Anderson Lebbad
Department of Mechanical Engineering,
Villanova University,
Villanova, PA 19085
e-mail: alebba01@villanova.edu

Garrett M. Clayton
Associate Professor
Mem. ASME
Center for Nonlinear Dynamics & Control,
Department of Mechanical Engineering,
Villanova University,
Villanova, PA 19085
e-mail: garrett.clayton@villanova.edu

C. Nataraj
Moritz Endowed Professor
Life Fellow ASME
Villanova Center for Analytics of
Dynamic Systems,
Department of Mechanical Engineering,
Villanova University,
Villanova, PA 19085
e-mail: c.nataraj@villanova.edu

Convolutional Networks for Classification of Mortars

The task of classifying unexploded mortars is critical in both humanitarian and military explosive ordnance disposal (EOD) operations. Classification needs to be completed quickly and accurately and is the first step toward disarming the ordnance because it provides information about the fuzing mechanism, or the stage in the arming cycle that the ordnance is currently in. To assist EOD technicians with mortar identification, this article presents an automated image-based algorithm and the database of images used in its development. The algorithm utilizes convolutional networks with variations to training to improve performance for ordnance found in varying states of disassembly. The classifier developed was found to be 98.5% accurate for these lab condition photos; future work will focus on more cluttered environments. [DOI: 10.1115/1.4053886]

Keywords: machine learning, neural networks

1 Introduction

This article presents a classification tool to assist an explosive ordnance disposal (EOD) technician in identifying mortars so that they can be disarmed/disposed (after the ordnance has been detected/unearthed). The presented method uses a convolutional network trained on images of mortars that could be used as a support system for field operators and could prove especially useful if their previous experience or training is limited, as is the case with EOD technicians in developing countries.

There are several technologies that have been developed to characterize ordnance by explosive content, for example, by using vapor pressure, x-rays, thermal neutron analysis, gamma spectroscopy [1], or Raman spectroscopy [2]. There have also been methods developed to characterize ordnance by the housing, such as magnetic analysis methods [3,4]. Most of these methods require very expensive hardware, and several are limited to chemical munitions, which exclude high explosive (HE) unexploded ordnance (UXO). Several types of these technologies do not work for artillery rounds [5]. These methods often have practical limitations in the field, such as trace gas analysis—where the soil and wind conditions may affect the sensing [5].

The current ubiquity of cameras make an image-based solution desirable, and developments in machine learning and computer vision [6–8] make such a solution potentially feasible. Importantly, this type of approach, which minimizes the physical interaction necessary with the UXO, reduces the chance of accidental detonation. Recent work has been done using non-image-based methods to discriminate between types of UXO [1,9], as well as detection of UXO using computer vision methods [10–12], but the automated discrimination of UXO using computer vision is still largely an unexplored topic.

Convolutional networks (convnets) are the most promising method for this task [13–18]. Convnets are an extension of neural networks and use convolutional filters. Generally, neural networks require *a priori* feature extraction. In contrast, the raw data are directly fed into convnets [19]. This removes the heuristic layer of feature selection, using layers of filters for end-to-end optimization on the algorithm. This is much more efficient than using a traditional feed-forward neural network to produce similar results, while the theoretical best performance is only slightly better [20].

This article presents an image acquisition method used to build a dataset of UXO images and a convnet classifier to discriminate between the different models of mortars within the database. The images used within the dataset were captured as a part of this study by the authors. They came from the ordnance collections of Golden West Humanitarian Foundation, and Cambodia Mine Action Centre (through collaboration with Golden West Humanitarian Foundation). To the authors' knowledge, UXO discrimination using convnets has not previously been attempted.

2 Image Acquisition

This method utilized a multiple camera setup to collect images from multiple angles at once. This caused a high number of photos to be produced of each piece of UXO. Photos were taken of many different types of ordnance, but this initial analysis is on discrimination of a subset of the images of mortars collected.

2.1 System Design. The configuration of eight cameras tethered via USB cables to the same computer was ideal for acquiring photos from several different perspectives simultaneously. These perspectives offered views of opposing sides of the ordnance and differing conditions of lighting. The cameras were setup in four stereo pairs around the ordnance. The pairs allowed the option for stereo analysis for dimensional analysis, which was not used in this work. The configuration can be seen in Fig. 1.

2.2 Acquisition Procedure. A controlled background was laid out for the ordnance to be photographed on. The cameras were staged so the entire view was within the bounds of the control background, and all the photos only contained a reference object and the subject ordnance piece. Focal length, aperture settings, and focus were all manually set to ensure all ordnance that would be

Manuscript received April 18, 2021; final manuscript received February 1, 2022; published online March 22, 2022. Assoc. Editor: Reza Tafreshi.

Fig. 1 This figure shows the configuration of the eight-camera image acquisition setup. Each of the cameras is zoomed and focused on the piece of ordnance and checkerboard (used as a reference for scale). The diagram in the bottom right of the figure shows the relative positions of the cameras, as some are difficult to distinguish from the foliage in the background.

photographed during a session would be maximally visible, while staying entirely within the frame along with the reference object.

Each piece of UXO was captured in at least ten different positions by the eight-camera setup. These variations in position were focused on exposing portions of the ordnance that were previously unseen or causing the light to hit the ordnance differently. Up to ten pieces of each model of ordnance were captured, if that many pieces were available. So, for a given model of ordnance, this method would capture up to 800 images.

2.3 Example Photos. Once acquired, the images are cropped and scaled to only contain the ordnance and to produce the standard 224×224 pixel size used in training the classifier. Images taken from the same camera of ten ordnance pose variations can be seen on the left side of Fig. 2. On the right side of Fig. 2, one of the shots is enlarged to show details of the ordnance.

3 Dataset

The ordnance images used in this work are from three ordnance libraries located in Cambodia—Cambodian Mine Action Centre's (CMAC's) [21] Demining Unit 4 library in Siem Reap, CMAC's Phnom Penh library, and Golden West Humanitarian Foundation's [22] Kampong Chhnang site. This collection included 11,968 total images of 150 distinct pieces, which included 28 unique models of ordnance. These models ranged in type from antipersonnel mines, to grenades, to mortars, including both illumination and high explosive charges.

3.1 Classification Set. A subset of the database developed was used for the classification discussed in this article. The subset was selected by choosing ordnance of the same type (mortars) and

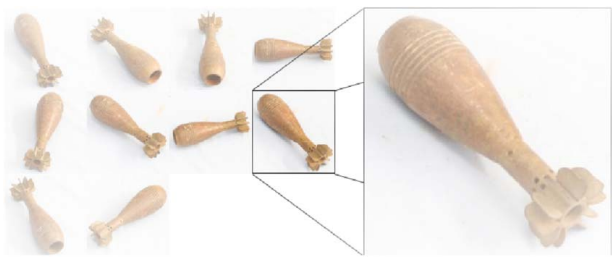

Fig. 2 Images of variations of pose of ordnance with respect to the camera. All images are of the same piece of ordnance taken from the same camera. This figure shows how ten images of each piece of ordnance were collected from each camera, to get 80 images per piece of ordnance.

Table 1 Models of ordnance by class

Class	Country	Model	Size	Images
M1	China	Type 22	60 mm	800
M2	China	Type 30	82 mm	808
M3	USA	M374 series, HE	81 mm	800
M4	USSR	F-843	120 mm	552
M5	USSR	S832SM, Illum.	82 mm	152
M6	Vietnam	Prop. mdl. UNK	82 mm	400
M7	Yugoslavia	M74, HE	82 mm	160

specific models of the type that had enough pieces available to be photographed to have a significant testing set. The subset contains seven different models of ordnance (in varying conditions of assembly and disrepair), totaling 46 individual pieces of UXO and 3672 photographs. The distribution of the images among the seven classes are detailed in Table 1, along with information on their countries of origin and their sizes (specifically diameter). The training and validation sets for each model of ordnance were taken from images of a single piece (with images from multiple angles), as some ordnance libraries only contain only one piece of each model, with many (sometimes thousands) different models available.

3.2 Example Images. Figure 3 shows photos of all the models of mortars used in the classifier. Some were slightly damaged (M4 has bent tail fins), and some were more assembled than others (M3 is the only piece with a fuze attached and M6 is an illumination round that has a nose cone but is not actually fuzed). In the field, ordnance can be found in various states of damage, disrepair, and disassembly. Our focus is hence to derive an algorithm to be robust to these kinds of variations.

4 Classification Method

A convnet classifier was trained using this subset of the database. There are four operations that convnets are primarily composed of: convolution, nonlinearity, pooling, and classification (fully connected).

An image can be numerically represented by channels of two-dimensional arrays. A gray-scale image has one channel (intensity), while a color image has three (red, green, and blue).

Convolution is completed by convolving a mask (also known as a kernel) with the image. In regions of the image where the kernel matches more, the result of the convolution is a higher value. On the lower layers of convnet, these kernels tend to optimize toward detecting edges and corners, while the higher layers find the combinations of edges and corners that identify more recognizable shapes. The output of a convolution layer is often referred to as a feature map, with the kernel also being known a as a feature detector. All the convolutions used in this convnet structure were 3×3 arrays.

A nonlinear activation function after every layer (except for pooling, as it is a nonlinear process itself) is necessary to maintain effectiveness of the layers. This is because taking a linear function of a linear function would result in a linear function, effectively negating the effect of any hidden layers. The rectified linear unit is the most often used nonlinearity function for convnets, which is also commonly referred to as ReLU. If the input value was positive, the ReLU simply passes that value forward. If the input value was negative, then the ReLU returns a zero. There are other nonlinearities that could be used, such as sigmoid or hyperbolic tangents, but the ReLU performs typically as well as these other options in the context of convnets and offers several advantages—it creates data sparsity through the zero returns, it does not suffer from vanishing gradient, and it is very simple to compute.

The pooling operation downsizes the feature maps. Pooling does this by analyzing a region of the image and then shifting the operation by a value known as the stride. This effectively splits the

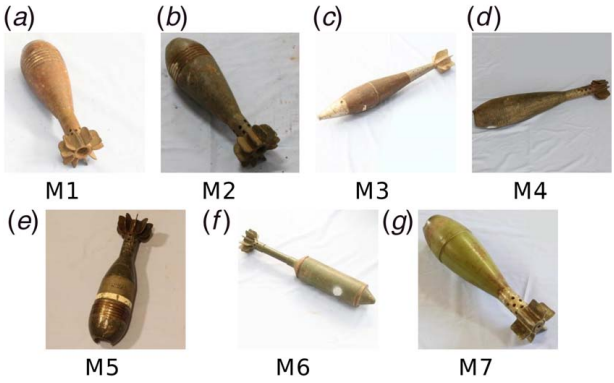

Fig. 3 This figure shows a training image for each of the classes of ordnance detailed in Table 1: (*a*) M1, (*b*) M2, (*c*) M3, (*d*) M4, (*e*) M5, (*f*) M6, and (*g*) M7

image into regions of interest and consolidates the information within these regions. This structure uses max pooling, which returns the highest value within the region of interest. The structure evaluates the max pool on 2×2 regions, with a stride of 2 between points of evaluation—so every maxpool layer divides the results of the previous layer into a grid of 2×2 values and selects the maximum value within each region. The result of the maxpool is half the size of the input in both dimensions, so a 224×224 input would have a 112×112 output.

Classification layers are also known as fully connected layers. The resultant data from the previous layer are rearranged into a vector and run through a traditional multilayer perceptron. The outputs of this final layer are run through a soft-max, which normalizes the results so that the sum of the prediction values equal one.

4.1 Structure. The 16-layer structure from Ref. [23] was the basis of the network structure for this classifier. Between each of the layers of this structure was a ReLU activation function. Each of the convolutional layers used small, 3×3 kernels, with a stride of 1. To preserve spatial resolution through the convolutional layers, the input image to each layer was padded by 1 pixel. The pooling layers were 2×2 max pooling, with a stride of 2.

The structure takes an RGB image of 224×224 pixels as an input and then runs the image through layers of convolution filters (each three dimensional to handle the three separate color channels), followed by pooling until the dimensionality is reduced enough to vectorize the output and run through a fully connected network. The structure is as follows: $2\times$ layers of 64 filters, pool, $2\times$ layers of 128 filters, pool, $3\times$ layers of 256 filters, pool, $3\times$ layers of 512 filters, pool, 3x layers of 512 filters, pool, fully connected layer of 4096 nodes, fully connected layer of 4096 nodes, fully connected layer of 7 nodes (for each of the different output classes), and a soft-max layer to normalize the outputs.

4.2 Training. Photos of a single piece of ordnance were used to train each class, with 60 photos used for training each class (for a total set size of 420 images) and 20 photos used for validation of each class (totaling 140 validation images). Training convergence of the classifier was defined as when the loss metric (mean squared error) of the validation set changed less than 0.0001 for three consecutive epochs. This resulted in perfect performance of the validation set, as detailed in Table 2.

5 Results

5.1 Metrics. The remaining 3112 images were used as the testing set for the classifier. The classifier took 71 epochs for the training to converge (by the definition in Sec. 4.2). The following terms are defined to quantify the effectiveness of the classifier: true positive (TP)—actual and predicted are both positive; true negative (TN)—both actual and predicted are negative; false positive (FP): actual is negative, predicted is positive; false negative (FN)–actual is positive, predicted is negative.

Table 2 Confusion matrix, validation set

		Predicted						
		M1	M2	M3	M4	M5	M6	M7
Actual	M1	20	0	0	0	0	0	0
	M2	0	20	0	0	0	0	0
	M3	0	0	20	0	0	0	0
	M4	0	0	0	20	0	0	0
	M5	0	0	0	0	20	0	0
	M6	0	0	0	0	0	20	0
	M7	0	0	0	0	0	0	20

By using these terms, many metrics can be defined to evaluate the effectiveness of the algorithm. We define and use two relevant ones here: *recall* (TP/(TP+FN)) and *precision* (TP/(TP+FP)).

Recall measures how likely an image of a particular class will be correctly identified, and precision measures how likely images identified as belonging to a certain class actually are of objects belonging to that class. The F_1 score combines these two into a more general metric, as follows.

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{1}$$

The F_1 score of a classification is in the range of 0–1, where a value of 1 would have both perfect precision and recall.

The statistics for the classifier are presented in Table 3. These are drawn from the confusion matrix, Table 4. Most classes (M1, M2, M4, M5, M6, and M7) had over 85% recall. The high number of misclassifications of M3 ordnance leads to very low recall for M3 and lower precision for M1, M4, M5, and M7 (although the precision statistics for M5 and M7 are heavily affected by the low number of images of each of those classes available in the testing set).

Table 3 Classification statistics

	Recall	Precision	F_1
M1	0.946	0.722	0.819
M2	0.984	0.923	0.952
M3	0.386	1.000	0.557
M4	0.877	0.793	0.833
M5	1.000	0.686	0.814
M6	1.000	0.970	0.985
M7	0.988	0.500	0.664

Table 4 Confusion matrix, testing set

		Predicted						
		M1	M2	M3	M4	M5	M6	M7
Actual	M1	681	0	0	0	0	1	38
	M2	0	716	0	12	0	0	0
	M3	262	2	278	95	33	9	41
	M4	0	58	0	414	0	0	0
	M5	0	0	0	0	72	0	0
	M6	0	0	0	0	0	320	0
	M7	0	0	0	1	0	0	79

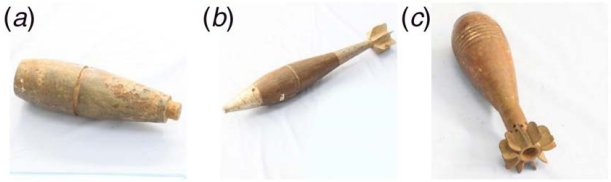

Fig. 4 Misclassification can be related to different states of assembly: (*a*) A mortar of class M3. This image does not have a tail or fuze attached, unlike the piece used in the training set, which can be seen in (*b*). (*a*) was incorrectly classified as class M1 (which can be seen in (c)).

5.2 Misclassifications. Images of class M3 had the most misclassifications, with a high percentage of the images incorrectly classified as belonging to class M1. As discussed, some of the ordnance pieces were partially disassembled or damaged. A large portion of the images of class M3 were of pieces that neither had a tail nor a fuze attached (as seen in Fig. 4(*a*)), unlike the images used in training the classifier (Fig. 4(*b*)). The majority of the images of these disassembled pieces of ordnance were found to be misclassified as M1, which is shown in Fig. 4. The tail and fuze of M3 appear to have identifying features that were likely used by the classifier to discern it among the other classes of ordnance. This highlighted a limitation of the classifier, a lack of robustness to disrepair and disassembly of the ordnance.

5.3 Variations. Due to the high amount of misclassification associated with varying states of disassembly, training sets of different states of disassembly were tested to see if performance could be increased. To classify the effectiveness of a classifier over all classes within a testing set, *accuracy* is defined as the sum of the correctly classified images for all the classes, over the total number of images in the full testing set. For the classifier trained with the original set (O), the accuracy was determined to be 82.3%. This led to creating alternative training sets to improve performance.

To improve classification on disassembled pieces of class M3, the V training set was created, where M3 training images had neither the tail nor the fuze attached. As presented in Table 5, the classifier from the V set improved minimally in performance compared to the original set. The classification results of the testing set of M3 images effectively swapped, where now the images of the disassembled ordnance were correctly classified, but the fully assembled pieces were generally confused with class M6. This classifier was found to have an accuracy of 89.8% across the testing set.

The classifier was then trained on a new set (B) that contained both the training images of O and V. This led to a significant improvement in performance. Whether this caused the classifier to find more general trends between the states of disassembly or if it effectively just trained for two different patterns to return the same result is unclear. Either way, the B set classifier had 98.0% accuracy on the testing set.

To see if the increased number of training images was necessary, or if it was simply the variety of the set, another training set (H) was created. Half of the training images from O and half from V were included in this set. The classifier from the H set performed very similarly to the classifier of B, where twice as many images of class M3 were used in training. The accuracy of the H classifier was found to be 98.5%.

Table 5 F_1 scores of variations

	O	V	B	H
M1	0.819	0.992	0.990	0.994
M2	0.952	0.978	0.971	0.977
M3	0.557	0.759	0.997	0.998
M4	0.833	0.966	0.955	0.962
M5	0.814	0.986	1.000	0.980
M6	0.985	0.714	0.982	0.998
M7	0.664	0.894	0.957	0.955

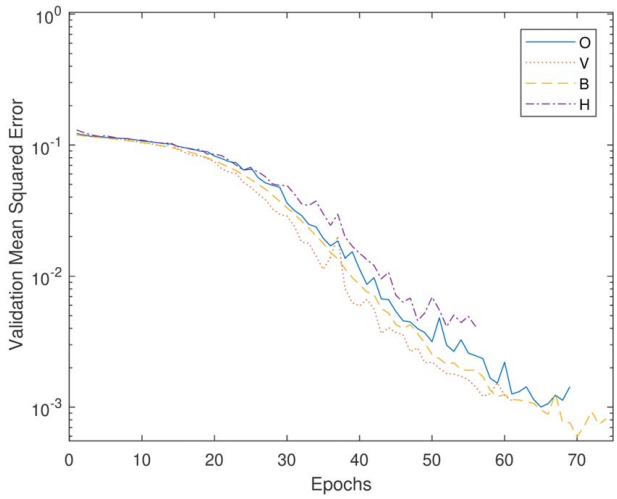

Fig. 5 Validation loss per epoch for each of the four classifiers over the course of training

Table 6 Variation *H* confusion matrix, testing set

		Predicted						
		M1	M2	M3	M4	M5	M6	M7
Actual	M1	712	1	2	0	2	1	2
	M2	0	708	0	19	1	0	0
	M3	0	0	720	0	0	0	0
	M4	0	13	0	459	0	0	0
	M5	0	0	0	0	72	0	0
	M6	0	0	0	0	0	320	0
	M7	0	0	1	4	0	0	75

As shown in Fig. 5, the validation loss at the end of training is actually higher for the H classifier. This was a result of the increased variety in the training and validation sets. The O and V classifiers were trained and validated on more uniform sets, so the validation results were more weakly correlated with the testing results than B or H. This is presented in Table 5, which shows the F_1 scores of each class for each of the variations of the classifier.

The testing results of the H classifier is presented in Table 6. The classification of type M3 mortars has drastically improved (to perfect classification of the testing set). Now the largest discrepancy is between class M2 and M4, which by visual inspection (looking at Fig. 3) appear very similar. Upon review of the misclassifications between M2 and M4, only 3 of the 32 were found to not have the correct class as the second highest prediction result. The biggest difference between the two models of ordnance is scale, which cannot be inferred from the images alone. Future work will be done to include other kinds of information (such as measurements of length and other kinds of expert data) to be able to discern between similar classes such as these.

6 Conclusion

The focus of this article was to develop an automated tool to assist field operators (especially those with limited experience) to determine the model of a mortar to reduce risk through the EOD process. The approach taken was to create an image-based classifier using convolutional networks. To facilitate this, a method of data

acquisition to rapidly image multiple perspectives of ordnance using an eight-camera configuration was developed. This method was then used to build a dataset of images of 150 different pieces of ordnance, totaling 11,968 photos. These pieces were accessed and photographed in the facilities of Cambodian Mine Action Centre and Golden West Humanitarian Foundation. From this set, a subset was selected of mortars that had enough photos taken to provide a significant testing set. The classifier was trained and validated using only photos of a single piece of ordnance of each model (as many other pieces in the dataset are limited to a single piece), and the classifier was found to have a testing set accuracy of 82.3%.

The classifier was noticed to be underperforming on testing samples that were more or less assembled than the piece used in the training set (as in when many pieces of class M3 did not have the fuze or tail attached in the testing set, unlike the training set). A classifier was trained with half the images in the training set being either assembled or disassembled, and the performance of the classifier significantly increased—having a testing set accuracy of 98.5%. Further research needs to be done to explore using the classifier with varying backgrounds, and work is currently being done toward this effort. Future work in building larger databases of ordnance could potentially be improved by staging the ordnance in different states of assembly (imaging with different combinations of removable parts attached).

Acknowledgment

This work would not have been possible without the access to inert ordnance and technical support provided by Golden West Humanitarian Foundation and Cambodian Mine Action Center. We are also grateful for the financial support provided by Golden West Humanitarian Foundation and the U.S. Department of State, Office of Weapons Removal and Abatement.

Conflict of Interest

There are no conflicts of interest.

References

[1] Bruschini, C., 2001, "Commercial Systems for the Direct Detection of Explosives for Explosive Ordnance Disposal Tasks," Subsurface Sens. Technol. Appl., **2**(3), pp. 299–336.
[2] Axelsson, M., Friman, O., Johansson, I., Nordberg, M., and Ostmark, H., 2013, "Detection and Classification of Explosive Substances in Multispectral Image Sequences Using Linear Subspace Matching," 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Vancouver, BC, Canada, May 26–31, pp. 3492–3496.
[3] Chen, X., O Neill, K., Barrowes, B. E., Grzegorczyk, T. M., and Kong, J. A., 2004, "Application of a Spheroidal-Mode Approach and a Differential Evolution Algorithm for Inversion of Magneto-Quasistatic Data in UXO Discrimination," Inverse Prob., **20**(6), p. S27.
[4] Prouty, M., George, D. C., and Snyder, D. D., 2011, "Metalmapper: A Multi-Sensor TEM System for UXO Detection and Classification," Geometrics Inc., San Jose, CA, Technical Report, Accession No. ADA578954.
[5] McFee, J. E., and Das, Y., 1980, "The Detection of Buried Explosive Objects," Can. J. Remote Sens., **6**(2), pp. 104–121.
[6] Szeliski, R., 2010, *Comuputer Vision: Algorithms and Applications*, Springer, Berlin/Heidelberg, Germany.
[7] Zeiler, M. D., and Fergus, R., 2014, "Visualizing and Understanding Convolutional Networks," European Conference on Computer Vision, Zurich, Switzerland, Sept. 6–12, Springer, pp. 818–833.
[8] Liu, Q., and Liu, C., 2017, "A Novel Locally Linear KNN Method With Applications to Visual Recognition," IEEE Trans. Neural Netw. Learn. Syst., **28**(9), pp. 2010–2021.
[9] Billings, S. D., 2004, "Discrimination and Classification of Buried Unexploded Ordnance Using Magnetometry," IEEE Trans. Geosci. Remote Sens., **42**(6), pp. 1241–1251.
[10] Amer, S., Shirkhodaie, A., and Rababaah, H., 2008, "UXO Detection, Characterization, and Remediation Using Intelligent Robotic Systems," Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XIII, Vol. 6953 of PROCSPIE, Orlando, FL, Mar. 17–20.
[11] Katartzis, A., Vanhamel, I., Chan, J., and Sahli, H., 2004, "Remote Sensing Minefield Area Reduction: Model-Based Approaches for the Extraction of Minefield Indicators," Theory and Applications of Knowledge-Driven Image Information Mining With Focus on Earth Observation, Madrid, Spain, Mar. 17–18.
[12] Shirkhodaie, A., and Rababaah, H., 2007, "Visual Detection, Recognition, and Classification of Surface-Buried UXO Based on Soft-Computing Decision Fusion," Detection and Remediation Technologies for Mines and Minelike Targets XII, Orlando, FL, Apr. 11–12, p. 655328.
[13] LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D., 1989, "Backpropagation Applied to Handwritten Zip Code Recognition," Neural Comput., **1**(4), pp. 541–551.
[14] Venieris, S. I., and Bouganis, C., 2019, "fpgaconvnet: Mapping Regular and Irregular Convolutional Neural Networks on FPGAS," IEEE Trans. Neural Netw. Learn. Syst., **30**(2), pp. 326–342.
[15] Kang, Y., Chen, S., Wang, X., and Cao, Y., 2019, "Deep Convolutional Identifier for Dynamic Modeling and Adaptive Control of Unmanned Helicopter," IEEE Trans. Neural Netw. Learn. Syst., **30**(2), pp. 524–538.
[16] Kim, J., Nguyen, A., and Lee, S., 2019, "Deep CNN-Based Blind Image Quality Predictor," IEEE Trans. Neural Netw. Learn. Syst., **30**(1), pp. 11–24.
[17] Sakhavi, S., Guan, C., and Yan, S., 2018, "Learning Temporal Information for Brain-Computer Interface Using Convolutional Neural Networks," IEEE Trans. Neural Netw. Learn. Syst., **29**(11), pp. 5619–5629.
[18] Liu, N., Han, J., Liu, T., and Li, X., 2018, "Learning to Predict Eye Fixations Via Multiresolution Convolutional Neural Networks," IEEE Trans. Neural Netw. Learn. Syst., **29**(2), pp. 392–404.
[19] LeCun, Y., Kavukcuoglu, K., and Farabet, C., 2010, "Convolutional Networks and Applications in Vision," Proceedings of 2010 IEEE International Symposium on Circuits and Systems (ISCAS), Paris, France, May 30–June 2.
[20] Krizhevsky, A., Sutskever, I., and Hinton, G. E., 2012, *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds., Curran Associates, Inc., Red Hook, NY, pp. 1097–1105.
[21] 2021, "La mine d argent de potosi: la technique au centre d une lutte de pouvoir entre les incas et les espagnols (xvie-xviie siècles)."
[22] 2020, "Golden West Humanitarian Foundation."
[23] Simonyan, K., and Zisserman, A., 2014, "Very Deep Convolutional Networks for Large-Scale Image Recognition," CoRR.