

Convolutional Network Analysis of Optical Micrographs for Liquid Crystal Sensors

Alexander Smith^a, Nicholas Abbott^b, and Victor M. Zavala^{a *}

^aDepartment of Chemical and Biological Engineering

University of Wisconsin-Madison, 1415 Engineering Dr, Madison, WI 53706, USA

^bSmith School of Chemical and Biomolecular Engineering

Cornell University, 113 Ho Plaza, Ithaca, NY 14853

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 that can shed some light on the underlying phenomena that governs and distinguishes LC responses. Previous work demonstrated that, by using features extracted from AlexNet, micrographs of different LC responses can be classified with an accuracy of 99%. Reaching such high levels of accuracy, however, required use of a large number of features (on the order of thousands), which was computationally intensive and which clouded the physical interpretability of the dominant features. To address these issues, here we report a study of the effectiveness of using features extracted from color images 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 on the same dataset used by Cao and co-workers, while reducing the number of features to less than a hundred. The number of features is further reduced to ten via recursive feature elimination with minimal loss in classification accuracy (5-10%). This feature reduction procedure reveals that differences in spatial color patterns are developed within seconds in the LC response. The results thus reveal that hue histograms provide an informative set of features that can be used to characterize LC micrographs of the sensor response. We also hypothesize that differences in 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. This latter proposal hints at fresh approaches for the design of LC-based sensors based on characterization of spontaneous fluctuations in orientation (as opposed to changes in time-average orientation).

Keywords: liquid crystals; chemical sensors; automated; machine learning; fast

*Corresponding Author: victor.zavala@wisc.edu




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

Liquid crystals (LCs) provide a versatile platform for sensing of air contaminants (chemical sensing) [1,2] and 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 diffusive transport of the analyte through the LC phase and displacement of the mesogen at the surface, triggering rich space-time optical responses (Figure 1). The response time of the 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 in [1].

A primary challenge for the development of LC sensors (and other sensing technologies) is their potential sensitivity to interfering chemical species. For instance, some LC sensors designed for detection of dimethyl methylphosphonate (DMMP) will exhibit similar optical responses when exposed to humid nitrogen [4]. LC sensor responses to DMMP and humid nitrogen are illustrated in Figure 2. Although the selectivity of LC sensors can be 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 designs of LC sensors.

Machine learning techniques have been recently 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 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/describe

an 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, rich features that capture multi-scale spatial patterns can be extracted from CNNs that have been pre-trained 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 [6], the pre-trained CNN Alexnet [8] is used to classify skin lesions. In the work of [9], textures extracted from the CNN VGG16 [10] are used to predict material properties. The principle behind the exploitation of feature information from pre-trained CNNs is known as *transfer learning* [11].

Cao and coworkers 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, they 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 to the chemical environment (DMMP and water) are sufficient to classify the environment. 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 build from the previous results of Cao and coworkers by analyzing features extracted from VGG16, which is a CNN that embeds a smaller set of convolutional filters than Alexnet. Moreover, in the current study, we use RGB 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 *perfect* classification accuracy for the same dataset used by Cao and coworkers, while reducing the number of features to approximately one hundred. We demonstrate that the number of features can be further reduced to ten 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 LC tilt direction play a key role in sensor selectivity and responsiveness. Our analysis also reveals that *hue histograms* provide an effective set of features to characterize LC responses.




Figure 2: Optical responses of liquid crystals under gaseous N_2 -water (30% relative humidity) and N_2 -DMMP (10 PPM) environments. Microwells are 3mm in diameter. LCs were deposited into microwells. LC responses were recorded at room temperature.

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). In other words, our framework is focused on binary classification. We use the same dataset reported by Cao and coworkers [12] but focus on patterns developed within the first 30 seconds of the LC response. Details regarding the experimental system and data preparation methods can be found in our prior report [12]. In summary, the dataset 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 (micrograph) from each video every 3.3 seconds. 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. Examples of snapshot sequences collected during the microwell responses are presented in Figure 2.

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 dataset 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 3. LSVM is a 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 non-informative 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.




Figure 3: Illustration of a linear support vector machine. Lines that are orthogonal to the separation hyperplane represent the margin.

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 not only finding a separating hyperplane

that best classifies the images but is also required to do so 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:

$$\underset{w_0, w}{\text{minimize}} \sum_{i=1}^n \left[1 - y_i \left(w_0 + \sum_{j=1}^q w_j x_{ij} \right) \right] + \lambda \|w\|_1 \quad (0.1)$$

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 , and $\lambda \in \mathbb{R}_+$ is a hyperparameter for the penalization of the l_1 norm [13]. The solution of problem (0.1) 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) dataset. This process is repeated five times, each time with a new training and validation set. This allows for a robust testing of the effectiveness of the classification model on the entire dataset.

Feature Extraction

In order to train the LSVM classifier, we first need to identify features that best explain each image. Cao and coworkers previously used Alexnet to conduct feature extraction from LC micrographs. Alexnet is a CNN that has been pre-trained using the Imagenet database [14]. This database is a collection of millions of images that contains over 1,000 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 re-training 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 5,000 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 physical LC behavior.

In this work, we consider a different pre-trained 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 physical behavior of LCs. A primary consideration is the length scale that characterizes the LC reponse. 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, we assume that the patterns created by the LCs need to be captured with 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 the VGG16 CNN, which has been pre-trained by the visual geometry

group at Oxford [10]. This CNN has also been trained using the `Imagenet` database but utilizes the smallest possible convolutional filter size (3x3), which should be best for capturing small-scale structural patterns in the images. Moreover, VGG16 is a much shallower CNN than AlexNet. A simplified representation of the VGG16 architecture is shown in Figure 4.

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, each 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 decision block in Figure 4. More details on this can be found in [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 the convolutional filter, is present within a neighborhood of an image (and repeating the process by spanning all neighborhoods of the image). An example of applying a convolution filter to an image (expressed as a matrix) is illustrated in Figure 6. Convolutional filters provide a quantifiable approach for identifying patterns within a given neighborhood (different filters identify different types of patterns). The larger the value of the filter output, the more similar the given neighborhood is to the pattern that filter is attempting to find. Optimal filter matrices that best classify a set of images can be found by training the CNN directly on the dataset. This training procedure, however, involves a highly computationally expensive procedure. Filters extracted from training over a given set, however, can also be reused to seek for similar patterns in a different dataset. In other words, pre-trained filters are used on a different image set with the sole purpose of obtaining feature information. While the filters are not optimal for the new dataset, this procedure is often effective at detecting general patterns in images and the obtained featured information can be used in an external classifier such as LSVM.

In the example provided in Figure 6, 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 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 equally 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




Figure 4: Schematic of VGG16 architecture.

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). 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 5. Feature extraction and network modification were performed using Keras [15] and Tensorflow [16]. 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.

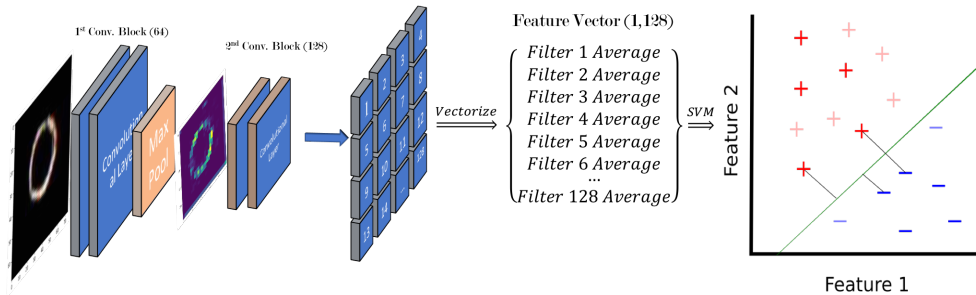


Figure 5: Schematic of feature extraction and classification framework.

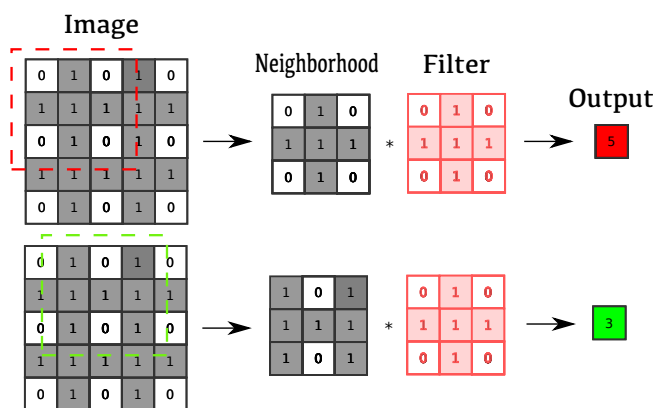


Figure 6: Illustration of the application of a convolution filter to a neighborhood of two different images.

Results and Discussion

Classification and Feature Reduction

We were able to classify water and DMMP micrographs with *100% accuracy* using VGG16 features and LSVM. Notably, these results were obtained for micrographs collected within 30 seconds 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 response* are highly 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 an 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 pre-trained 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 use in generic databases such as ImageNet). In Table 2 we observe that the 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 *two orders of magnitude*. This reduction facilitates the interpretation of the LC features.

Layer	Features	Accuracy	Std.Dev.
2 nd Conv.	128	100 %	± 0 %
2 nd Conv.	10	93 %	± 2 %
1 st Conv.	64	98 %	± 1 %
1 st Conv.	10	90 %	± 3 %

Table 1: Five-fold cross validation of SVM classification using VGG16 features.

Time	2 nd Conv.	Std.Dev.	1 st Conv	Std.Dev.
3 seconds	100 %	± 0 %	96 %	± 2 %
6 Seconds	100 %	± 0 %	95 %	± 2 %
9 Seconds	100 %	± 0 %	94 %	± 2 %
12 Seconds	100 %	± 0 %	96 %	± 1 %
15 Seconds	100 %	± 0 %	94 %	± 2 %
18 Seconds	100 %	± 0 %	95 %	± 1 %
21 Seconds	100 %	± 0 %	95 %	± 2 %
24 Seconds	100 %	± 0 %	96 %	± 2 %
27 Seconds	100 %	± 0 %	96 %	± 2 %

Table 2: Five-fold cross validation of select time SVM classification using VGG16 features.

To validate our classification results, we conducted principal component analysis (PCA) to project the 128 dimensional feature space of the second layer into two dimensions [17]. The results of the projection are visualized in Figure 7. The clustering and separation of the water and DMMP features indicates that there are perceptible differences in the CNN features of water and DMMP. These PCA results indicate that the CNN features are indeed highly informative but the existence of a significant overlapping region also highlights that an accurate classification between micrographs requires more than two features.

The highly classification accuracy achieved, while having high importance from a sensor design stand-point, is not the only goal of our analysis. Specifically, we are interested in assigning physical interpretation of the extracted features. To do so, we analyzed the features extracted from the first convolutional layer of the VGG16 network (visualized in Figure 5). These features are basic, highly informative, and do not depend on previous layers of convolution. Consequently, the features of the first layer are generalizable and more suitable for physical analysis. We recall also that the features extracted from the first layer of VGG16 are the average outputs of 64 different filters. The LSVM hyperplane feature weight (shown in Table 3) help us identify which of these 64 filters are most dominant.




Figure 7: Principal component analysis of VGG16 features.

Filter Number	Filter Weight Percent	Filter Association
Filter 8	16.8 %	Water
Filter 4	16.5%	Water
Filter 52	14.3%	DMMP
Filter 38	14.2%	DMMP
Filter 17	12.3%	Water
Filter 18	9.3%	DMMP
Filter 6	8.0%	DMMP
Filter 37	5.0%	Water
Filter 43	3.6%	Water
Filter 10	0.1%	DMMP

Table 3: Optimal LSM weight vector obtained from training set (using ten features from first CNN layer).

Maximally Activating Textures

To obtain some insight into the spatial patterns (textures) that the most dominant VGG16 filters identified 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 8. Visualizations of the top five maximally activating textures for water are presented in Figure 9 and for DMMP are presented in Figure 10. 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 color and the texture (spatial pattern). The hue is of particular interest in the analysis of LCs as different hues are a result of different orientations of the liquid crystals within the film [18, 19] (assuming that the film thickness is relatively uniform). Moreover, hue covers a spectrum of color, so it is preferred over RGB channels. The maximally activating textures in Figures 10 and 9, reveal that DMMP and water

have a distinct set of hues. From this observation, we hypothesize that *hue plays an important role in characterizing both water and DMMP responses*.

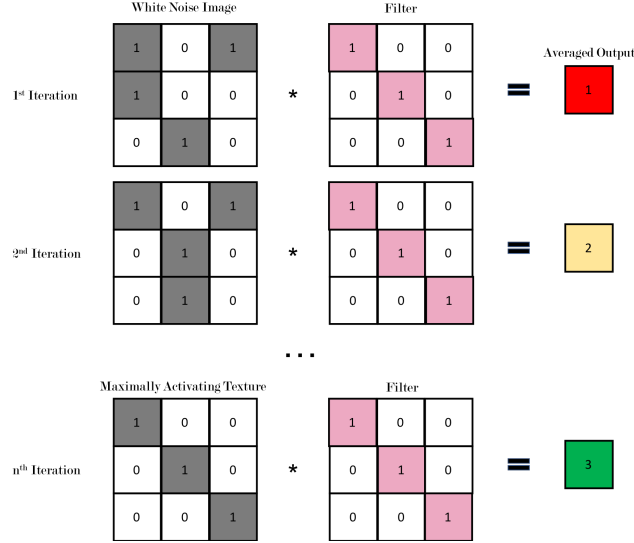


Figure 8: Finding maximally activating textures. To find the spatial pattern that is being maximized by a given filter, we feed different synthetic patterns and identify the one that maximizes the output.




Figure 9: Maximally activating textures (top) and activations (bottom) for top water filters.

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 image. Specifically, we analyzed the normalized histogram of the image hues. This histogram, which is split into 100 bins, captures the distribution of hue within each sample image. Each image is then represented as a 100-dimensional H vector in which each element h_i represents the probability of finding a pixel in a given point of the hue spectrum. An example hue histogram for water and DMMP are shown in Figure 11, along with their cumulative distribution functions (CDFs) in 12. From the hue histograms we see that the intensity peak at a hue value of 65 (yellow to orange) is much stronger for DMMP than for water. The CDF comparison reveals that DMMP exhibits no activity in the hue range of 20-60 (blue




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

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 heterogenous coverage of the hue spectrum.

We used LSVM and hue histogram information to understand the efficacy of hue in differentiating a water and DMMP responses. In Table 4 we can see that an accuracy of nearly 88% can be achieved by using hue histograms alone. These results reveal that *hue 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. Moreover, our results suggest that differences in color develop early in the response. The lower classification accuracy obtained with hue histograms (compared with CNN features) are attributed to the fact that hue histograms do not capture spatial pattern information (while CNN features do).




Figure 11: Hue histograms for representative water (left) and DMMP (right) micrographs.




Figure 12: Comparison of the hue cumulative distributions for water and DMMP.

Feature Type	Features	Accuracy	Std.Dev.
Hue Histogram	100	88 %	± 8 %

Table 4: Five-fold cross validation of LSVM classification using hue histogram.

Layer	Features (Grayscale)	Accuracy	Std.Dev.
2 nd Conv.	128	94 %	± 2 %
2 nd Conv.	10	75 %	± 3 %
1 st Conv.	64	87 %	± 3 %
1 st Conv.	10	83 %	± 3 %

Table 5: Five-fold cross validation of LSVM classification using grayscale VGG16 features.

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%. This further supports the observation that color is an 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 (0.2) and the coefficients are taken from Table 3. The linear combinations of the grayscale patterns for DMMP and water are shown in Figure 13.

$$\text{DMMP Texture} = \text{Filter } 52 \left(\frac{0.143}{\text{Total Weight} = 0.459} \right) + \text{Filter } 38 \left(\frac{0.142}{0.459} \right) + \text{Filter } 18 \left(\frac{0.093}{.459} \right) + \dots$$

$$\text{Water Texture} = \text{Filter } 8 \left(\frac{0.168}{\text{Total Weight} = 0.541} \right) + \text{Filter } 4 \left(\frac{0.165}{0.541} \right) + \text{Filter } 17 \left(\frac{0.123}{.541} \right) + \dots \quad (0.2)$$




Figure 13: Textures for water (left) and DMMP (right). Textures are linear combinations of maximally activating filters.

The representative textures for both DMMP and water are used to summarize and understand differences in spatial patterns. The water texture possesses 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. Moran's I coefficient, which is a measure of global spatial autocorrelation, and given by:

$$\text{Moran's I} = \frac{N}{\sum_i \sum_j w_{ij}} \frac{\sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2} \quad (0.3)$$

Here, N represents the size of the neighborhood being analyzed, x_i represents the intensity of pixel i , \bar{x} represents the average intensity in neighborhood N , and w_{ij} represents the inverse distance weighting matrix in neighborhood N .

Texture	Moran's I	P Analysis
Water	0.54	$P < 0.00001$
DMMP	0.40	$P < 0.00001$

Table 6: Global Moran's I coefficient values.

The Moran's I coefficients reveal that both DMMP and water patterns have positive spatial autocorrelation with high confidence (Table 6) but that autocorrelation in water is of 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 3x3 pixel neighborhood. The resulting correlation fields, shown in Figure 14, indicate that the DMMP texture has higher variance and that areas of both positive and negative autocorrelation are clustered. For the water texture, on the other hand, we see a more uniform autocorrelation and with higher overall magnitude (confirming the observations obtained with the global Moran's I coefficient).




Figure 14: Local Moran's I analysis for water (left) and DMMP (right).

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 high accuracy. We hypothesize that the differences in correlation length 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 that characterization of fluctuations in orientation by using VGG16 may be a useful future strategy to explore in experiments.

Overall, analysis of both the grayscale spatial patterns and hues provide new insight into possible physical mechanisms that underlie the ability of VGG16 to differentiate the response 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 conducted a detailed feature extraction analysis that provides high classification accuracy 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 autocorrelation structure of the textures has perceptible differences. Specifically, the water texture is more uniform and has higher autocorrelation. The CNN analysis also reveals that color (hue) is an important feature that develops early in LC responses. This information lead 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 further this hypothesis, physics based molecular simulations that capture the dynamics of LC and spatio-temporal 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 timescales. If proven true, our findings may provide 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).

Acknowledgments

We acknowledge funding from the U.S. National Science Foundation under BIGDATA grant IIS-1837812.

References

- [1] R. R. Shah and N. L. Abbott, "Principles for measurement of chemical exposure based on recognition-driven anchoring transitions in liquid crystals," *Science*, vol. 293, no. 5533, pp. 1296–1299, 2001.
- [2] D. Mulder, A. Schenning, and C. Bastiaansen, "Chiral-nematic liquid crystals as one dimensional photonic materials in optical sensors," *Journal of Materials Chemistry C*, vol. 2, no. 33, pp. 6695–6705, 2014.

- [3] P. Ireland and T. Jones, "Liquid crystal measurements of heat transfer and surface shear stress," *Measurement Science and Technology*, vol. 11, no. 7, p. 969, 2000.
- [4] K.-L. Yang, K. Cadwell, and N. L. Abbott, "Use of self-assembled monolayers, metal ions and smectic liquid crystals to detect organophosphonates," *Sensors and Actuators B: Chemical*, vol. 104, no. 1, pp. 50–56, 2005.
- [5] E. I. Zacharaki, S. Wang, S. Chawla, D. Soo Yoo, R. Wolf, E. R. Melhem, and C. Davatzikos, "Classification of brain tumor type and grade using mri texture and shape in a machine learning scheme," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 62, no. 6, pp. 1609–1618, 2009.
- [6] J. Kawahara, A. BenTaieb, and G. Hamarneh, "Deep features to classify skin lesions," in *Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on*, pp. 1397–1400, IEEE, 2016.
- [7] P. C. Collins, S. Koduri, B. Welk, J. Tile, and H. L. Fraser, "Neural networks relating alloy composition, microstructure, and tensile properties of α/β -processed timetal 6-4," *Metallurgical and Materials Transactions A*, vol. 44, no. 3, pp. 1441–1453, 2013.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- [9] J. Ling, M. Hutchinson, E. Antono, B. DeCost, E. A. Holm, and B. Meredig, "Building data-driven models with microstructural images: Generalization and interpretability," *Materials Discovery*, vol. 10, pp. 19–28, 2017.
- [10] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [11] S. J. Pan, Q. Yang, *et al.*, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [12] Y. Cao, H. Yu, N. L. Abbott, and V. M. Zavala, "Machine learning algorithms for liquid crystal-based sensors," *ACS sensors*, vol. 3, no. 11, pp. 2237–2245, 2018.
- [13] J. Zhu, S. Rosset, R. Tibshirani, and T. J. Hastie, "1-norm support vector machines," in *Advances in neural information processing systems*, pp. 49–56, 2004.
- [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*, pp. 248–255, Ieee, 2009.
- [15] F. Chollet *et al.*, "Keras," 2015.
- [16] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, "Tensorflow: a system for large-scale machine learning.," in *OSDI*, vol. 16, pp. 265–283, 2016.

- [17] I. Jolliffe, "Principal component analysis," in *International encyclopedia of statistical science*, pp. 1094–1096, Springer, 2011.
- [18] D. S. Miller, R. J. Carlton, P. C. Mushenheim, and N. L. Abbott, "Introduction to optical methods for characterizing liquid crystals at interfaces," *Langmuir*, vol. 29, no. 10, pp. 3154–3169, 2013.
- [19] V. Belyaev, A. Solomatin, and D. Chausov, "Phase retardation vs. pretilt angle in liquid crystal cells with homogeneous and inhomogeneous lc director configuration," *Optics Express*, vol. 21, no. 4, pp. 4244–4249, 2013.