


Researchers use AI to identify similar materials in images

May 23 2023, by Adam Zewe



Our synthetic dataset. We show samples from our synthetic training dataset. Top: three sample rendered images. Bottom: The corresponding material id maps, ids are mapped to random colors. Credit:

https://prafullsharma.net/materialistic/materialistic_camera_ready.pdf

A robot manipulating objects while, say, working in a kitchen, will benefit from understanding which items are composed of the same materials. With this knowledge, the robot would know to exert a similar amount of force whether it picks up a small pat of butter from a shadowy corner of the counter or an entire stick from inside the brightly lit fridge.

Identifying objects in a scene that are composed of the same material, known as material selection, is an especially challenging problem for machines because a material's appearance can vary drastically based on the shape of the object or lighting conditions.

Scientists at MIT and Adobe Research have taken a step toward solving this challenge. They developed a technique that can identify all pixels in an image representing a given material, which is shown in a [pixel](#) selected by the user.

The method is accurate even when objects have varying shapes and sizes, and the machine-learning [model](#) they developed isn't tricked by shadows or lighting conditions that can make the same material appear different.

Although they trained their model using only "synthetic" data, which are created by a computer that modifies 3D scenes to produce many varying images, the system works effectively on real indoor and outdoor scenes it has never seen before. The approach can also be used for videos; once the user identifies a pixel in the first frame, the model can identify objects made from the same material throughout the rest of the video.

In addition to applications in scene understanding for robotics, this method could be used for image editing or incorporated into computational systems that deduce the parameters of [materials](#) in images. It could also be utilized for material-based web recommendation

systems. (Perhaps a shopper is searching for clothing made from a particular type of fabric, for example.)

"Knowing what material you are interacting with is often quite important. Although two objects may look similar, they can have different material properties. Our method can facilitate the selection of all the other pixels in an image that are made from the same material," says Prafull Sharma, an [electrical engineering](#) and computer science graduate student and lead author of a paper on this technique.

Sharma's co-authors include Julien Philip and Michael Gharbi, [research scientists](#) at Adobe Research; and senior authors William T. Freeman, the Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL); Frédo Durand, a professor of electrical engineering and [computer science](#) and a member of CSAIL; and Valentin Deschaintre, a research scientist at Adobe Research. The research will be presented at the SIGGRAPH 2023 conference.

A new approach

Existing methods for material selection struggle to accurately identify all pixels representing the same material. For instance, some methods focus on entire objects, but one object can be composed of multiple materials, like a chair with wooden arms and a leather seat. Other methods may utilize a predetermined set of materials, but these often have broad labels like "wood," despite the fact that there are thousands of varieties of wood.

Instead, Sharma and his collaborators developed a machine-learning approach that dynamically evaluates all pixels in an image to determine the material similarities between a pixel the user selects and all other regions of the image. If an image contains a table and two chairs, and the

chair legs and tabletop are made of the same type of wood, their model could accurately identify those similar regions.

Before the researchers could develop an AI method to learn how to select similar materials, they had to overcome a few hurdles. First, no existing dataset contained materials that were labeled finely enough to train their machine-learning model. The researchers rendered their own synthetic dataset of indoor scenes, which included 50,000 images and more than 16,000 materials randomly applied to each object.

"We wanted a dataset where each individual type of material is marked independently," Sharma says.

Synthetic dataset in hand, they trained a [machine-learning model](#) for the task of identifying similar materials in real images—but it failed. The researchers realized distribution shift was to blame. This occurs when a model is trained on [synthetic data](#), but it fails when tested on real-world data that can be very different from the training set.

To solve this problem, they built their model on top of a pretrained computer vision model, which has seen millions of real images. They utilized the prior knowledge of that model by leveraging the visual features it had already learned.

"In machine learning, when you are using a neural network, usually it is learning the representation and the process of solving the task together. We have disentangled this. The pretrained model gives us the representation, then our [neural network](#) just focuses on solving the task," he says.

Solving for similarity

The researchers' model transforms the generic, pretrained visual features

into material-specific features, and it does this in a way that is robust to object shapes or varied [lighting conditions](#).

The model can then compute a material similarity score for every pixel in the image. When a user clicks a pixel, the model figures out how close in appearance every other pixel is to the query. It produces a map where each pixel is ranked on a scale from 0 to 1 for similarity.

"The user just clicks one pixel and then the model will automatically select all regions that have the same material," he says.

Since the model is outputting a similarity score for each pixel, the user can fine-tune the results by setting a threshold, such as 90 percent similarity, and receive a map of the image with those regions highlighted. The method also works for cross-image selection—the user can select a pixel in one image and find the same material in a separate image.

During experiments, the researchers found that their model could predict regions of an image that contained the same material more accurately than other methods. When they measured how well the prediction compared to ground truth, meaning the actual areas of the image that are comprised of the same material, their model matched up with about 92 percent accuracy.

In the future, they want to enhance the model so it can better capture fine details of the objects in an image, which would boost the accuracy of their approach.

"Rich materials contribute to the functionality and beauty of the world we live in. But computer vision algorithms typically overlook materials, focusing heavily on objects instead. This paper makes an important contribution in recognizing materials in images and video across a broad

range of challenging conditions," says Kavita Bala, Dean of the Cornell Bowers College of Computing and Information Science and Professor of Computer Science, who was not involved with this work. "This technology can be very useful to end consumers and designers alike. For example, a home owner can envision how expensive choices like reupholstering a couch, or changing the carpeting in a room, might appear, and can be more confident in their design choices based on these visualizations."

More information: Materialistic: Selecting Similar Materials in Images. prafullsharma.net/materialistic_camera_ready.pdf

Provided by Massachusetts Institute of Technology

Citation: Researchers use AI to identify similar materials in images (2023, May 23) retrieved 25 April 2024 from <https://techxplore.com/news/2023-05-ai-similar-materials-images.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--