



2022 Annual International Conference on Brain-Inspired Cognitive Architectures for Artificial Intelligence: The 13th Annual Meeting of the BICA Society

Robust Classification of Contraband Substances using Longwave Hyperspectral Imaging and Full Precision and Neuromorphic Convolutional Neural Networks

Kyung Chae Park^a, Jeremy Forest^a, Sudepto Chakraborty^a, James T. Daly^b, Suhas Chelian^a, Srini Vasani^{a,*}

^aQuantum Ventura, 1 S. Market St., Suite 1715, San Jose, CA 95125, USA

^bBodkin Design and Engineering, 77 Oak St., Suite 201, Newton Upper Falls, MA 02464, USA

Abstract

Several agencies such as the US Department of Homeland Security (DHS) seek to improve the detection of illegal threats and materials passing through Ports of Entry (POE). A combined hardware/software solution that is portable, non-ionizing, handheld, low cost, and fast would represent a significant contribution towards that goal as existing systems do not fulfil many or all of these requirements. To design such a system, Quantum Ventura partnered with Bodkin Design and Engineering to combine long-wave infrared (LWIR) hyperspectral imaging (HSI) with convolutional neural networks (CNNs), implemented on full precision GPUs and neuromorphic computing modules. Our capability study showed that our system can accurately detect and classify contraband in a variety of situations, including varied backgrounds, temperatures, and purities. With a small size, weight, power and cost (SWaP-C) envelope, neuromorphic computing implementations of CNNs showed promising results, though not as well as full precision results.

© 2022 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the 2022 Annual International Conference on Brain-Inspired Cognitive Architectures for Artificial Intelligence: The 13th Annual Meeting of the BICA Society.

Keywords: Contraband detection, long-wave infrared (LWIR) hyperspectral imaging (HSI), convolutional neural networks (CNNs), neuromorphic computing

1. Introduction

Improving the detection of illegal threats and materials passing through Customs and Border Protection (CBP)

* Corresponding author. Tel.: +1 424-227-1417

E-mail address: srini@quantumventura.com

Ports of Entry (POE) is sought by several agencies such as the US Department of Homeland Security (DHS) [1]. The large volume of conveyances, material and cargo entering the United States daily, however, pose a challenge. Concerns for both materials and persons at these legal border crossings limit the type of inspection tools that can be used safely. A combined hardware/software solution that is portable, non-ionizing, handheld, weighs less than 4.5 kg (10 lbs.), costs less than \$50,000 USD that allows cargo to be scanned and assessed with results provided in seconds while scanning either from a moving platform or containers moving at speeds of up to 4.5 m/s (10 mph) would represent a significant contribution to high throughput non-ionizing screening capability for identifying contraband.

As part of DHS Phase 1 Small Business Innovation Research (SBIR) contract, Quantum Ventura, Inc. partnered with Bodkin Design & Engineering (BD&E)--hyperspectral imaging specialists with over three decades of proven experience--to design a high throughput non-ionizing screening capability based on long-wave infrared (LWIR) hyperspectral imaging (HSI), and convolutional neural networks (CNNs) implemented on full precision GPUs and reduced precision neuromorphic computing modules. Our capability study demonstrated that LWIR HSI and CNNs can detect and classify contraband in a variety of conditions such as different backgrounds, temperatures and purities with good accuracy. Furthermore, neuromorphic computing implementations of CNNs offer promising results, though less than full precision results, with a low size, weight, power and cost (SWaP-C) envelope.

2. Methods and Materials

2.1 Hyperspectral Imaging

We began with the adoption of long-wavelength infrared (LWIR) hyperspectral imaging (HSI) for chemical analysis and identification because the LWIR absorption spectra are unique for different molecules. LWIR occurs between 8 and 14 μm . Figure 1 shows the spectra for a range of contraband and potential background materials. On the left, contraband chemicals include fentanyl and surrogates for organophosphate nerve agents methyl salicylate and SF96; nerve agent Vx is also shown. On the right, background materials include lignin, a major component of cardboard; plastic which occurs in shipping containers, suitcases, etc.; and polyester which occurs in some clothing are shown.

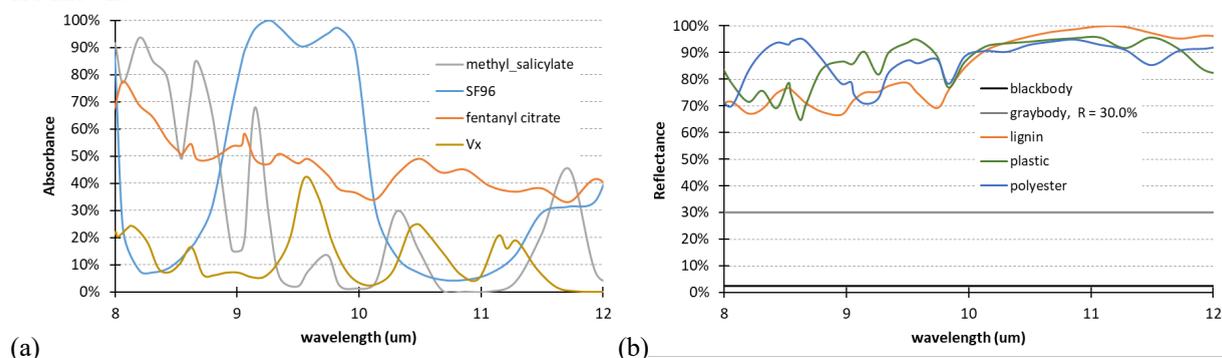


Figure 1. Absorption spectra (a) contraband chemicals of methyl salicylate, SF96, fentanyl citrate and Vx and (b) background chemicals of lignin, plastic, polyester, and a spectrally flat ($R=30\%$) graybody.

2.2 Simulation of Hyperspectral Imaging (HSI) Spectra

Due to scarcity of contraband/controlled substance hyperspectral imaging (HSI), we created custom code to generate synthetic HSI data from emittance spectra of these chemicals. This code enables us to create synthetic HSI of arbitrary chemicals with user specified backgrounds, temperatures, and purities (coverage amounts). This capability will also allow us to produce synthetic data to address new and emerging threats. We generated synthetic HSI data of target chemicals: fentanyl citrate, methyl salicylate, and SF96, and background chemicals: lignin, polycarbonate (bisphenol A aka BPA) and polyester. Lignin is a major component of cardboard; polycarbonate occurs in shipping containers, suitcases, etc.; and polyester which occurs in some clothing.

For chemicals and backgrounds, we used spectra from the several datasets [2, 3, 4, 5, 6, 7]. For temperature, the radiance of a background will be given by the product of the material's emissivity as a function of wavelength times the Planck function for that temperature:

$$L(\lambda, \tau) = \varepsilon(\lambda) \cdot B(\lambda, \tau) \quad (1)$$

where $\varepsilon(\lambda)$ is the background emissivity vs wavelength, λ . $B(\lambda, \tau)$ is the Planck radiance at temperature τ . For purities, we consider a foreground residue of the contraband chemical which either completely covers (100%) or partially covers (<100%) its background. The foreground residue does not present as a film of uniform thickness, but rather as clusters of various sizes and coverage. To account for this, we implement a coverage or filling factor, f , that scales the relative contributions of the clean and contaminated surfaces. Thus the effective radiance of the contaminated surface becomes:

$$L_{cont}^e = fL_{cont} + (1 - f)L_{clean} \quad (2)$$

The resulting radiance change is then:

$$L_{cont}^e - L_{clean} = f[L_{cont} - L_{clean}] = f[(1 - R_0)\Delta B - (R_0 - R)(B - L_d)] \quad (3)$$

where L 's are radiance values, f is the coverage ratio, R_0 is the reflectance of the clean surface, and B is the Planck radiance. Note that this equation assumes each cluster has the same reflectance.

Furthermore, we synthesized 27 artificial HSI datacubes. A datacube is a 3D structure with two spatial dimensions (rows and columns) and one spectral dimension as shown in Figure 2a. Each datacube has the three contraband chemicals side by side, on three specific backgrounds, at three specific temperatures, with three specific coverage amounts ($3 \times 3 \times 3 = 27$). (In the field, the three contraband chemicals are unlikely to be next to each other but this was done to create a more compact dataset.) Spectra were synthesized by assuming a normal distribution at each wavelength with a nominal standard deviation of $1e-5$ representing camera noise. There were 61 bands with a width of 50 nm between 7.8 and 10.8 μm . This is based on the BDE LWIR-62 camera which has 62 bands [8]. However, we removed the top band because it often has high noise characteristics. The synthetic datacubes consist of small rectangular pixel patches of fentanyl citrate (cyan, left), methyl salicylate (green, middle), and SF96 (yellow, right) on a lignin, BPA, or polyester backgrounds (dark blue) as in Figure 2b. Temperatures used were 278 K (5°C, 41°F), 293 K (20°C, 68°F), and 308 K (35°C, 95°F). This is to show that our system can work on a cold day along the US-Canada border, in moderate temperatures, or a hot day along the US-Mexico border. Coverage amounts used were 100%, 50% and 25%.

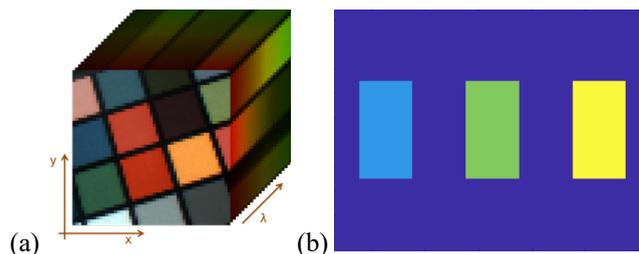


Figure 2. (a) Datacubes were composed of two spatial dimensions and one spectral dimension. Apparent rotation of image is due to packing arrangement for 40 pixel long spectral streaks at spatial format of 90x75 on a SXGA detector array. (b) Synthetic images were composed of small rectangular 4x2 patches of fentanyl citrate (cyan, left), methyl salicylate (green, middle), and SF96 (yellow, right) in various 10x12 backgrounds (dark blue).

All datacubes were used to train and test the same CNN but 20% of pixels were used for training and 80% were used for testing. The CNN was trained like so because we are interested in implementing a multi-class contraband substance classification system to satisfy DHS's specifications.

2.3 Full precision GPU Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs) are the latest incarnation of neural networks. Compared to existing machine learning algorithms, CNNs are known to provide significantly more robustness against variability in targets and the environments in which they occur. CNNs can be implemented on GPUs using 16-bit precision, which we refer to as full precision or on neuromorphic processors which use less precision (e.g., 4 or 8 bits). Neuromorphic implementation is described in the next section.

In one of our recent projects involving hyperspectral imaging (HSI) to detect far-off unmanned aerial vehicles flying at a distance of 10 km (6.2 mi.) [9], we successfully demonstrated the detection and classification capabilities of a supervised 3D CNN model with a Spectral-Spatial Residual Network (SSRN) architecture which takes raw HSI data cubes as input. Here 3D refers to two spatial dimensions and one spectral dimension, not three spatial dimensions (see Figure 2a).

SSRNs consecutively extract discriminative spectral and spatial features for classification [10]. SSRNs outperform other CNNs due to 1) the residual (ResNet) blocks that connect every other convolutional layer through identity mapping facilitating the back propagation of gradients, and 2) the batch normalization of the output of every layer to regularize the learning process. The spectral residual blocks (Figure 3, top left) comprise n convolutional kernels used in successive filter banks h^{p+1} and h^{p+2} for p^{th} and $(p + 1)^{\text{th}}$ layers. The two convolutional layers build a residual function $F(X^p; \theta)$ directly mapping input 3D feature cubes X^p using a skip connection. The spatial residual block (Figure 3, bottom left) extracts spatial features from input 3D feature cubes X^q using n 3D convolutional kernels in the successive filter banks H^{q+1} and H^{q+2} for the two successive layers. The full SSRN architecture includes sequentially connected spectral and spatial feature learning sections, an average pooling layer, and a fully connected (FC) layer (Figure 4).

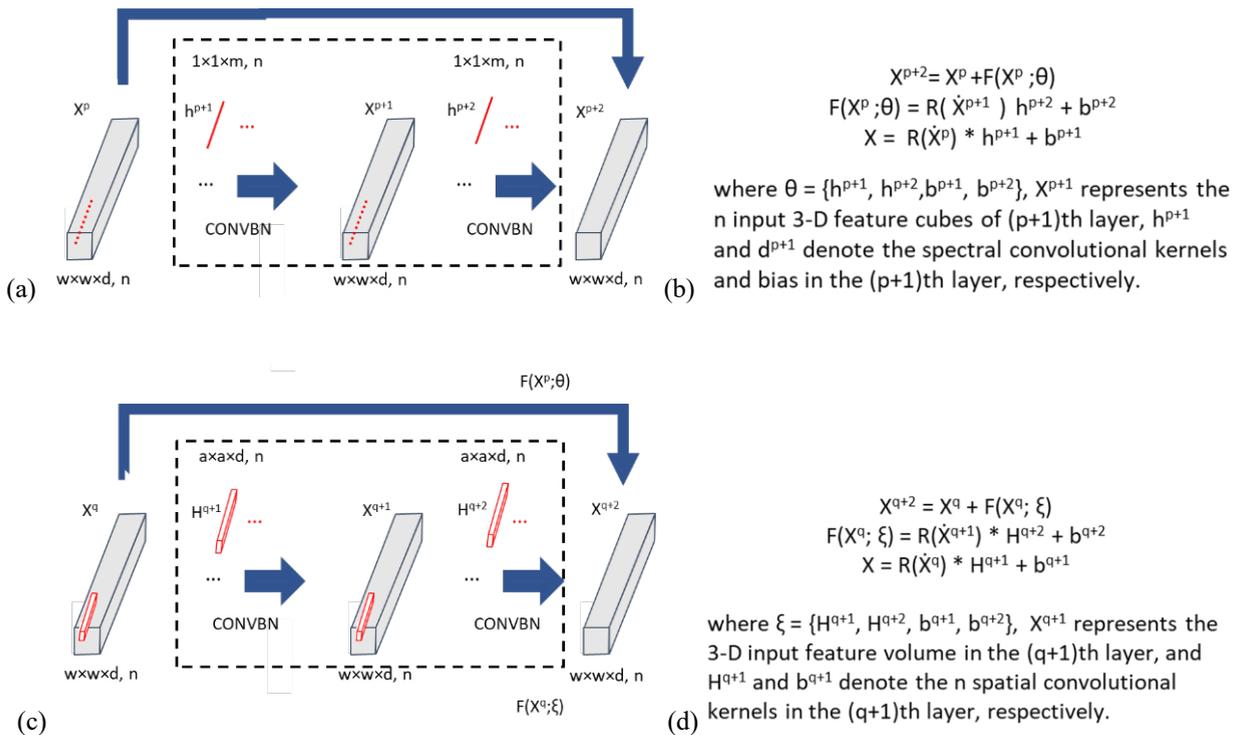


Figure 3. (a) Schematic representation and (b) computational representation of the spectral residual CNN block. (c) Schematic representation and (d) computational representation of the spatial residual CNN.

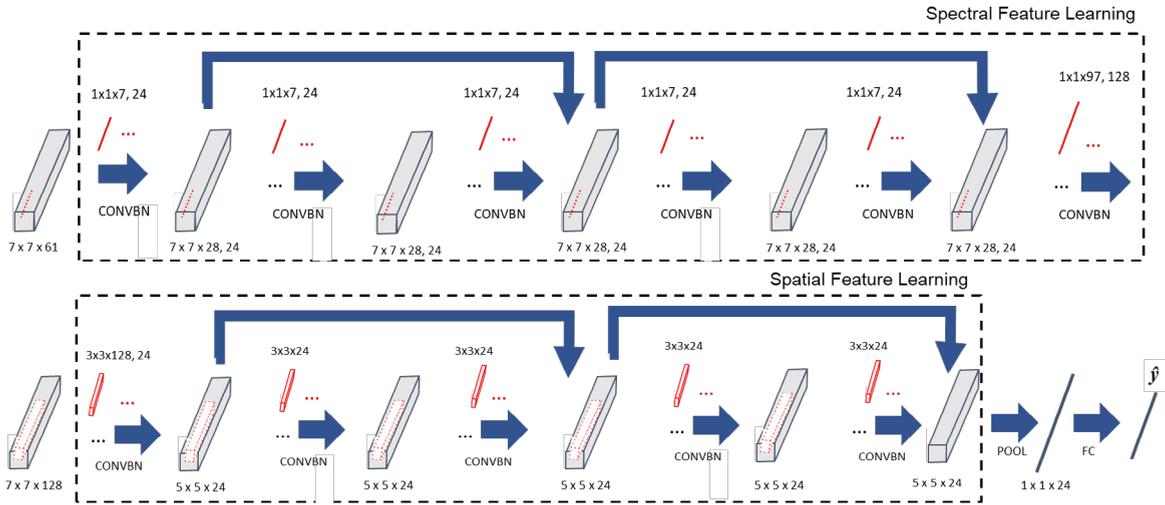


Figure 4. Spectral-spatial residual network (SSRN) with consecutive spectral and spatial residual blocks.

In Table 1, we show the details of the CNN (layer, kernel sizes, etc.). We also experimented with different input patch sizes and other parameters. Training parameters were: 200 epochs using the RMSprop optimizers with a learning rate of 0.0003 and a batch size of 16.

Table 1. Layer and kernel sizes of the CNN. BN is batch normalization. L is the number of output classes; four in our case—one for each contraband class and one for background or other.

| Layer | Type | Kernel size | Stride | Input shape | Activation | Output shape |
|-------|--|--------------|---------|-----------------|------------|--------------|
| 1 | Conv3D(BN) | (1,1,7)x24 | (1,1,2) | (7,7,61) | ReLU | (7,7,28,24) |
| 2 | Conv3D(BN) | (1,1,7)x24 | (1,1,1) | (7,7,28,24) | ReLU | (7,7,28,24) |
| 3 | Conv3D(BN) | (1,1,7)x24 | (1,1,1) | (7,7,28,24) | ReLU | (7,7,28,24) |
| 4 | Conv3D(Residual from layer 2 as an input+BN) | (1,1,7)x24 | (1,1,1) | (7,7,28,24) x 2 | ReLU | (7,7,28,24) |
| 5 | Conv3D(BN) | (1,1,7)x24 | (1,1,1) | (7,7,28,24) | ReLU | (7,7,28,24) |
| 6 | Conv3D(Residual from layer 4 as an input+BN) | (1,1,97)x128 | (1,1,1) | (7,7,28,24) x 2 | ReLU | (7,7,128) |
| 7 | Conv3D(BN) | (3,3,128)x24 | (1,1,1) | (7,7,128) | ReLU | (5,5,24) |
| 8 | Conv3D(BN) | (3,3,24) | (1,1,1) | (5,5,24) | ReLU | (5,5,24) |
| 9 | Conv3D(Residual from layer 7 as an input+BN) | (3,3,24) | (1,1,1) | (5,5,24) x 2 | ReLU | (5,5,24) |
| 10 | Conv3D(BN) | (3,3,24) | (1,1,1) | (5,5,24) | ReLU | (5,5,24) |
| 11 | Conv3D(Residual from layer 9 as an input+BN) | (3,3,24) | (1,1,1) | (5,5,24) | ReLU | (5,5,24) |
| 12 | Pooling Layer(average) | | | (5,5,24) | x | (1,1,24) |
| 13 | Fully Connected Layer | | | (1,1,24) | softmax | (1,1,L) |

2.4 Neuromorphic Convolution Neural Networks (CNNs)

Neuromorphic processors implement CNNs with dramatically reduced size, weight, power and cost (SWaP-C) compared to GPU versions. Here we describe converting the 3D CNN into a format that can be run on neuromorphic platforms. We had early access to BrainChip’s software developer kit (SDK) and simulator thus we focused our efforts

on this. We now have access to Intel Neuromorphic Research Consortium and are using it for other projects [11].

BrainChip can support many features of CNNs but not all. For example, it can only accept grayscale or RGB images, not hyperspectral images (HSIs) for convolutional input layers. (For regular input layers, it may be possible to input HSIs but only 4-bit precision can be used at this time.) Because of this, we had to remap the 61 bands of the HSI image into separate “grayscale” input channels and then fuse across input channels in groups. Furthermore, the skip connections in the original 3D CNN are implemented by copying activation values from one neural processor unit (NPU) to another, and then copying them to the original NPU with identical weights of the value 1. This was the recommendation from BrainChip. The Akida™ chip has 80 NPUs so using a handful of extra NPUs to implement the skip connections would not prevent neuromorphic implementation [12]. In Figure 5, we show the translated CNN compatible with the BrainChip hardware.

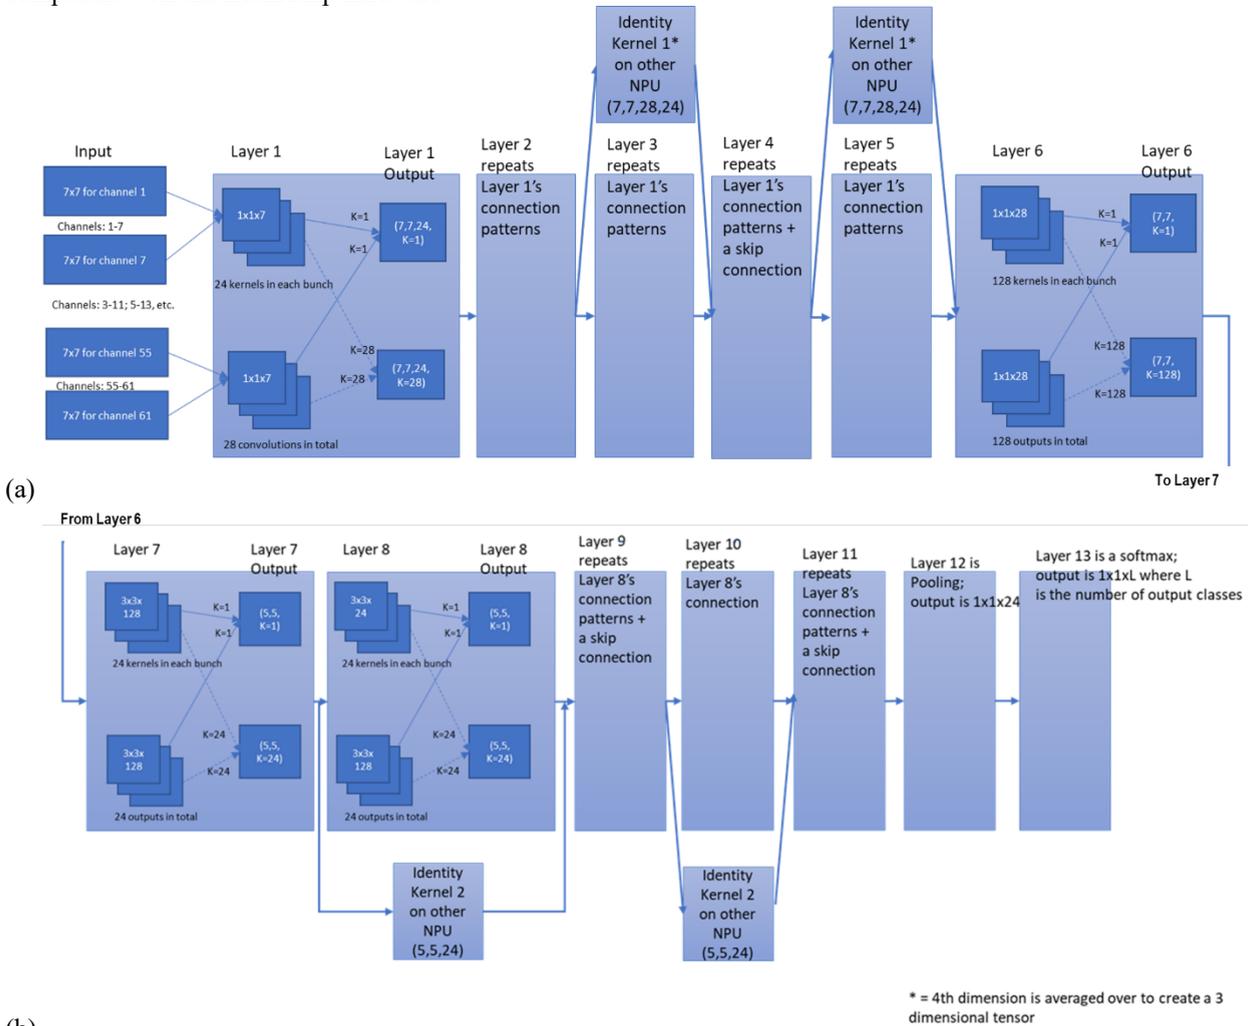


Figure 5. BrainChip translation of the 3D CNN. (a) Layers 1 to 6 and (b) layers 7 to 13.

Using the BrainChip CNN2SNN simulator, TensorFlow layers in the original 3D CNN were simulated with BrainChip compatible versions. That allows us to simulate the hardware on CPU, without the actual hardware. We then trained the network and tested it on HSI images of fentanyl, methyl salicylate and SF96 with different backgrounds, temperatures and purities or coverage amounts as in section 2.2. As before, twenty percent of the data was used for training and the remainder for testing. We compared this result with the train/test results from the original full precision CNN performing classification. We used a different sample of data to estimate the variability of the full

precision CNN.

3. Results

3.1 Simulation of Hyperspectral Imaging (HSI) Spectra

The emittance spectra for the various chemicals at various conditions are shown in Figure 6. Furthermore, some sample synthetic images are shown in Figure 7. Please note that the temperatures in Figure 6 are representative of cold, moderate and hot temperatures. We also have the ability to simulate very cold or very hot temperatures.

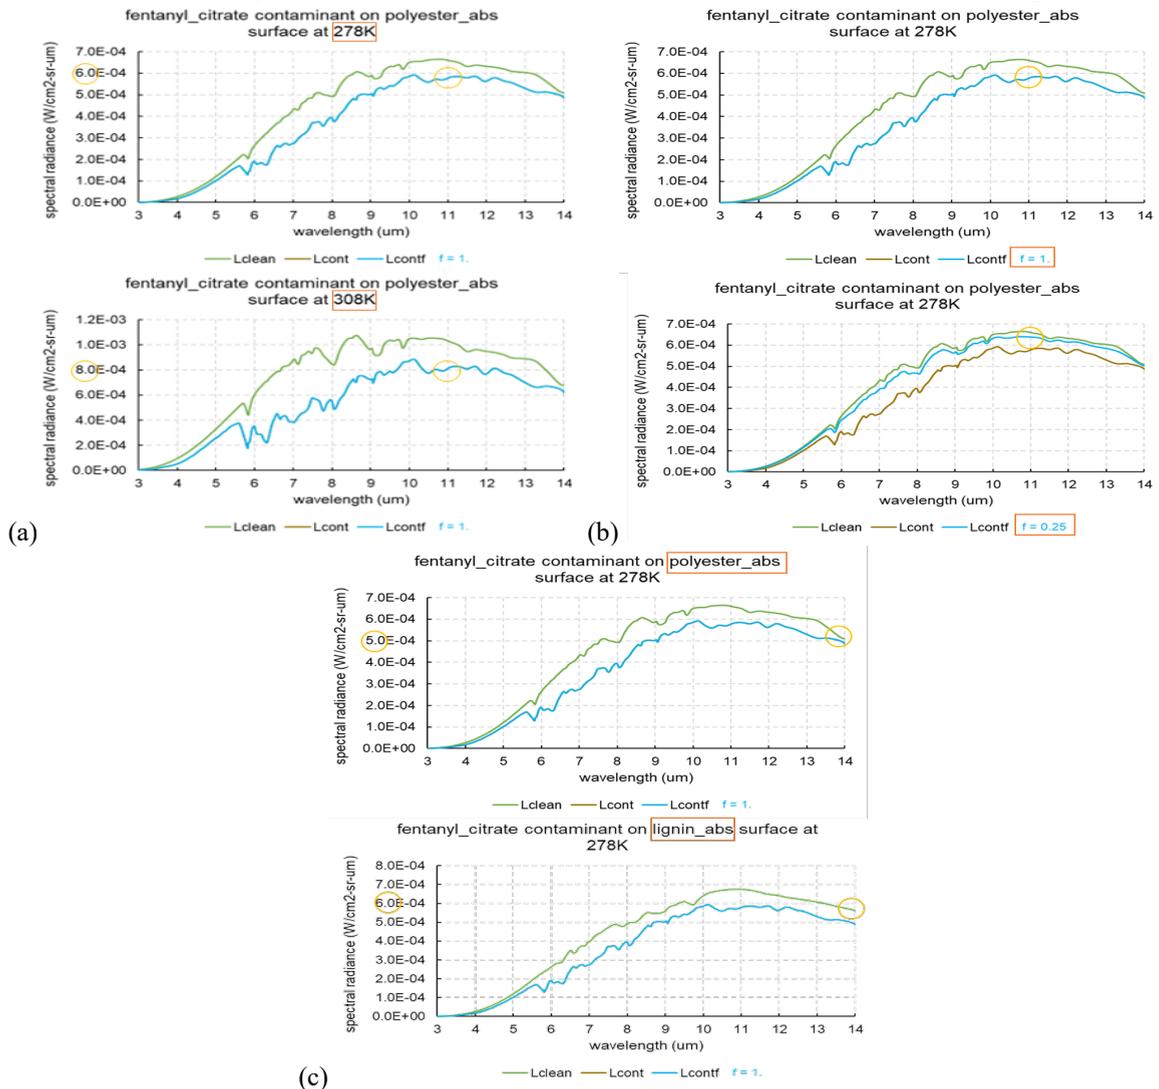


Figure 6. (a) Simulation of chemicals at different temperatures. On the top, we show fentanyl at 278 K and on the bottom, 308 K. Notice that at 11 um, on the top, the reflectance of fentanyl is $\sim 6 \cdot 10^{-4}$ and on the bottom, $\sim 8 \cdot 10^{-4}$. (b) Simulation of chemicals at different purities or coverage factors. On the top, we show fentanyl at 100% coverage and on the bottom, 50% coverage. Notice that at 11 um, on the top, the reflectance of fentanyl $\sim 5.5 \cdot 10^{-4}$ and on the bottom, $\sim 6.5 \cdot 10^{-4}$. (c) Simulation of chemicals on different backgrounds. On the top, we show fentanyl on polyester and on the bottom, on lignin, a major component of cardboard. Notice that at 14 um, on the top, the reflectance of fentanyl $\sim 5 \cdot 10^{-4}$ and on the bottom, $\sim 6 \cdot 10^{-4}$. On each plot, only one spectral difference is described although there are many more which can be seen with close inspection.

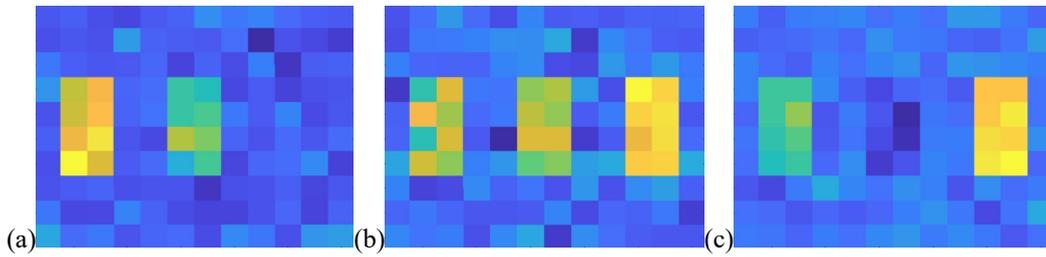


Figure 7. Synthetic images generated from emittance spectra shown in Figure 6, assuming a normal distribution with a nominal standard deviation of $1e-5$ representing camera noise. False color images of (a) lignin background, temperature 278 K, wavelength 7200 nm; (b) polycarbonate (BPA) background, temperature 293 K, wavelength 8000 nm; and (c) polyester background, temperature 308 K, wavelength 9500 nm.

3.2 Full Precision GPU Convolutional Neural Networks (CNNs)

The results are presented in Table 2 to Table 4. Results averaged over the 27 conditions are promising: 64% correct for fentanyl, 60% for methyl salicylate, and 81% for SF96.

Table 2. Average classification accuracy by background. Results are roughly the same across backgrounds: 70.5% average percent correct across three contraband chemicals for lignin, 67.9% for polyester and 66.3% for plastic.

| Contraband\Background | Lignin | Polyester | Plastic |
|-----------------------|--------|-----------|---------|
| Fentanyl | 65.6% | 57.5% | 68.8% |
| MS | 62.8% | 67.2% | 50.0% |
| SF96 | 83.1% | 78.7% | 80.2% |

Table 3. Average classification accuracy by temperature. Results decrease at high temperatures: 70.3% average percent correct across three contraband chemicals for 278 K, 72.8% for 298 K, and 62% for 308 K.

| Contraband\Background | 278 K | 293 K | 308 K |
|-----------------------|-------|-------|-------|
| Fentanyl | 64.0% | 70.4% | 57.5% |
| MS | 62.2% | 63.9% | 53.9% |
| SF96 | 84.5% | 84.1% | 73.4% |

Table 4. Average classification accuracy by purity (coverage). Results are roughly the same for different purities: 69.8% average percent correct across three contraband chemicals for 100% purity, 66.6% for 50% purity and 68.2% for plastic.

| Contraband\Background | 100% purity | 50% purity | 25% purity |
|-----------------------|-------------|------------|------------|
| Fentanyl | 63.5% | 65.1% | 63.5% |
| MS | 67.8% | 48.3% | 63.9% |
| SF96 | 78.3% | 86.5% | 77.3% |

3.3 Neuromorphic Convolutional Neural Networks (CNNs)

On a second dataset, we found neuromorphic algorithms produced an average accuracy of 65%, 56% and 55% for fentanyl, methyl salicylate and SF96 in that order, all of which are promising. The full precision CNN, on the other hand, got 70%, 86% and 75% correct for the same chemicals. In Table 5, we show more detailed information via confusion matrices.

Table 5. Performance comparison of full precision implementation versus the neuromorphic version. Rows are the true label and columns are predicted label. Rows 1 to 4 are background pixels, fentanyl, MS and SF96 in that order. (E.g., for the first matrix, there was a total of 4086 correctly labelled background pixels; however, 15 background pixels were labelled as fentanyl, 8 as MS, and 7 as SF96.)

| Confusion matrix of full precision GPU implementation | | | | Confusion matrix of neuromorphic implementation | | | |
|---|-----|-----|-----|---|-----|-----|-----|
| 4086 | 15 | 8 | 7 | 4085 | 1 | 1 | 2 |
| 0 | 394 | 38 | 43 | 0 | 365 | 141 | 176 |
| 0 | 114 | 508 | 96 | 1 | 60 | 332 | 85 |
| 0 | 38 | 40 | 441 | 0 | 135 | 120 | 324 |

We roughly estimate the power draw of one chip to be between 100 mW and 1 W based on BrainChip’s implementation of a network to process the CIFAR dataset. In addition, BrainChip’s processor may fit into USB keys, making them amenable to deployment in handheld/portable units which DHS seeks. On the other hand, the 3D CNN run on a GPU would require much more SWaP-C. In Table 6, we detail the SWaP-C differences between CPU/GPU platforms and neuromorphic versions.

Table 6. Comparison of SWaP-C profiles for CPU/GPU platforms and neuromorphic versions

| | CPU/GPU platforms | Neuromorphic versions |
|------|---|---|
| SWaP | NVIDIA A100 26.7 long x 11.2 tall x 3.5 wide cm (10.5 x 4.4 x 1.4 in) 250 W | USB key form factor 5.1 long 1.3 tall x 0.6 wide (2 x 0.5 x 0.3 in) 1 W |
| Cost | \$30,000 USD (est.) | \$50 USD (est.) |

4. Discussion and Conclusion

We have demonstrated proof of concept for a high throughput non-ionizing screening capability based on long-wave infrared (LWIR) hyperspectral imaging (HSI), convolutional neural networks (CNNs) and neuromorphic computing as part of DHS Phase 1 SBIR. With good accuracy, LWIR HSI and CNNs can detect and classify contraband in a range of situations, including various backgrounds, temperatures, and purities. Furthermore, neuromorphic computing implementations of CNNs offer promising results with a low size, weight, power and cost (SWaP-C) envelope but can have less accuracy than full precision versions. This change in performance is likely due to the limited precision of weights and neurons in the neuromorphic chips and can be mitigated by hyperparameter tuning, which we are exploring in a DOE cybersecurity project [11].

One could pursue improvements along a number of dimensions. For example, on the sensing end, compressive sensing techniques could be used to reduce SWaP-C of the imager. Non-contact ultrasound could be used to examine the interior of objects without ionizing samples. (Terahertz imaging or RF-backscatter are other options however they cannot see through metal such as car doors, which is common in Ports of Entry). In terms of HSI synthesis, one could also use generative adversarial networks (GAN) [13, 14]. To deal with other types of purity, we could detect constituent functional groups with molecules and additional machine learning [15]. Finally our full precision and neuromorphic CNN results are preliminary. There is, for example, degradation in the classification accuracies at high temperatures, which merits further investigation. Hyperparameter optimization of the full precision and neuromorphic CNNs is likely to yield better results [11]. Testing on other chemicals and environmental conditions could also be conducted. Finally we mention the work Chelian et al. who simulated multispectral image processing from retinal [16] and thalamic [17] circuits in the spiking domain under the DARPA SyNAPSE program. These architectures provide spatial, spectral and temporal features such as noise suppression, ratios of spectral bands and early motion processing. These features are useful in multispectral or hyperspectral classification, and were informed by studies in the rate coded domain [18, 19]. Their work is especially well suited for low SWaP-C neuromorphic platforms.

Acknowledgements

This material is based upon work supported by the U.S. Department of Homeland Security under Award Number should DHS211-009-0011. The views and conclusions contained hereon are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DHS, or the U.S. Government.

References

- [1] Department of Homeland Security. (2021, May 14). Realtime, rapid detection of illegal substances using hyperspectral and neuromorphic technologies. <https://www.sbir.gov/node/2005799>.
- [2] Japanese Pharmacopoeia. (2011). Society of Japanese Pharmacopoeia.
- [3] SpectraBase. (2021, Aug. 15). SF96-100. <https://spectrabase.com/compound/IVKJTHjmjvz>.
- [4] NIST. (2021, Aug. 15). Methyl salicylate. <https://webbook.nist.gov/chemistry/name-ser/>.
- [5] Stark, N., D. Yelle, and U. Agarwal. (2015). “Techniques for characterizing lignin.” In O. Faruk and M. Sain (Eds.). *Lignin in polymer composites* (pp. 49-66). Elsevier.
- [6] Shekhawat, N., A. Sharma, S. Aggarwal, and K. Nair. (2011). Refractive index engineering in polycarbonate implanted by 100 keV N⁺ ions. *Optical Engineering*, 50 (4), 044601.
- [7] SpectraBase. (2021, Aug. 15). Polyester - Database of ATR-FT-IR spectra of various materials. <https://spectra.chem.ut.ee/textile-fibres/polyester/>.
- [8] Bodkin, A., A. Sheinis, A. Norton, J. Daly, S. Beaven, and J. Weinheimer. (2009). Snapshot hyperspectral imaging: the hyperpixel array camera. In *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XV* (Vol. 7334, pp. 164-174). SPIE.
- [9] Department of Defense - Navy. (2021, Jan 12). Multispectral detection of small UAVs. <https://www.sbir.gov/node/2174741>.
- [10] Zhong, Z., J. Li, Z. Luo, and M. Chapman. (2017). Spectral-spatial residual network for hyperspectral image classification: A 3-D deep learning framework. *IEEE Transactions on Geoscience and Remote Sensing*, 56 (2), 847-858.
- [11] Zahm, W., T. Stern, M. Bal, A. Sengupta, A. Jose, S. Chelian, and S. Vasan (2022). Cyber-Neuro RT: Real-time Neuromorphic Cybersecurity. In *Proceedings of the 13th Annual International Conference on Biologically Inspired Cognitive Architecture (BICA 2022)*, Guadalajara, Mexico.
- [12] BrainChip, Inc. (2022, April 25). CNN2SNN Toolkit. https://doc.brainchipinc.com/user_guide/cnn2snn.html.
- [13] Audebert, N., B. Le Saux, and S. Lefèvre (2018). Generative adversarial networks for realistic synthesis of hyperspectral samples. In *IGARSS 2018 IEEE International Geoscience and Remote Sensing Symposium* (pp. 4359-4362). IEEE.
- [14] Hennessy, A., K. Clarke, and M. Lewis (2021). Generative adversarial network synthesis of hyperspectral vegetation data. *Remote Sensing*, 13 (12), 2243.
- [15] Xu, M., C. Wang, A. Terracciano, A. Masunov, and S. Vasu. (2020). High accuracy machine learning identification of fentanyl-relevant molecular compound classification via constituent functional group analysis. *Scientific Reports*, 10 (1), 1-10.
- [16] Chelian, S. and N. Srinivasa. (2015). System, method, and computer program product for multispectral image processing with spiking dynamics. (U.S. Patent No. 9,111,182). U.S. Patent and Trademark Office.
- [17] Chelian, S. and N. Srinivasa (2016). Neuromorphic image processing exhibiting thalamus-like properties. (U.S. Patent No. 9,412,051). U.S. Patent Trademark Office.
- [18] Carpenter, G.A. and S. Chelian. (2013). DISCOV (DIimensionless Shunting COlor Vision): A neural model for spatial data analysis. *Neural Networks*, 37, 93-102.
- [19] Chelian, S. (2006). *Neural models of color vision with applications to image processing and recognition*. Boston University. <https://www.proquest.com/openview/41ac5c06c80fffd0631e94c26245f00d/1>.