

Convolutional Network Analysis of Optical Micrographs for Liquid Crystal Sensors

Published as part of *The Journal of Physical Chemistry* virtual special issue "Machine Learning in Physical Chemistry".

Alexander D. Smith, Nicholas Abbott, and Victor M. Zavala*



Cite This: *J. Phys. Chem. C* 2020, 124, 15152–15161



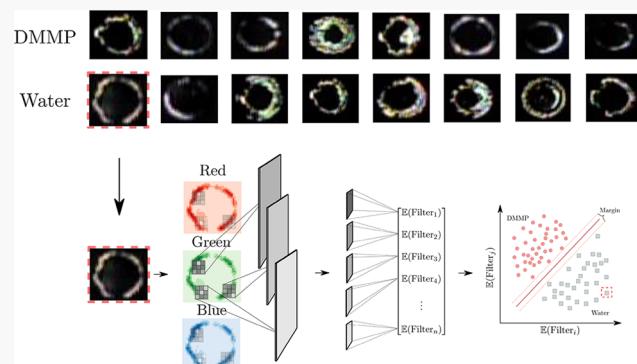
Read Online

ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: We provide an in-depth convolutional neural network (CNN) analysis of optical responses of liquid crystals (LCs) when exposed to different chemical environments. Our aim is to identify informative features that can be used to construct automated LC-based chemical sensors and shed some light on the underlying phenomenon that governs and distinguishes LC responses. Previous work demonstrated that, by using features extracted from AlexNet, grayscale micrographs of different LC responses can be classified with an accuracy of 99%. Reaching such high levels of accuracy, however, required the use of a large number of features (on the order of thousands), which was computationally intensive and clouded the physical interpretability of the dominant features. To address these issues, here we report a study on the effectiveness of using features extracted from color micrographs using VGG16, which is a more compact CNN than Alexnet. Our analysis reveals that features extracted from the first and second convolutional layers of VGG16 are sufficient to achieve a perfect classification accuracy while reducing the number of features to less than 100. The number of features is further reduced to 10 via recursive elimination with a minimal loss in classification accuracy (5–10%). This reduction procedure reveals that differences in spatial color patterns are developed within seconds in the LC response. From this, we conclude that hue distributions provide an informative set of features that can be used to characterize LC sensor responses. We also hypothesize that differences in the spatial correlation length of LC textures detected by VGG16 with DMMP and water likely reflect differences in the anchoring energy of the LC on the surface of the sensor. Our results hint at fresh approaches for the design of LC-based sensors based on the characterization of spontaneous fluctuations in the orientation (as opposed to changes in time-average orientations reported in the literature).



INTRODUCTION

Liquid crystals (LCs) provide a versatile platform for the sensing of air contaminants (chemical sensing)^{1,2} and for the sensing of heat transfer and shear stress (mechanical sensing).³ In the context of chemical sensing, LC sensors can be designed to change their orientational ordering and optical birefringence upon exposure of the LC to a certain targeted chemical environment. For instance, an LC sensor can be prepared by supporting a thin LC film (thickness of micrometers) on a chemically functionalized surface. Typically, the molecules within the LC film (the mesogen) bind to the surface and assume a homeotropic (perpendicular) orientation that provides an initial optical signal. Subsequent exposure of the LC film to an analyte leads to the diffusive transport of the analyte through the LC phase and the displacement of the mesogen at the surface, triggering rich space–time optical responses (Figure 1). The response time of the spatial-average

brightness of the optical signal has been shown to be strongly correlated to the differential binding energy between the analyte and mesogen to the surface. The physicochemical principles of LC chemical sensors are explained in detail elsewhere.¹

A primary challenge for the development of LC sensors (as in other sensing technologies) is their potential sensitivity to interfering chemical species. For instance, LC sensors designed for the detection of dimethyl methylphosphonate (DMMP),

Received: March 4, 2020

Revised: June 5, 2020

Published: June 9, 2020



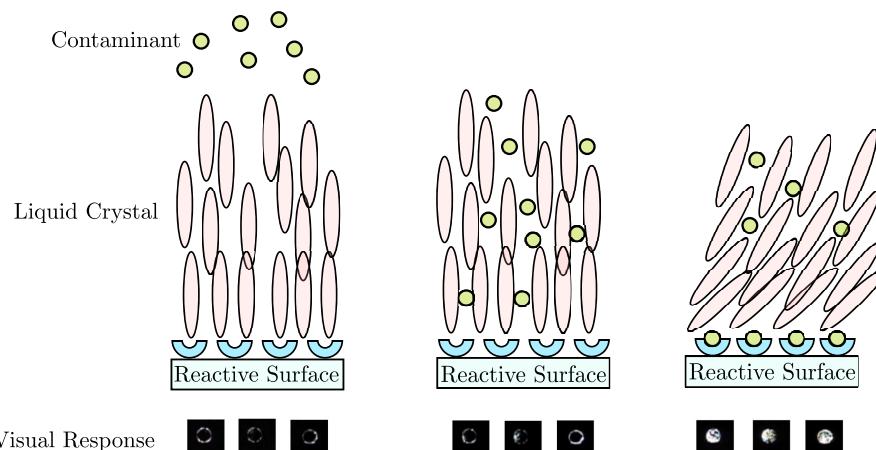


Figure 1. Working design principles of a liquid crystal chemical sensor.

$\text{CH}_3\text{PO}(\text{OCH}_3)_2$, might exhibit similar optical responses when exposed to humid nitrogen.⁴ Moreover, LC responses can also be slow, as these require the diffusion of the air contaminant through the LC film and the displacement of the mesogen at the surface. Sluggish responses limit the applicability of the LC sensor (e.g., when detecting highly toxic chemicals). These issues are illustrated in the experimental responses shown in Figure 2. Although the selectivity of LC sensors can be

classification.^{5,6} The goal of a classification strategy is to separate different images by using numerical features (descriptors) that characterize such images. Features are projections of the original image into an information space that seek to best summarize or describe an image (features are characteristic patterns of the image). Certain features can be strongly correlated to physical phenomena that govern a system; for instance, image features such as textures are often correlated to structural properties of materials.⁷ Interestingly, informative features that capture multiscale spatial patterns can be extracted from CNNs that have been pretrained using generic images (that are not directly related with the application at hand). Such features can then be used in an external classification engine such as a fully connected network, logistic regression, or support vector machine. For instance, in the work of Kawahara et al.,⁶ the pretrained CNN Alexnet⁸ was used to classify skin lesions. In the work of Ling et al.,⁹ textures extracted from the pretrained CNN VGG16¹⁰ are used to predict material properties. The principle behind the exploitation of pretrained CNNs is known as transfer learning.¹¹

Cao and co-workers recently used Alexnet to characterize optical LC responses (as grayscale images) and demonstrated that spatial features of the LC response can be used to discern the chemical environment.¹² Specifically, the authors demonstrated that spatial features extracted from the deep layers of AlexNet can be used to achieve classification accuracy levels of 99%. Notably, they also observed that snapshots taken within three seconds of exposing the LC are sufficient to classify the environment (either DMMP or humid nitrogen). Unfortunately, reaching such high levels of accuracy required an extremely large number of features (on the order of thousands), which resulted in computational issues and clouded the physical interpretability of the dominant features. In particular, features extracted from deep CNN layers, while informative, are difficult to interpret.

In this work, we extend the results of Cao and co-workers by analyzing LC response features extracted from VGG16, which is a CNN that embeds a smaller set of convolutional filters than that of Alexnet. Moreover, in the current study, we use RGB (color) images directly (in previous work grayscale images were used). Our findings demonstrate that features extracted from the first and second convolutional layers of VGG16 allow for a perfect classification accuracy for the same data set studied by Cao and co-workers while reducing the

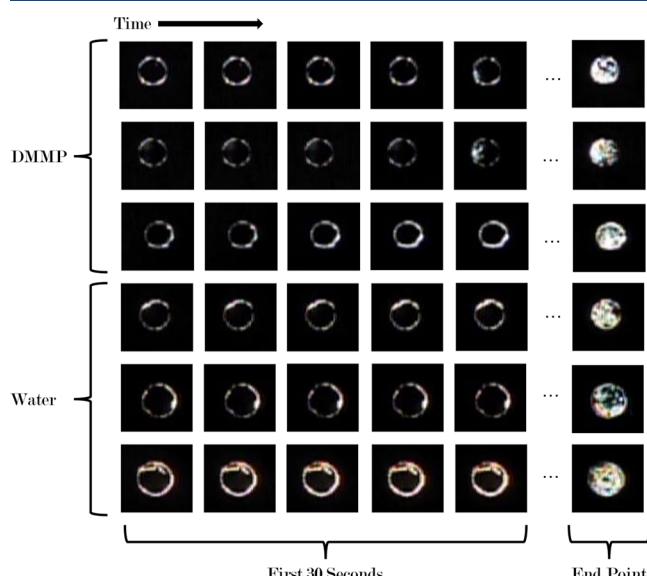


Figure 2. Optical responses of liquid crystals under gaseous N_2 -water (30% relative humidity) and N_2 -DMMP (10 ppm) environments. LCs were deposited into microwells with a diameter of 3 mm to enable high-throughput data collection. LC responses were recorded at room temperature.

optimized by chemical design to largely eliminate the effects of humidity, a natural step is to determine whether or not one can unravel hidden patterns in the optical responses that can help discern between chemical species. The identification of such patterns can also help reduce detection times and simplify the design of LC sensors.

Machine learning techniques are actively being used for pattern recognition in diverse branches of science and engineering. Specifically, convolutional neural networks (CNNs) have been used for brain tumor and skin lesion

number of features to approximately one hundred. We demonstrate that the number of features can be further reduced to 10 via recursive feature elimination with minimal losses in sensor accuracy. This feature reduction procedure reveals that complex spatial color patterns are developed within seconds in the LC response, which leads us to hypothesize that differences in spontaneous fluctuations in the LC tilt orientation (angle) play a key role in sensor selectivity and responsiveness. Our analysis also reveals that hue distributions provide an effective set of features to characterize LC responses.

METHODS

Experimental Methods. As described in our prior publication,¹² we recorded six videos that show the response of LCs to N_2 –DMMP at 10 ppm (the length of each video ranges from 4 to 13 min) and six videos that show the response of LCs to N_2 –water (the length of each video ranges from 7 to 30 min). The experimental system is sketched in Figure 3.

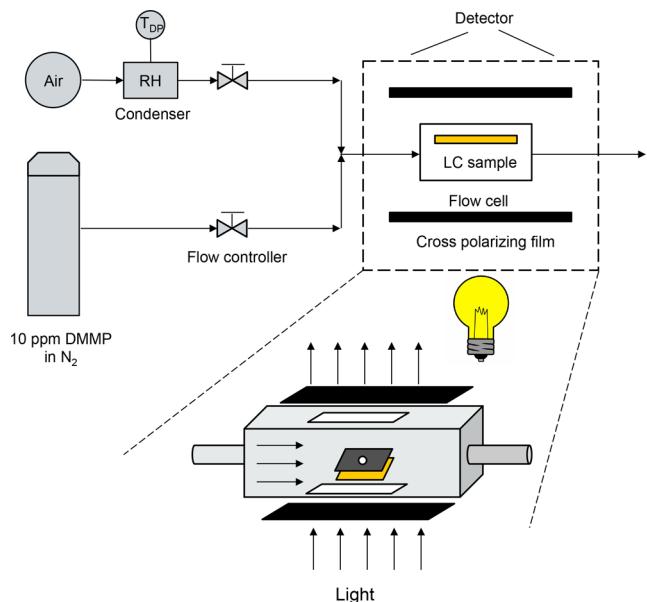


Figure 3. Sketch of experimental system used for collecting LC response data. Reproduced with permission from ref 12. Copyright 2018, American Chemical Society.

Each video tracks the dynamic evolution of multiple independent microwells (the total number of microwells recorded was 391). We captured a frame (micrograph) from each video every 3.3 s. We split each frame into several images, each containing a single microwell at a specific time. The total number of microwell images (snapshots) generated was 75 081, and each image is resized to 60 pixels \times 60 pixels (see Figure 2 for some example micrographs). The experimental procedure followed to obtain the LC response data was the following:

- Formation of thin films of LC supported on metal–salt-decorated surfaces: Aluminum perchlorate salts (50 μ L, 10 mM) in ethanolic solution were deposited by spin-coating (3000 rpm for 30 s) onto the glass surfaces at the bottom of the polymeric microwells. Next, 2 μ L of 5CB (4-cyano-4'-pentylbiphenyl) was deposited into the polymeric microwells¹³ with a depth of 5 μ m using a

micropipette. The excess LC was removed from the array by wicking it into a microcapillary.

- Optical characterization of LC films: The optical appearance of the LC was characterized by using an Olympus BX-60 polarizing-light microscope in transmission mode (Olympus, Japan). Conoscopic imaging of the LC films was performed by inserting a Bertran lens into the optical path of a polarized-light microscope to confirm the homeotropic orientation.¹⁴

- Ordering transitions induced by DMMP and humid N_2 : The LC-filled microwells were exposed to a stream of dry N_2 containing DMMP (10 ppmv) within a flow cell¹⁵ with glass windows that permitted the characterization of the optical appearance of the LC using a polarized optical microscope. The gas containing DMMP was delivered to the flow cell at 300 mL/min by using a rotameter (Aalborg Instruments and Control, Orangeburg, NY). For experiments performed to evaluate the response of the LCs to water vapor, nitrogen containing 30% relative humidity was delivered to the flow cell at 300 mL/min with the same rotameter (we call this mixture N_2 –water). The optical appearance of the LC film was recorded using an Olympus camera (Olympus C2040Zoom, Melville, NY) and WinTV software (Hauppauge, NY).

Computational Methods. In this section, we summarize the machine learning methods used to analyze optical micrographs of LCs. We focus on classifying whether an LC sensor has been exposed to DMMP or humid air (we call this water, for convenience). In other words, our framework is focused on binary classification. We use the same data set reported by Cao and co-workers¹² but focus on patterns developed within the first 30 s of the LC response. Details regarding the experimental system and data preparation methods can be found in¹².

In summary, the data set analyzed was obtained from six videos that show the response of LCs to a gaseous stream of N_2 containing 10 ppm DMMP and six videos that show the response of LCs to a gaseous stream of N_2 containing 30% relative humidity (both at room temperature). Each video tracks the dynamic evolution of multiple independent microwells (the total number of microwells recorded was 391). We captured a frame (a micrograph) from each video every 3.3 s. We split each frame into several images, each containing a single microwell at a specific time. The total number of microwell snapshots generated was 75 081 (the data set analyzed is extensive).

Examples of snapshot sequences collected during the microwell responses are presented in Figure 2. Our machine learning analysis treats snapshots as time-independent; this type of analysis is more challenging than the analysis of time-dependent sequences, and it is more desirable from a sensor design perspective; this is because we want to detect a contaminant by ignoring its response history. Specifically, our aim is to show that machine learning techniques can detect a contaminant by just looking at a snapshot at any time (by exploiting the spatial pattern of the response).

Classification. In an ideal setting, in which an image can be characterized using highly informative features, classification can be performed using a linear hyperplane, where the dimension of the hyperplane is equal to the number of features minus one. For instance, if an image can be

characterized using two features, the hyperplane will be a line. This hyperplane provides a decision boundary under which every image on one side is considered a member of one class and every image on the opposite side is considered a member of the contrasting class. In most settings, these classes are provided a numerical label of +1 or -1. In our setting, water is considered the +1 class and DMMP is considered the -1 class.

The classification engine used for the LC data set is a linear support vector machine (LSVM), which is trained using image features extracted from the CNN VGG16. An illustration of the LSVM method is presented in Figure 4. LSVM is a

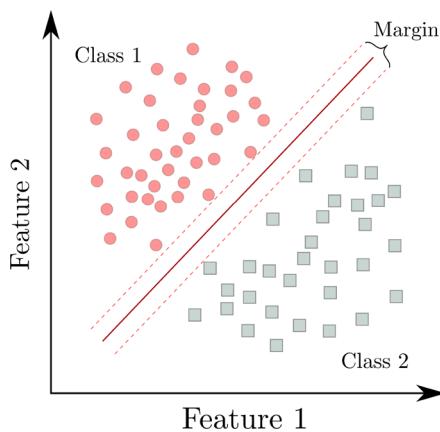


Figure 4. Illustration of a linear support vector machine.

classification method that builds a linear decision boundary between observations. This is done by finding a hyperplane that maximizes the margin between the set of closest images to the hyperplane (known as the support vectors) and the hyperplane itself. The hyperplane is a weighted linear combination of all the CNN features representing each observation. The magnitude of each feature weight represents its relative importance (a proxy for information content); in other words, a feature that is highly informative (explains differences in the images well) will tend to have a large weight while a noninformative feature will tend to have a small weight. The images that are closest to the margin are the most difficult to classify (difficult to distinguish), while the ones that are farthest away from the margin can be easily classified (easy to distinguish). The support vectors are the images that define the separation boundary.

The identification of relevant features can be achieved by penalizing the l_1 norm of the weights of the LSVM classifier. This penalization term seeks to sparsify the weight vector (have few nonzero entries). Consequently, a penalized LSVM classifier is tasked with finding a separating hyperplane that best classifies the images, and the classifier is also required to

complete this task with a minimal number of features (this set of features are interpreted as the ones that provide most information). The mathematical formulation of the LSVM problem is

$$\min_{w_0, w} \sum_{i=1}^n \left[1 - y_i \left(w_0 + \sum_{j=1}^q w_j x_i \right) \right] + \lambda \|w\|_1 \quad (1)$$

Here, n is the number of images (observations). m is the image feature dimension. $w \in \mathbb{R}^m$ represents the feature weights. $x_i \in \mathbb{R}^m$ are the features of observation i . $y_i \in \{-1, +1\}$ represents the label for observation i . $\lambda \in \mathbb{R}_+$ is a hyperparameter for the penalization of the l_1 norm.¹⁶ The solution of eq 1 is often called the training phase and the images used for its solution are often called the training set. Once the classifier has been trained, one uses the optimal hyperplane weights w^* identified in the training phase to predict the label of a new image that is not in the original training set. The new images are known as the test (validation) data set. This process is repeated five times, each time with a new training and validation set (five-fold validation). This allows for a robust testing of the effectiveness of the classification model on the entire data set.

Feature Extraction. In order to train the LSVM classifier, we first need to identify features that best explain each image. Cao and co-workers previously used Alexnet to conduct feature extraction from LC micrographs. Alexnet is a CNN that has been pretrained using the Imagenet database.¹⁷ This database is a collection of millions of images that contains over 1000 categories. The original goal of Alexnet was to work as a classifier;⁸ however, one can also use features extracted by Alexnet to train an external classifier such as an LSVM (transfer learning). This approach avoids retraining the CNN, which can be highly computationally expensive. Cao and co-workers demonstrated that the transfer learning approach can be effectively used to classify optical micrographs of LCs using LSVM; their analysis, however, used over 5000 features to explain each micrograph. Moreover, in their approach, the micrographs were transformed into grayscale images; as we will see, this transformation leads to significant losses of information and hides the physical LC behavior.

In this work, we consider a different pretrained CNN that we hypothesize may be better suited to our given application. We sought to merge our understanding of working principles of CNNs with our knowledge of the physical behavior of LCs. A primary consideration is the length scale that characterizes the LC response. Specifically, we know that nematic ordering and interfacial interactions within LCs give rise to optical patterns of orientation on the micrometer scale; because of this, the patterns created by the LCs need to be captured with

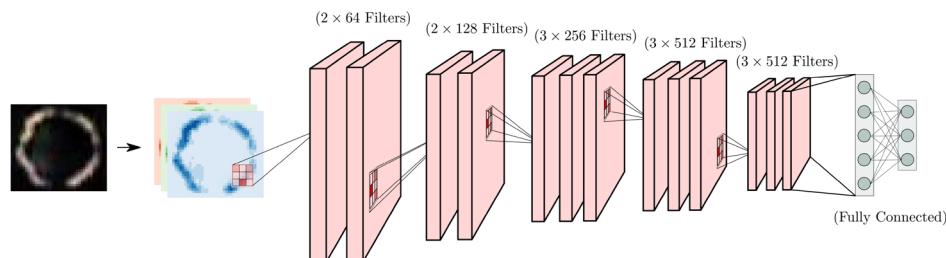


Figure 5. Schematic of VGG16 architecture.

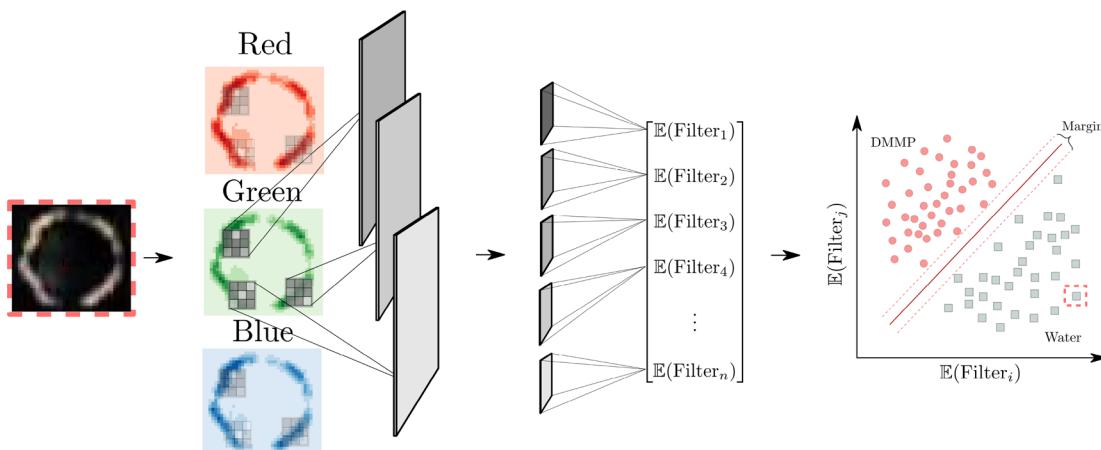


Figure 6. Schematic of feature extraction and classification framework ($E(\cdot)$ represents spatial average).

a small observation lense. Moreover, the interference colors created by the LCs are an indicator of their tilt angles (orientation), and thus, the CNN selected should be trained using RGB images directly (as opposed to grayscale images). A CNN that fits these requirements is VGG16, which has been pretrained by the visual geometry group at Oxford.¹⁰ The VGG16 CNN has been trained on the Imagenet database. The structure and optimal weight values for the trained VGG16 network are freely available through the Keras software and are what is used during feature extraction.¹⁸ VGG16 utilizes the smallest possible convolutional filter size (3×3), which should be best for capturing small-scale structural patterns in images. Moreover, VGG16 is a much shallower CNN than AlexNet, and thus, its features are easier to interpret. A simplified representation of the VGG16 architecture is shown in Figure 5.

The basic idea behind feature extraction using a CNN such as VGG16 is to reduce a given input image into a small set of numerical values that can be used to best summarize and classify the image. Each image is represented by a set of input channels, and each set of input channels is expressed as a two-dimensional pixel field (a matrix). The input channels are typically the red, green, and blue (RGB) channels of an image. Image reduction is performed through a sequence of matrix convolution operations in which spatial information is extracted from the image using filters (matrices with specific patterns). Subsequent convolutions compress this image to the point where a decision (e.g., classification or regression) can be made, which is represented by the fully connected layer in Figure 5. More details on this procedure can be found in ref 10.

A convolution is a manipulation of an image matrix with a filter matrix. Specifically, a convolution is the process of finding the extent to which a given pattern (defined by a convolutional filter) is present within a neighborhood of an image (and repeating the process by spanning all neighborhoods of the image). In other words, the convolution seeks to identify to what extent a specific spatial morphology and/or correlation structure (defined by the filter matrix) is present in the image. An example of applying a convolution filter to an image is illustrated in Figure 7. Convolutional filters provide a quantifiable approach for identifying multiple spatial structures within a given neighborhood (different filters identify different types of patterns). The larger the value of the filter output is, the more similar the given neighborhood is to the pattern that the filter is attempting to find. Optimal filter matrices that best

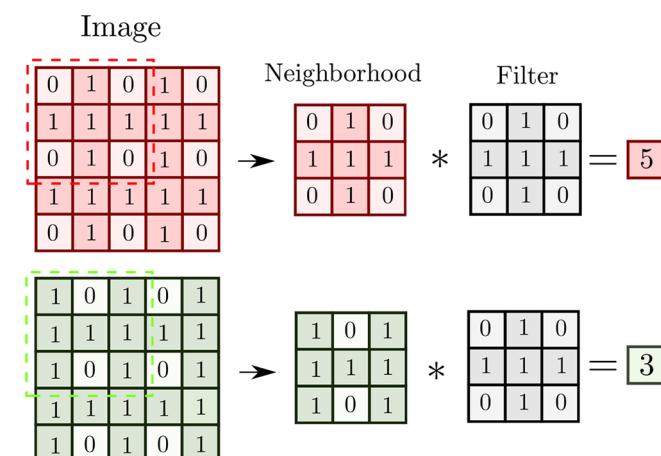


Figure 7. Illustration of the application of a convolution filter to a neighborhood of two different images.

classify a set of images (optimal patterns) can be found by training the CNN directly on the data set. Specifically, the training process aims to compute the entries of the filter matrices that best separate the images. Matrices are high dimensional objects, and as such, training a CNN involves a highly computationally expensive procedure. Filters extracted from training over a given set, however, can also be reused to search for similar patterns in a different data set. In other words, pretrained filters (preidentified spatial patterns) can be used on a different image set with the sole purpose of obtaining feature information. While the filters are not optimal for the new data set, this procedure is often effective at detecting general patterns in images, and the obtained feature information can be used in an external classifier such as LSVM.

In the example provided in Figure 7, we see that the convolutional filter is seeking to match the neighborhood to a cross-pattern, and thus, the top neighborhood has a higher output (perfect match) than that of the lower output (imperfect match). In the CNN, the matching is applied to every pixel in the image, and thus, there is a convolution value for every pixel neighborhood (resulting in a matrix of filter outputs). In our approach, the entire set of outputs for each filter are averaged and utilized as a feature for the LSVM classification. This is done in order to ensure that the features are spatially invariant. Spatial invariance allows for images that are not of a uniform size or perfectly centered to be treated as

similarly as possible. This practice also forces the classifier to seek meaningful and generalizable features associated with the sensors rather than arbitrary features based upon the location of the sensor in the given frame (thus leading to more consistency in the results).

The next decision to be made is what feature information should be extracted from VGG16. The VGG16 network has been trained to classify highly complex images, and the deepest layers have been carefully tuned to differentiate such images. The early layers of the network, however, are the most general and are easier to interpret (they are less evolved). Accordingly, in our approach, we use the outputs of the first and second convolutional blocks to inform features for LSVM classification. Feature extraction is conducted by feeding a given image into VGG16. We modified the network so that the only output it provides is information extracted from the first and second convolutional blocks. This information is extracted in the form of convolutional filter activations via convolutions.

In summary, the CNN used here provides a number of features equal to the number of convolutional filters used for each image. In our case, the total number of features reaches 192 (64 for the first block and 128 for the second block). Note that the number of features increases with the depth of the layer, which precisely reinforces our desire to focus on the first layers. A visual representation of this process for the first and second convolutional blocks is provided in Figure 6. Feature extraction and network modification were performed using Keras¹⁸ and Tensorflow.¹⁹ The VGG16 network and trained weights are made available in the Keras software, which allows for easy manipulation of the VGG16 network so that this process may be completed for any number of image sets. With the extraction of the features from the first two layers of the VGG16 network, analysis of the classification may be conducted.

RESULTS AND DISCUSSION

We now describe our findings when applying CNN techniques to analyze LC micrographs. All scripts and data needed to reproduce the results are available at https://github.com/zavalab/ML/tree/master/LC_CNN_Color.

Classification and Feature Reduction. Our machine learning (ML) framework using VGG16 features and LSVM was able to classify water and DMMP micrographs with 100% accuracy. Notably, these results were obtained for micrographs collected within 30 s of exposing the LCs to the chemical environments. This result was achieved when using all of the 128 features of the second convolutional layer. Table 1 reports the results for a five-fold cross-validation; here, we can see that an accuracy of 98% is obtained when we use the 64 features of the first convolutional layer. These results indicate that LC features developed early in the sensor response are highly

Table 1. Five-Fold Cross-Validation of SVM Classification Using VGG16 Features

conv ^a layer	feature	accuracy	SD
second	128	100%	±0%
second	10	93%	±2%
first	64	98%	±1%
first	10	90%	±3%

^aThe abbreviation conv stands for convolutional layer.

informative and sufficient to discriminate among chemical environments.

From Table 1, we also see that it is possible to drastically reduce the feature set (this is done by selecting the features with the largest LSVM weights) while retaining a high accuracy level of 90–93%. The fact that we can obtain such high levels of accuracy with a reduced feature set can be attributed to the fact that the VGG16 network was pretrained using highly complex images, which suggests that many of the features extracted may be redundant or unnecessary (i.e., optical LC micrographs are simpler images than those used in generic databases such as ImageNet). In Table 2, we observe that the

Table 2. Five-Fold Cross-Validation of Select Time SVM Classification Using VGG16 Features

time (s)	second conv ^a	SD	first conv	SD
3	100%	±0%	96%	±2%
6	100%	±0%	95%	±2%
9	100%	±0%	94%	±2%
12	100%	±0%	96%	±1%
15	100%	±0%	94%	±2%
18	100%	±0%	95%	±1%
21	100%	±0%	95%	±2%
24	100%	±0%	96%	±2%
27	100%	±0%	96%	±2%

^aThe abbreviation conv stands for convolutional.

performance of the classifier is independent of the time at which the samples are collected. This reinforces our observation that differences in LC features develop early in the response, and they seem to persist. Our results achieve a reduction in the number of required features reported in previous work by 2 orders of magnitude. This reduction facilitates the physical interpretation of the LC features.

To validate the classification results of our ML framework, we compared our results against those of the classification achieved with principal component analysis (PCA). Here, we use PCA to project the 128-dimensional feature spaces of the second layer into two dimensions.²⁰ The results of the projection are visualized in Figure 8. The clustering and separation of the water and DMMP features indicate that perceptible differences in the CNN features of water and DMMP exist. These PCA results indicate that the features extracted from the CNN are indeed highly informative, but the

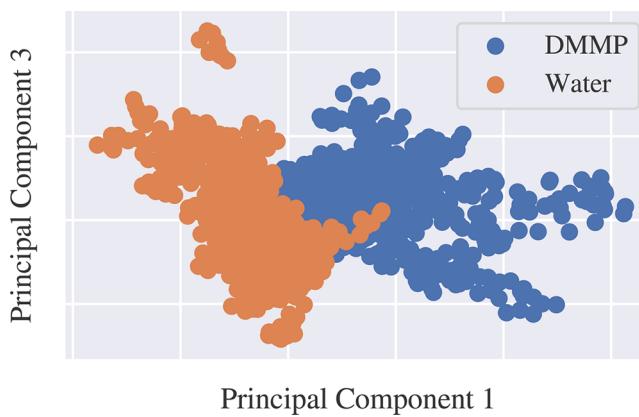


Figure 8. Classification using principal component analysis of VGG16 features.

existence of a significant overlapping region also highlights that an accurate classification between micrographs requires more than two features.

Achieving a high classification accuracy, while having a high importance from a sensor design standpoint, is not the only goal of our analysis. Specifically, we are interested in assigning the physical interpretation of the extracted features. To do so, we analyzed the features extracted from the first convolutional layer of VGG16 (visualized in Figure 6). These features are basic and highly informative, and they do not depend on previous layers of convolution. Consequently, the features of the first layer are generalizable and more suitable for physical analysis. We also recall that the features extracted from the first layer of VGG16 are the average outputs of 64 different filters. The LSVM hyperplane feature weights (shown in Table 3) help us identify which of these 64 filters are most dominant. Here, we see that Filters 8, 4, 52, and 38 are the most dominant ones.

Table 3. Optimal LSVM Weight Vector Obtained from the Training Set^a

filter number	filter weight percent	filter association
8	16.8%	water
4	16.5%	water
52	14.3%	DMMP
38	14.2%	DMMP
17	12.3%	water
18	9.3%	DMMP
6	8.0%	DMMP
37	5.0%	water
43	3.6%	water
10	0.1%	DMMP

^aUsing 10 features from the first CNN layer.

Maximally Activating Textures. To obtain some insight into the spatial patterns (textures) that the most dominant VGG16 filters (identified in Table 3) are capturing, we generated synthetic textures and identified the ones that maximized the average output for the different filters. This was done by feeding white noise images into VGG16 and modifying the image to maximize the output of each filter. We refer to these textures as the maximally activating textures. A visualization of this process is seen in Figure 9. Visualizations of the top five maximally activating textures for water are presented in Figure 10, and those for DMMP are presented in Figure 11. Here, we also show the activation fields on the input image associated with each filter. Two important aspects to consider when evaluating maximally activating textures are the color and the texture (spatial pattern). The hue (color) is of particular interest in the analysis of LCs, as different hues are a result of different orientations of the LCs within the film^{14,21} (assuming that the LC film thickness is relatively uniform). Moreover, the hue covers a spectrum of color, so it is preferred over RGB channels. In other words, a hue value captures the three values of RGB associated with a color. The maximally activating textures in Figures 10 and 11 reveal that DMMP and water have a distinct set of hues. From this observation, we conclude that the hue plays an important role in characterizing both water and DMMP responses.

Hue Analysis. In order to understand the importance of hue in the characterization of the DMMP and water responses, we developed a simple (but interpretable) feature set for each

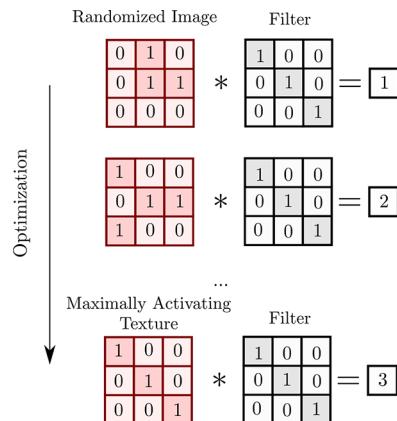


Figure 9. Graphic describing how to find maximally activating textures. To find the spatial pattern that is being maximized by a given filter, we feed different synthetic patterns and identify the one that maximizes the output.

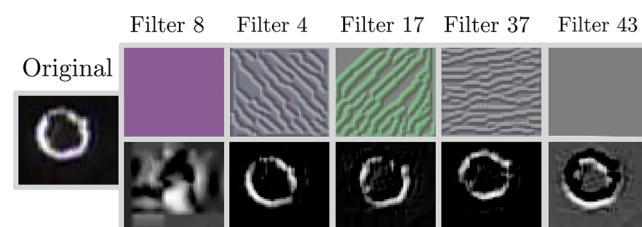


Figure 10. Maximally activating textures (top) and activations (bottom) for top water filters.

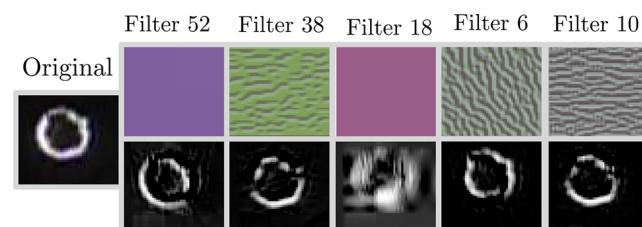


Figure 11. Maximally activating textures (top) and activations (bottom) for top DMMP filters.

image. Specifically, we analyzed the normalized distribution of the image hues. This distribution, which is split into 100 bins, captures the distribution of the hue within each sample image (the distribution of color). Each image is then represented as a 100-dimensional H vector in which each element h_i represents the probability (frequency) of finding a pixel in a given point of the hue spectrum.

Examples of hue distributions for water and DMMP are shown in Figure 12, and with their cumulative distribution functions (CDFs) are shown in Figure 13. From the hue distributions, we see that the intensity peak at a hue value of 65 (yellow to orange) is much stronger for DMMP than that for water. The CDF comparison reveals that DMMP exhibits no activity in the hue range 20–60 (blue to yellow), while water does. The CDFs also indicate that water micrographs have a more homogeneous coverage of the hue spectrum (reflected as a smoother CDF curve), while DMMP micrographs have a more heterogeneous coverage of the hue spectrum.

We used LSVM and hue distribution information to understand the efficacy of using hue in differentiating water and DMMP responses. In Table 4, we can see that an accuracy

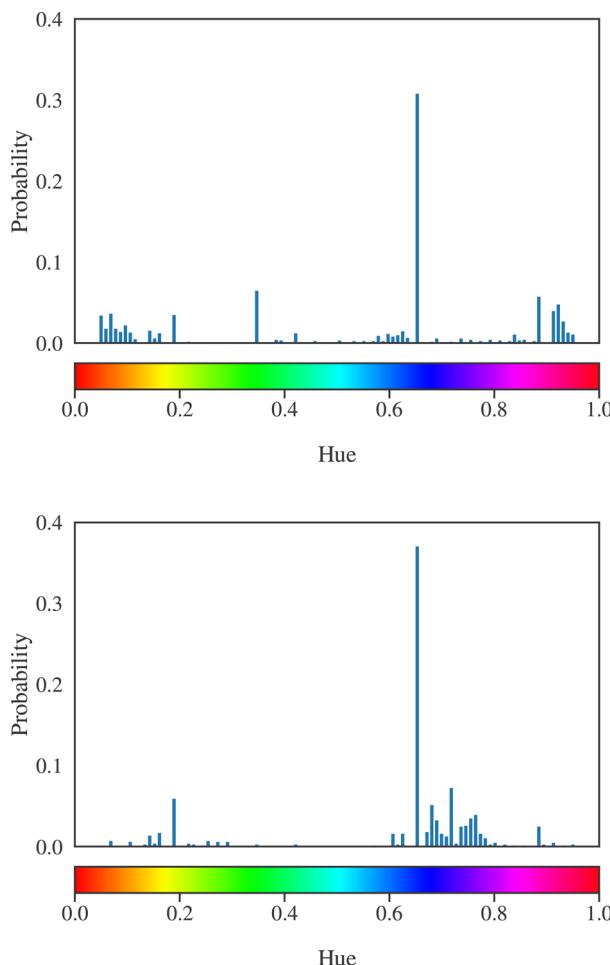


Figure 12. Hue distributions for representative water (top) and DMMP (bottom) micrographs.

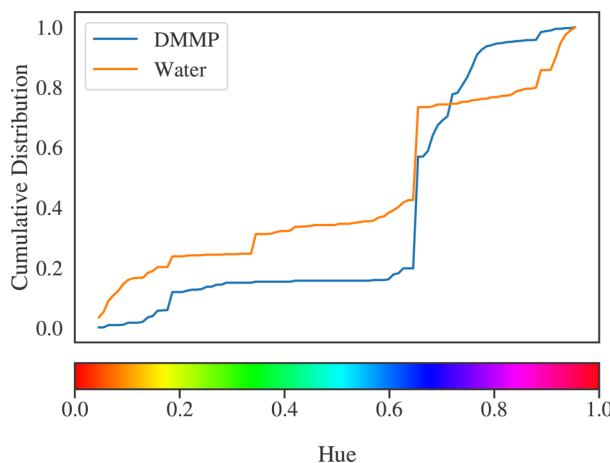


Figure 13. Comparison of the hue cumulative distributions for water and DMMP.

Table 4. Five-Fold Cross-Validation of LSVM Classification Using Hue Distribution

feature type	feature	accuracy	SD
hue distribution	100	88%	±8%

of nearly 88% can be achieved by using hue distributions of the images alone. These results reveal that hue (color) is an

informative feature for classification. Moreover, this result suggests that water and DMMP contain different hue distributions, which is most likely a result of differing LC orientations within the sensor film. Our results suggest that differences in color develop early in the response. These results make sense, because the optical properties of liquid crystals are known to be highly sensitive to stimuli. The lower classification accuracy obtained with hue distributions (compared with that of CNN features) is attributed to the fact that hue distributions do not capture spatial pattern (correlation) information (while CNN features do).

Grayscale Analysis. To understand the information content that can be attributed to color and to pure spatial patterns, we used VGG16 feature information extracted from grayscale images (ignoring color). From this analysis, we found that the classification accuracy was reduced by 6–12% (Table 5). This further supports the observation that color is an

Table 5. Five-Fold Cross-Validation of LSVM Classification Using Grayscale VGG16 Features

conv ^a layer	feature (grayscale)	accuracy	SD
second	128	94%	±2%
second	10	75%	±3%
first	64	87%	±3%
first	10	83%	±3%

^aThe abbreviation conv stands for convolutional.

important source of information but also that the spatial patterns found within the filters cannot be ignored. In order to analyze the grayscale patterns, we created a single texture that is a linear combination of the maximally activating textures. The linear combination was created by using the hyperplane weights obtained with LSVM. The linear combination is shown in eq 2, and the coefficients are taken from Table 3. The linear combinations of the grayscale patterns for DMMP and water are shown in Figure 14.

$$\begin{aligned}
 \text{DMMP texture} &= \text{filter 52} \left(\frac{0.143}{\text{total weight} = 0.459} \right) + \text{filter 38} \\
 &\quad \left(\frac{0.142}{0.459} \right) + \text{filter 18} \left(\frac{0.093}{0.459} \right) + \dots \\
 \text{water texture} &= \text{filter 8} \left(\frac{0.168}{\text{total weight} = 0.541} \right) + \text{filter 4} \\
 &\quad \left(\frac{0.165}{0.541} \right) + \text{filter 17} \left(\frac{0.123}{0.541} \right) + \dots
 \end{aligned} \tag{2}$$

The representative textures for both DMMP and water are used to summarize and understand differences in spatial patterns. The water texture possess a larger spatial correlation between the light and dark pixels, while the DMMP texture appears more randomized. We confirmed this observation quantitatively by analyzing the spatial autocorrelation of the textures. This is done by using Moran's I coefficient, which is a measure of global spatial autocorrelation, and is given by

$$\text{Moran's I} = \frac{N}{\sum_i \sum_j w_{ij}} \frac{\sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2} \tag{3}$$

Here, N represents the size of the neighborhood being analyzed. x_i represents the intensity of pixel i . \bar{x} represents the

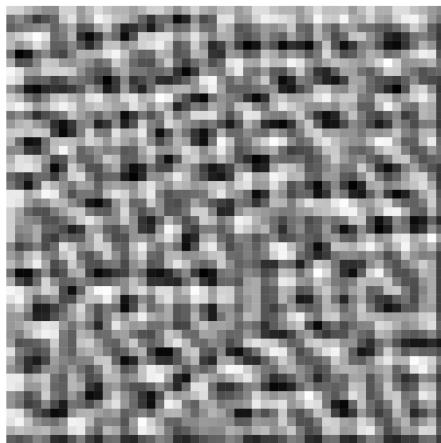
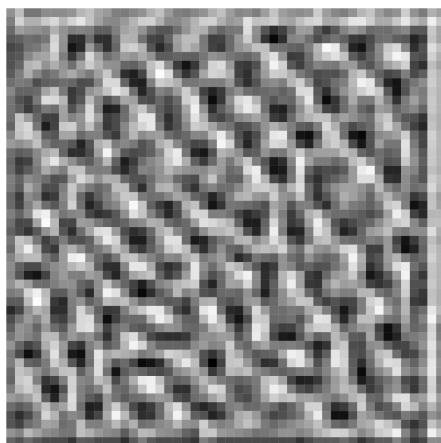


Figure 14. Textures for water (top) and DMMP (bottom). Textures are linear combinations of maximally activating filters.

average intensity in neighborhood N . w_{ij} represents the inverse distance weighting matrix in neighborhood N .

The Moran's I coefficients reveal that both DMMP and water patterns have a positive spatial autocorrelation with high confidence (Table 6), but the autocorrelation in water is of a

Table 6. Global Moran's I Coefficient Values

texture	Moran's I	P analysis
water	0.54	$P < 0.00001$
DMMP	0.40	$P < 0.00001$

longer range. This result may be further validated by calculating the local Moran's I coefficient values for every pixel in the image in a 3 pixel \times 3 pixel neighborhood. The resulting correlation fields, shown in Figure 15, indicate that the DMMP texture has a higher variance, and the areas of both positive and negative autocorrelation are clustered. For the water texture, on the other hand, we see a more uniform autocorrelation with a higher overall magnitude (confirming the observations obtained with the global Moran's I coefficient).

Our analysis indicates that VGG16 is capable of unraveling spatial patterns that result from exposure of the LC sensor to either DMMP or water. Moreover, we conclude that perceptible changes in spatial patterns are sufficient for the LSVM to discern between two chemical environments with a

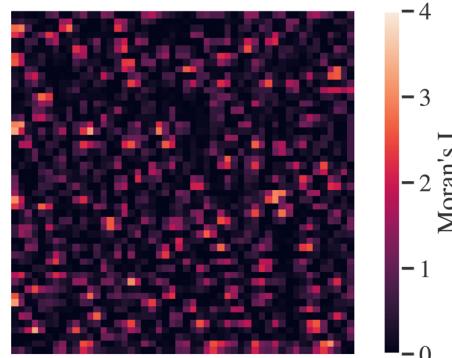
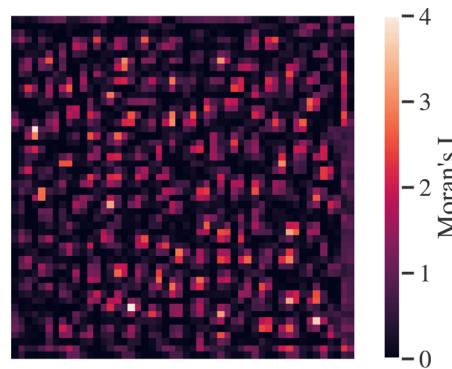


Figure 15. Local Moran's I analyses for water (top) and DMMP (bottom).

high accuracy. We hypothesize that the differences in the correlation lengths of the LC textures detected by VGG16 with DMMP and water reflect differences in the anchoring energy of the LC on the surface of the sensor. Specifically, a high anchoring energy will suppress LC orientational fluctuations and lead to a small correlation length. This result suggests that one key influence of water on the LC is to lower the anchoring energy at the metal salt-coated surface used in the LC sensor. The result also suggests that macroscopic orientational transitions may not be necessary in order to detect targeted chemical species using LCs, but the characterization of fluctuations in the orientation by using VGG16 may be a useful future strategy to explore in experiments.

Overall, analyses of both the grayscale spatial patterns and hues provide new insight into possible physical mechanisms that underlie the ability of VGG16 to differentiate the responses of the LC sensors to water and DMMP. Moreover, an additional important finding of our study is that perceptible changes in both color and spatial patterns can be detected with VGG16 within seconds of exposure of the LC film to the chemical environments (a thin bright ring is only perceptible by human vision early in the response).

Conclusions and Future Work. We have developed a machine learning framework to obtain high classification accuracies for optical micrographs of LC-based sensors. The features are outputs of the convolutional filters over a given image, which were extracted from the first and second layers of the VGG16 network. The selection of the VGG16 network was based upon its high performance on the Imagenet database and its use of small convolutional filters, which are believed to be able to extract small-scale texture and hue differences between LC sensors. The total number of features (filters) was

recursively reduced to ten, and these features were analyzed through the creation of a linear combination that represented water and DMMP textures. Analysis of these spatial patterns indicates that the correlation structure of the textures has perceptible differences. Specifically, the water texture is more uniform and has a higher correlation. The CNN analysis also reveals that color (hue) is an important feature that develops early in LC responses, and the use of the simplified grayscale images results in a large loss of information and a reduction in classification performance. This information led to the hypothesis that changes in the optical micrographs detected by VGG16 arise from spontaneous fluctuations in the LC orientations that reflect changes in LC anchoring energies. In order to explore this hypothesis further, physics-based molecular simulations that capture the dynamics of LC and spatiotemporal data analysis techniques are needed. Such techniques can be used to determine if the evolution of the spatial patterns follow different dynamic structures and/or present different time scales. If proven true, our findings may provide a fresh design principle for LC-based sensors in which the mesogen and surface are tuned to maximize differences in textures and hue fluctuations, as opposed to maximizing differences in average response times (which has been the main design principle behind LC-based sensors).

AUTHOR INFORMATION

Corresponding Author

Victor M. Zavala — Department of Chemical and Biological Engineering, University of Wisconsin-Madison, Madison, Wisconsin 53706, United States;  orcid.org/0000-0002-5744-7378; Email: victor.zavala@wisc.edu

Authors

Alexander D. Smith — Department of Chemical and Biological Engineering, University of Wisconsin-Madison, Madison, Wisconsin 53706, United States
Nicholas Abbott — Smith School of Chemical and Biomolecular Engineering, Cornell University, Ithaca, New York 14853, United States;  orcid.org/0000-0002-9653-0326

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jpcc.0c01942>

Notes

The authors declare the following competing financial interest(s): Nicholas Abbott has a financial interest in Platypus Technologies LLC, a company that sells analytical technologies based on liquid crystals.

All scripts and data needed to reproduce the results (when applying CNN techniques to analyze LC micrographs) are available at https://github.com/zavalab/ML/tree/master/LC_CNN_Color.

ACKNOWLEDGMENTS

We acknowledge funding from the U.S. National Science Foundation under BIGDATA grant IIS-1837812. Partial financial support for this work was also provided by the National Science Foundation through a grant provided to the UW-Madison Materials Research Science and Engineering Center (MRSEC, DMR-1720415).

REFERENCES

- Shah, R. R.; Abbott, N. L. Principles for measurement of chemical exposure based on recognition-driven anchoring transitions in liquid crystals. *Science* **2001**, *293*, 1296–1299.
- Mulder, D.; Schenning, A.; Bastiaansen, C. Chiral-nematic liquid crystals as one dimensional photonic materials in optical sensors. *J. Mater. Chem. C* **2014**, *2*, 6695–6705.
- Ireland, P.; Jones, T. Liquid crystal measurements of heat transfer and surface shear stress. *Meas. Sci. Technol.* **2000**, *11*, 969.
- Yang, K.-L.; Cadwell, K.; Abbott, N. L. Use of self-assembled monolayers, metal ions and smectic liquid crystals to detect organophosphonates. *Sens. Actuators, B* **2005**, *104*, 50–56.
- Zacharaki, E. I.; Wang, S.; Chawla, S.; Soo Yoo, D.; Wolf, R.; Melhem, E. R.; Davatzikos, C. Classification of brain tumor type and grade using MRI texture and shape in a machine learning scheme. *Magn. Reson. Med.* **2009**, *62*, 1609–1618.
- Kawahara, J.; BenTaieb, A.; Hamarneh, G. Deep Features to Classify Skin Lesions. In *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*; 2016; pp 1397–1400.
- Collins, P. C.; Koduri, S.; Welk, B.; Tiley, J.; Fraser, H. L. Neural networks relating alloy composition, microstructure, and tensile properties of α/β -processed TIMETAL 6–4. *Metall. Mater. Trans. A* **2013**, *44*, 1441–1453.
- Krizhevsky, A.; Sutskever, I.; Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Commun. ACM* **2012**, *55*, 1097–1105.
- Ling, J.; Hutchinson, M.; Antono, E.; DeCost, B.; Holm, E. A.; Meredig, B. Building data-driven models with microstructural images: Generalization and interpretability. *Materials Discovery* **2017**, *10*, 19–28.
- Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR 2015*, San Diego, CA, May 7, 2015; arXiv:1409.1556.
- Pan, S. J.; Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* **2010**, *22*, 1345–1359.
- Cao, Y.; Yu, H.; Abbott, N. L.; Zavala, V. M. Machine learning algorithms for liquid crystal-based sensors. *ACS sensors* **2018**, *3*, 2237–2245.
- Bedolla Pantoja, M. A.; Abbott, N. L. Surface-controlled orientational transitions in elastically strained films of liquid crystal that are triggered by vapors of toluene. *ACS Appl. Mater. Interfaces* **2016**, *8*, 13114–13122.
- Miller, D. S.; Carlton, R. J.; Mushenheim, P. C.; Abbott, N. L. Introduction to optical methods for characterizing liquid crystals at interfaces. *Langmuir* **2013**, *29*, 3154–3169.
- Hunter, J. T.; Abbott, N. L. Dynamics of the chemo-optical response of supported films of nematic liquid crystals. *Sens. Actuators, B* **2013**, *183*, 71–80.
- Zhu, J.; Rosset, S.; Tibshirani, R.; Hastie, T. J. 1-Norm Support Vector Machines. In *Advances in Neural Information Processing Systems*; 2004; pp 49–56.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*; 2009; pp 248–255.
- Géron, A. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*; O'Reilly Media, 2019.
- Abadi, M.; Agarwal, A. *Tensorflow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*; 2016; arXiv preprint, arXiv:1603.04467.
- Jolliffe, I. *International Encyclopedia of Statistical Science*; Springer, 2011; pp 1094–1096.
- Belyaev, V.; Solomatin, A.; Chausov, D. Phase retardation vs. pretilt angle in liquid crystal cells with homogeneous and inhomogeneous LC director configuration. *Opt. Express* **2013**, *21*, 4244–4249.