino. Reflectance hashing for material recognition. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 371–380, 2015. 2, 3
- [49] H. Zhang, K. Nishino, and K. Dana. Friction from Reflectance: Deep Reflectance Codes for Predicting Physical Surface Properties from One-Shot In-Field Reflectance. In *European Conference on Computer Vision*, pages 808–824, 2016. 3
- [50] H. Zhang, J. Xue, and K. Dana. Deep ten: Texture encoding network. *arXiv preprint arXiv:1612.02844*, 2016. 3