Synthetic Imagery Generation

Steven Webster
U.S. Army Combat Capabilities Development Command (DEVCOM) C5ISR Center
UNITED STATES OF AMERICA

Mark Jeiran
U.S. Army Combat Capabilities Development Command (DEVCOM) C5ISR Center
UNITED STATES OF AMERICA

Kimberly Manser
U.S. Army Combat Capabilities Development Command (DEVCOM) C5ISR Center
UNITED STATES OF AMERICA

Neil Pinto
U.S. Army Combat Capabilities Development Command (DEVCOM) C5ISR Center
UNITED STATES OF AMERICA

Susan Harkrider
U.S. Army Combat Capabilities Development Command (DEVCOM) C5ISR Center
UNITED STATES OF AMERICA

ABSTRACT

Artificial Intelligence (AI) technologies are increasingly being applied to military challenges. One of the many applications is Command and Control (C2) of maneuver and fires within a coordinated multi-domain battlespace. In order for this AI to be effective for both on-board and remote imaging sensing, the associated models must be trained and iteratively improved, often with limited knowledge of the imminent threat. Synthetic imagery generated using modeling and simulation fills the gap between the immense training dataset required for AI model development and the limited sensor data available from live collections.

We synthesize Infrared (IR) imagery by combining physics models, image generators and game engines with modern infrastructure. The Night Vision Integrated Performance Model (NVIPM), the foundational physics model, was automated to generate varying performance curves that underpin the imagery. We initially extended our Night Vision Image Generator (NVIG) with editing controls to create scenes with known terrain, clutter, environment, and targets. Subsequently, that IR representation was applied to Unity and Unreal game engines, leveraging those frameworks for similar scene composition. Improved application program interfaces (APIs) enabled greater imagery variations and counts with automation. These advances in simulation, modeling, and our "Sensor as a Service" platform design are discussed.

1.0 MOTIVATIONS

Artificial Intelligence (AI) is increasing used to solve challenges in the Military environment, particularly in Command and Control (C2), Effects, and Intelligence Surveillance and Reconnaissance (ISR) domains. One of the most prevalent uses for AI in civilian and military applications is image recognition, where items of interest are classified and ideally recognized. One military image recognition challenge lies within the Sensor-to-Shooter engagement, described as "detect, decide, deliver and assess," where AI detection reduces tactical workloads and improves execution timelines [1]. Addressing the sensing components within Sensor-to-Shooter, we perform Aided Target Recognition (AiTR) using thermal sources, since thermal imagery enables military operations at any time and within Degraded Visual Environments (DVE). We focus on imagery rather than video since algorithms supporting target detection and classification do not have a temporal correlation component. Eventually, additional tasks such as tracking will require temporal understanding, but the tools and techniques developed for single images should apply without significant modification.

AI model development requires thousands of images tagged with truth information, such as an outline of a tank within a background environment. The most representative source data comes from live collections, where for a given location, the thermal imaging sensors are emplaced, targets are positioned, the physical environment (temperature, humidity, time of day, ...) is recorded, and then imagery is collected. Afterwards, the imagery is tagged with the truth metadata, which involves human review and updates. While this approach results in high quality content, it is resource intensive as multiple professionals are involved, collection locations must be sourced, and logistics such as travel and shipping supported. In addition, the collection itself is at risk from weather, equipment failure, and any logistical challenges. Finally, the imagery collected represents a narrow set with limited targets, environments, and sensors. Therefore, multiple collections are needed before a significant variety in the AI data set is realized. Even though these collections have been ongoing for years, the vast quantities of imagery data for AI training and verification motivates additional data generation techniques.

The required characteristics of synthesized imagery varies greatly with the AI model under development. Before we can realize more generalized capabilities, we currently focus on low-level tool chains and the automation of those applications to shorten timelines for AI researchers and application developers. Other work pertaining to on-orbit visual navigation applications [2] takes a similar approach, where a pipeline comprises image generation tools, scales as required, and delivers appropriate imagery for deep learning model development. They emphasize commodity AI and compute from cloud providers, aiming for a portable solution with low financial requirements. Another approach [3] uses parallel computing clusters to train algorithms for target detection within satellite imagery. Their pipeline includes Digital Imaging and Remote Sensing Image Generation (DIRSIG) [4], a high fidelity model we also use to generate thermal imagery. However, their use of parallel clusters makes this type of pipeline less approachable since many researchers and developers cannot realize their compute performance without access to these specialized clusters. We align with the former approach, providing and automating image generation tooling, hostable at available scale, which leverages open source clustering technologies.

2.0 IMAGERY APPROACHES

2.1 Fully Synthetic

Fully synthetic imagery, created by image generators and gaming engines, is one method that addresses this data shortfall. Highly varied environments can be synthesized, limitless targets represented, and truth tagging integrated, making large scale source data generation theoretically possible. There are challenges in this approach, as AI trained and verified using synthesized content may not perform adequately when applied to a live context. On-going science and technology efforts are addressing the relevant characteristics within the generated content that will result in successful transference from synthetic to live domains. Those fidelity concerns will be addressed as discovered, so the focus for fully synthetic imagery is to build user friendly tooling and automating operations where applicable.

Early work generating fully synthetic imagery generation leveraged the internally developed the Night Vision Image Generator (NVIG) which is predicated upon the Night Vision Integrated Perform Model (NV-IPM). NV-IPM is a first principles physics model, which generated performance characteristics along with noise and blur kernels for imaging sensors [5]. NVIG is a custom scene generator parameterized from NV-IPM performance outputs, and it provides representative image and video captures with backgrounds for the NV-IPM modeled sensor [6]. Creators using NVIG can stage a 3D scene in an appropriate environment using models for targets of interest as well as clutter. NVIG renders this scene using sensor effects that are representative of a sensor type described by the associated NV-IPM model, and captures the resulting image. This process repeats for all desired target aspects, and the resulting image set is analogous to a live collection. The process repeats for all specified locations (backgrounds), target models, sensor types, and environment conditions. Larger scale production can be automated using a Domain Specific Language (DSL) NVIG declaration to define and iterate over these permutations. NVIG can also persist a model identifier, per pixel, per image, so that the truth tagging of valued targets and clutter is part of image capture. The appeal of using NVIG is its high fidelity thermal representation, which originates from its use of temperature-based textures. However, AI users have also used game engines for these purposes.

Many users apply game engines to generate synthetic imagery, however most work involves visible spectrum imagery. As a pathfinder project, we ported NVIG thermal sensor effects to the Unreal and Unity engines to understand the challenges associated with this capability. Both engines successfully rendered the targets under different thermal conditions, but we did not develop representative thermal backgrounds. NVIG sensor effects also rely upon temperature base textures, and those textures do not exist for game engine content. There were additional challenges with ambient and location based lighting, which is applicable for the engine's native visible spectrum rendering but not applicable to thermal sensor behaviors. Other scope that remained are User Interface (UI) controls and associated automation to stage, render, and capture with parametrically defined thermal sensor effects. Finally, packaging these thermal behaviors and UI as a game engine add-on became difficult because of the coupling between content (models, textures) and core game capabilities (lighting, rendering).

Other team members also synthesized thermal images using the Unreal engine with third party tooling [7]. While early work had leveraged manual runs of DIRSIG, the Unreal based capability enabled orders of magnitude improvements in runtime performance, however with a loss of fidelity. They have updated their Unreal engine representation with improvements in target and clutter models, dynamic range, anti-aliasing, and physical environment. In addition, the high adoption and availability of the Unreal engine facilitated the development of UIs, which eased parameterization of the models, environments, and runtime configurations. Automation of imagery generation was through Python scripting, so that concurrent runs were possible. Figure 1 shows sample imagery from their pipeline.

Figure 1: Sample photo-realistic images using the Unreal Engine 4 Pipeline.

Furthermore, a third-party algorithm developer was employed to systematically infuse Unreal Engine generated datasets into their image-classification algorithm's training pipeline. The quality of these datasets is evaluated based upon the algorithm's performance on detection or classification of real field imagery. While the initial results were somewhat disappointing, the Unreal Engine pipeline was drastically improved in a short period of time via feedback driven development. This process also helped to understand what provided the most training value to the algorithm. As illustrated in Figure 2, the initial unreal images suffered from aliasing and other synthetizing artefacts. While those artefacts were corrected for in V1.2 of the pipeline, these images still have a very narrow dynamic range. This limited the amount of information that the algorithm was able to extract from the image, hence limiting its performance. While variations in the target's thermal signature was highly desirable, as illustrated, the relative contrast of target to the background could fall out of the realistic range and cause the target to be too hot or too cold for the background. V2.0 of the pipeline had all of these inaccuracies corrected, but the scene clutter was very simplistic and lacked deeper variations. This was drastically improved in V3.0 of the dataset with a wide verity of realistic foliage, man-made objects, and even confuser targets such as civilian vehicles. Overall the introduction of synthetic data to the algorithm was shown to significantly improve its performance. Also, in cases of limited field data, synthetic data was shown to be an adequate replacement.

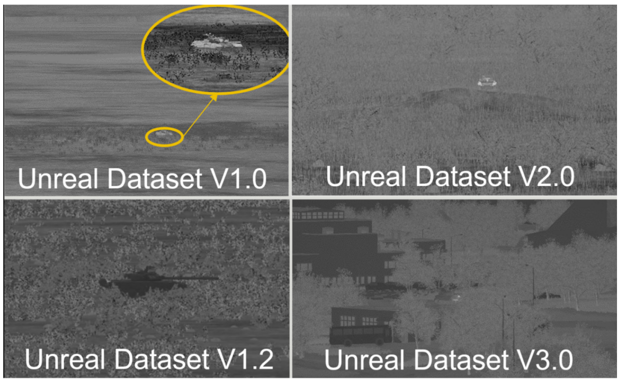

Figure 2: Sample Unreal Engine simulation fidelity and complexity improvements.

2.2 Hybrid, Augmentation, and Manipulation

In addition to full virtual image generation, composite or hybrid imagery is another approach to generating sufficient data sets for AI development. This method combines virtual and live components within an image, using virtual renderings for the target and inserting those pixels into a real background image. The live background is attractive because it includes proper variation of a large set of features that would be difficult to replicate synthetically. Conversely, the virtual targets are attractive because it becomes trivial to insert a large variety of types and poses into a scene, which is challenging to do in a live collection. Subtle variations can be included within the virtual target, such dents, impact marks, mud, or wear, improving the variation within the imagery set. The use of virtual models also allows target representations that are not accessible, such as future designs or platforms outside of our allies' possession. Finally, target truth data can be specific to target sub component, which enables deep learning algorithms to infer what components imply a given target.

Beyond the usage of synthetic targets, augmentation and manipulation of a source image results in additional variations. The ideal start is high fidelity pristine image, which we then degrade to represent an actual sensor capture. For example, post processing a higher fidelity image capture with additional noise and blur represents a less capable thermal image capture. This approach is similar to the fully virtual capture, where the sensor effects are applied to a pristine rendering of all the geometry in the scene. The difference is the source for the degradation, where the image manipulation / augmentation starts with a live target capture.

Figure 3 demonstrates this technique with live imagery, where a manipulated reference image represents sensor performance at range with various resolutions. Other useful manipulations and augmentations are common image transforms, such as resampling, rotation, and cropping. NVPyPM, a fully releasable application, implements these degradations, manipulations, and augmentations.

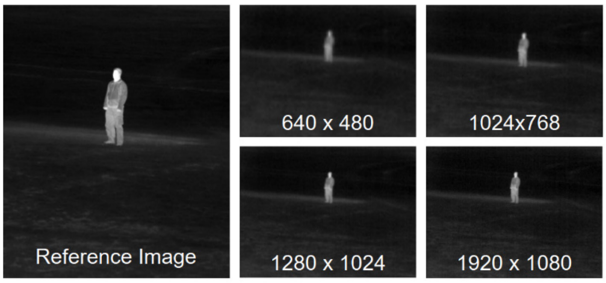

Figure 3: Degradation for resolution at range.

The composite image process uses segmented targets, either from synthetic or real source imagery, and inserts the target into a background. The background can be either real or synthetic, though our research has shown promising success using real backgrounds from multiple sources, including rural scenes and images from open-source databases such as COCO [8]. The target insertion process utilizes image processing techniques to improve relative contrast, add representational blur, and blend the edges of the target into the background effectively without introducing significant frequency artifacts. In some cases, simple cropping is also employed to simulate landscape occlusion, such as in the examples in Figure 4.

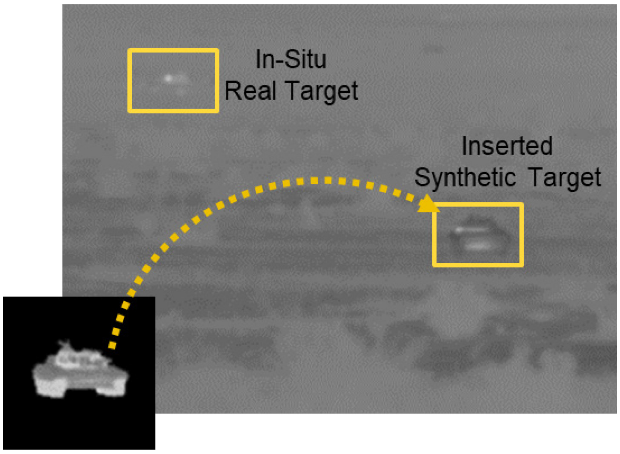

Figure 4: Example of synthetic target in real background.

2.3 Algorithm Training Results for Sensor Emulation

To provide context for the value of sensor emulation and the proposed "sensor as a service" capability, researchers at the C5ISR Center have conducted a sensitivity analysis regarding the amount and severity of degradations in training and test sets used for AI/ML algorithm evaluation. The analysis focuses on two modes of variation in the sensor settings: prevalence and severity. The prevalence of image degradation in a dataset refers to how many of the images are degraded, e.g. 25%, while keeping the severity of the degradation (size of the blur and magnitude of the noise) constant. A real-world analogy for this is an intermittent degraded visual environment (DVE) condition in the field, such as a sand storm. In this experiment, it's as if there is a sandstorm 25% of the time, and the sandstorm has a predictable and repeatable intensity. The severity of the image degradation is varied while keeping prevalence constant in the second part of the analysis. This is akin to having a sandstorm every day, but each sandstorm has varying intensity such that some DVE conditions are more severe than others during operation. For these experiments, Defense Systems Information Analysis Center's (DSIAC's) Automated Target Recognition (ATR) Algorithm Development Image Database [9] was used (which includes IR imagery of military-relevant targets) along with YOLOv5 [10] for evaluation.

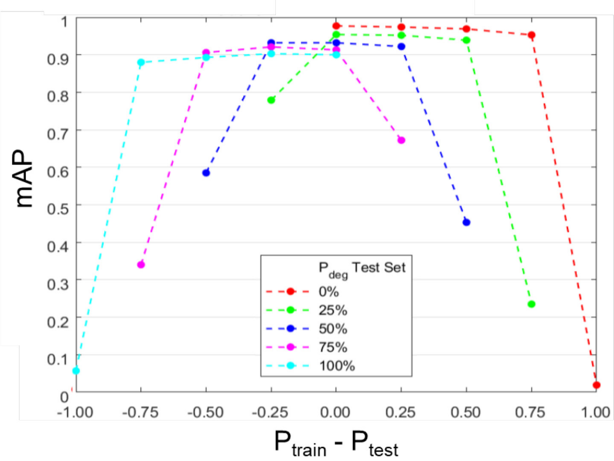

Figure 5: Prevalence mismatch of training and test set degradation.

Figure 5 shows the results for the prevalence mismatch experiment, where the x-axis is the difference between the prevalence of degradation in training set vs the test set, and the y-axis is the mean average precision of the algorithm. The differently colored lines represent different prevalence values, ranging from 0% in red to 100% in cyan. For negative values of x, the prevalence in the test set is higher than the prevalence in the training set, while the inverse is true for positive values of x. As the mismatch between the prevalence of degradation in the training and test sets increases, performance decreases, eventually reaching a steep dropoff on either end of the distribution. The best performance occurs when the prevalence of degradation in the training and test sets are the same, or when x=0.

Figure 6 shows the results for the severity mismatch experiment, where the x-axis is the difference between degradation severity in the training and test sets, and the y-axis is the mean average precision of the algorithm. The degradation is defined in terms of blur (simulation optical aberration) and noise (SigmaVH, or the standard deviation of noise in both vertical and horizontal channels). Blur and noise are paired variations, meaning that each degradation condition is defined by a pair of values that do not vary independently. Blur (aberration) varies from 0-0.2 wavefront error while noise varies from 0-30 digital counts root-mean-square (for 8-bit images). Each colored line represents a different blur and noise pair condition, ranging from pristine imagery in red, to the maximum degradation in cyan. For negative values of x, the test set degradation severity is greater than that of the training set, while the inverse is true for positive values of x. Unlike the prevalence mismatch in Figure 5, the results are not symmetrical for positive and negative values of x, though the maximum performance is still found at x=0. There is much more of a penalty on performance for negative values of x, where there is more severe degradation in the test set than in the training set. When the training set has a higher degradation severity, the performance drop is much smaller.

The conclusion to be drawn from this analysis is that the similarity of training and test sets, in terms of sensor emulation, is very impactful to algorithm performance. If the fielded environment of a sensor is expected to have occasional DVE of varying severity, it is vital to include representative examples of those conditions in the training set to optimize performance. This also applies to pristine synthetic data relative to real-world data. Synthetic data generated by gaming engines or high-fidelity physics simulators often output "perfect" imagery.

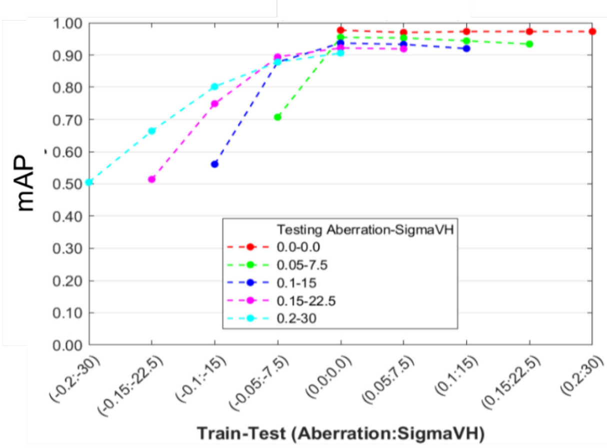

Figure 6: Severity mismatch of training and test set degradation.

3.0 SENSOR AS A SERVICE

"Sensor as a Service" is the collection of these tools that enable user access to the capabilities without requiring configuration of their workstations. Sensor as a Service includes tools as REpresentational State Transfer (REST) web services, after packaging as containerized applications. A Kubernetes [11] cluster hosts and orchestrates these containers, which is available as a cloud hosted service or buildable upon on premise infrastructure. The cluster enables elastic provisioning, scaling the component tools up and down tailored to user loads. While following commercial best practices, this approach also supports the NMSG-164 Modeling and Simulation as a Service (MSaaS) architectural principles, including "Facilitate simulation service composition," "Facilitate M&S resource pooling, "and "Use open standards." [12]

The objective Sensor as a Service composition, shown in Figure 7, supports multiple methods of imaging sensor modeling, simulation, and data production. The architecture addresses the synthetic, augmentation, and hybrid data generation needs previously described while also supporting real-time data services and video generation capabilities. A web client will enable administrative control of the supporting services and application UI to unprivileged users. A REST Application Programming Interface (API) will be available for third-party automation of sensor services.

The initial implementation of Sensor as a Service comprised the re-hosting of NV-IPM as a REST web service. This service reuses NV-IPM's Simplified Interface, Figure 8, by consuming a user provided JavaScript Object Notation (JSON) document, which replicates the associated NV-IPM dialog. Within this dialog, parameters are either mandatory or only used in conjunction with other parameters. Application logic within NV-IPM validates these parameter sets, whereas the REST web service uses JSON schema [13] validation to ensure well-formed input.

Like the NV-IPM application, the associated REST service generates sensor performance data products, including lookup tables for the ACQUIRE model used in One Semi-Automated Forces (OneSAF) and other constructive simulations. After the initial NV-IPM integration, the Sensor As A Service capability was extended to include a real time ACQUIRE implementation, so users can define their sensor characteristics and then utilize the computed performance parameters to generate sensor – target perception outcomes during simulation runtimes.

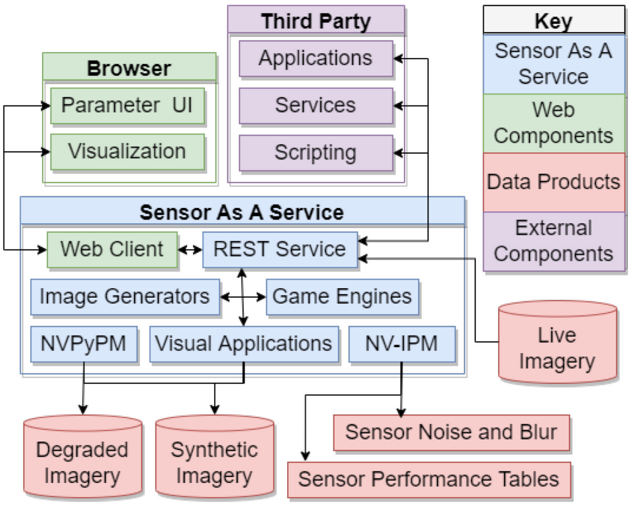

Figure 7: Sensor as a service architecture.

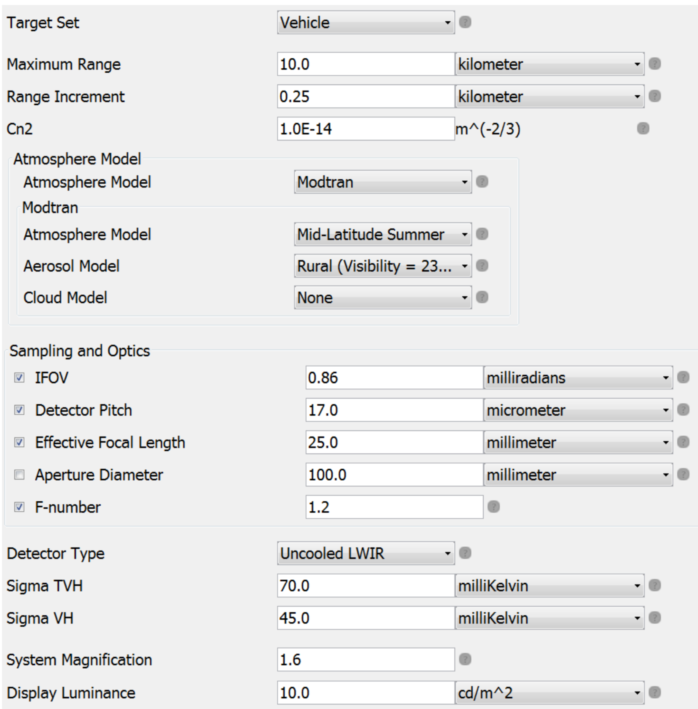

Figure 8: Simplified interface eases configuration.

To add the imagery hybrid composition, augmentation and manipulation capabilities, we will integrate NVPyPM in a similar manner as NV-IPM. That effort includes a re-hosting the application as a REST service, and an additional web UI to reference the source image, the desired manipulations and augmentations, and describe the desired permutations for output imagery. This UI will rely upon the NVPyPM REST API, which will make these same features available for third-party head-less automation.

The Kubernetes cluster scheduler will dynamically provision game engines and the NVIG on capable nodes. These applications require hardware acceleration for graphics processing, so Kubernetes “taint” and “toleration” annotations along with an Admission Controller will schedule containers appropriately [14].

Moreover, these applications require a graphics context, and that context is associated with a graphical user session when hosted on a traditional workstation. Some applications can also require a display, so these graphic applications may require additional integration and / or development efforts to support in the cluster. Finally, generated video services will need remote access for the consumer, which is another technology to layer within the graphical application container. The goal is to support graphic applications in the same manner as other sensor models and simulations so that containerized services benefit from elastic provisioning and resource sharing.

The NV-IPM, NVPyPM, NVIG, and gaming engines comprise the beginning of Sensor as a Service capabilities. The composition of Sensor as a Service will grow in response to user needs where relevant models and simulations will be adapted to the runtime environment. The architecture employs commercial best practices of open APIs and standards, such as REST, JSON and application containerization, which reduce impediments to adopting new applications. The dominant industry practice for container orchestration, Kubernetes, hosts the applications, enabling elastic provisioning and process concurrencies. AI model training and verification can be hosted in a similar manner, so co-hosting sensor models and simulations with the consuming AI technologies will reduce development timelines. Given the flexibility and ease of adaptation for models, simulation, and AI, the Sensor as a Service platform will support our evolving needs.

4.0 MILITARY APPLICATIONS

There is a lot of value across the military community in leveraging an AI powered and modular synthetic image generation capability. Besides training operational AiTR systems, Sensor as a service can be used to inform Acquisition decisions, train Warfighters at the point of need as well as inform and substantiate other sensor-to-shooter technologies such as computing and network communications.

4.1 Acquisition

One of the biggest hurdles in Acquisition is demonstrating value of highly technical and complex systems. Senior leaders need to understand the technologies that are present now, the increase in Warfighter capability with a potential new system and how that translates to cost savings or benefit to the Soldier. Military Capability Developers also have a similar gap in that they need to understand and write requirements based on current gaps while staying within the art of the possible. For example, the average time it takes the U.S. Army to approve requirements is two to three years [15]. Sensor as a service can be used to "show" differences in sensor systems to requirements writers and decision makers as well as the end users for an all-around benefit to the military.

The Enhanced Night Vision Goggle Binocular (ENVG-B) is good example of a revolutionary technology that could have benefitted from Sensor as a Service and its implementation. ENVG-B implements multiple enhancements to dismount night vision goggles including: high figure of merit (FOM) white phosphorus tubes allows the user to see images in a white field instead of the current green background, as shown in Figure 9 [16]. ENVG-B comprises higher resolution stereoscopic displays to allow for faster target acquisition by improving separation of targets from background and wirelessly transmitted weapon sight crosshair and thermal imagery from the Family of Weapon Sights-Individual (FWS-I) [17]. FOM is an abstract measure of image tube performance, derived from the number of line pairs per millimeter multiplied by the tube's signal-to-noise ratio (SNR).

While it may be intuitive that higher resolutions and SNRs are "better" it is not clear to a decision maker, requirements writer or a Warfighter how that actually translates to improved ability to move, shoot or communicate. As previously mentioned, a web client with a simplified interface designed around FOM can allow each of these roles to better understand improvements in performance. The decision maker can apply the sensor characteristics to a constructive simulation and understand the improvement in overall performance in the echelon of interest (platoon, company or higher).

Figure 9: Sample imagery, green vs white phosphor.

The non-quantitative changes in visual sensors and displays are also much harder to communicate without visual representations. For example, the white vs green phosphor and the integration of the thermal weapon sight into the ENVG-B is not obviously a better user experience until one sees results like Figure 9. The requirements writers and end user Warfighters can use video and/or image generation capabilities to understand technical metrics as a visual representation of what the sensor can see in different environments, day or night and with different target sets. These different use cases can all be developed, managed and run from different compositions of the Sensor as a Service Architecture.

It is apparent that multiple sensors must be run concurrently to show performance of visual sensors that include multiple modalities and combinations. Figure 10 [18] represents one such type of design, depicting FWS-I and ENVG-B's Picture in Picture (PnP) feature. As the number and types of sensors that can be viewed at once is increased and multiple target sets and environments are desired in each iteration, the need for scalability and elasticity is clearly demonstrated.

Figure 10: Sample imagery, ENVG-B with FWS-I.

4.2 Targets / Training

In all of the cases mentioned thus far, it has been from the perspective of the observer or end user of the sensor. The target of observation is also critical in understanding the sensor capabilities, whether that be friendly or enemy. A major use case for this is Warfighter training to enable users of sensor systems to better identify and classify observed targets. For example, in November 2021, it was announced that the US

Military would acquire 1,669 Joint Light Tactical Vehicles (JLTV) and of those, 125 vehicles will also be delivered to NATO and allied partners, including Brazil, Lithuania, Montenegro, and Slovenia [19]. The JLTV complies with the US Army's "Long Term Armor Strategy (LTAS) which seeks to build tactical trucks with an A-kit, B-kit modular armor approach, allowing [vehicles] to adjust their protection to the potential threats they will face in combat [20]." An immediate benefit of synthetic data generation would be to predictively understand what a new vehicle looks like in different sensor modalities especially with a variable visible look and varying thermal propagation properties due to customizable armor packages. As stated earlier, it would take a significant amount of personnel and resources to visually record and process all iterations of these vehicles.

Conflicts around the world can also actively prevent the collection of data for training. For example, there is significant NATO support for Ukraine after "Russia's unprovoked full-scale invasion on 24 February 2022 [21]". The Military Engineering Vehicles and Systems High Visibility Project is a coordinated NATO effort to develop and procure various military vehicles, subsystems and modular mission payloads. As systems are researched and developed, they could alter the look of current vehicles in various sensor bands without the ability to obtain ground truth on what they actually look like. Currently, French AMX-10s, German Marders and the U.S. M2 Bradley Fighting Vehicles will be used in Ukraine [22] and they must be integrated into current training systems and otherwise noted known friendly vehicles. This is a direct follow on to prior work relating to improvements in Recognition of Combatants - Vehicle (ROC-V). ROC-V is day/night optics and thermal sight training program that helps Warfighters learn to identify the thermal signatures of combat vehicles based on unique patterns and shapes of vehicle "hotspots," and overall vehicle shapes and characteristics [23]. Specifically, the Sensor as a Service enables "predictive modeling" to incorporate vehicle imagery that is unavailable due to cost or other higher priority uses. It also allows for known vehicles in previously unused environmental settings such as those seen in Ukraine.

4.3 Enabling Emerging Technology

Joint All Domain Command and Control (JADC2) is emerging as the preeminent operational concept. It is intended to improve situational awareness, improve abilities to direct forces across domains and services, and facilitate rapid decision making by leveraging distributed sensors, shooters, and data from all domains [24]. The Advanced Battle Management System (ABMS) is a command and control system that is the central system for JADC2. The ABMS digital infrastructure will move data via data tagging, standardization and other methods. The Sensor as a Service distributed, networked approach would easily facilitate stimulating the ABMS subcomponents by integrating directly into the operational software. The REST API from Sensor as a Service can be tied directly into input nodes to test overall system response, load test as well as provide augmentation to live sensors. For example, ABMS cloud-based command and control currently ingests over 750 radar feeds and fuses them within a single user interface [25]. Sensor as a Service could scale to meet the demand of this entire system without requiring the physical radars, geographic footprint and communications infrastructure to exercise the final system.

5.0 CONCLUSION

Artificial Intelligence is dependent on input data that has been curated, tagged with truth metadata, and includes sufficient variation and degradation for the considered algorithm. While there may not be viable ways to collect data for all sensor types and all algorithms, leveraging synthetic data generation can close the gap much faster while ground truth data is collected. Since this is heavily compute, storage and network driven, cloud resources such as containers with standard APIs will support this demand. Sensor as a Service, a scalable service orientated architecture, enables the flexibility to generate high fidelity one off images to visually show the characteristics of a sensor or generate millions of images with proper metadata to stimulate operational systems and algorithms. This ability to support a wide range of use cases demonstrates its importance to every part of Defense organizations around the world.

6.0 REFERENCES

[1] https://www.defensedaily.com/armys-project-convergence-proving-next-gen-sensor-shooter-network-possibilities/army/

[2] Schubert, C., Black, K., Fonseka, D., Dhir, A., Deutsch, J., Dhamani, N., ... & Akella, M. (2021, March). A Pipeline for Vision-Based On-Orbit Proximity Operations Using Deep Learning and Synthetic Imagery. In 2021 IEEE Aerospace Conference (50100) (pp. 1-15). IEEE.

[3] Han, S., Fafard, A., Kerekes, J., Gartley, M., Ientilucci, E., Savakis, A., ... & Rovito, T. (2017, May). Efficient generation of image chips for training deep learning algorithms. In Automatic Target Recognition XXVII (Vol. 10202, pp. 15-23). SPIE.

[4] DIRS Laboratory, \DIRSIG Reference Docs." http://www.dirsig.org/docs/new/intro.html. (Late Updated: 21 July 2016).

[5] Brian Teaney, Joseph Reynolds, "Next generation imager performance model," Proc. SPIE 7662, Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXI, 76620F (22 April 2010); https://doi.org/10.1117/12.850876

[6] May, Christopher; Harkrider, Susan; Paul, Douglas; Osborn, Walter; Tran, Vinh, "Target IR Signature Creation Used in NVIG Sensor Simulation" Proc. SimTecT 31 May - 2 June (2011).

[7] Mark Jeiran et al, "Unreal Engine 4 Synthetic Data Generation Pipeline for AiTR Training," in Military Sensing Symposia, Orlando, FL, 2023.

[8] https://cocodataset.org/#home

[9] https://dsiac.org/databases/atr-algorithm-development-image-database/

[10] https://ultralytics.com/yolov5

[11] https://kubernetes.io/

[12] North Atlantic Treaty Organization, STO-TR-MSG-164-RA, "Modelling and Simulation as a Service, MSaaS Technical Reference Architecture"

[13] http://json-schema.org/specification.html

[14] https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/#example-use-cases

[15] Purtiman, R. (2018, September). Soldier Lethality CFT Bringing Next Generation Technologies to Soldiers [Review of Soldier Lethality CFT Bringing Next Generation Technologies to Soldiers]. Infantry, 3–4. https://www.benning.army.mil/infantry/magazine/issues/2018/Jul-Sep/PDF/JUL-AUG18_INFMAG_final.pdf

[16] https://www.ainonline.com/aviation-news/business-aviation/2015-09-21/latest-night-eyes-white-phosphor-nvgs

[17] USAASC, Enhanced Night Vision Goggle - Binocular (ENVG-B). Retrieved from https://asc.army.mil/web/portfolio-item/enhanced-night-vision-goggle-binocular-envg-b/

[18] https://taskandpurpose.com/tech-tactics/china-military-augmented-reality-system-weapons/

[19] OSHKOSH DEFENSE (2021, November 29). OSHKOSH DEFENSE RECEIVES $591.6 MILLION JLTV ORDER TO CONTINUE FILLING CAPABILITY GAPS FOR U.S. MILITARY AND INTERNATIONAL ALLIES. Oshkosh Defense. https://oshkoshdefense.com/oshkosh-defense-receives-591-6-million-jltv-order-to-continue-filling-capability-gaps-for-u-s-military-and-international-allies/

[20] Osborn, Kris. Army crafts tailorable tactical wheeled vehicle acquisition strategy. (August 13, 2010). Retrieved from https://www.army.mil/article/43617/army_crafts_tailorable_tactical_wheeled_vehicle_acquisition_strategy/

[21] NATO. Military Engineering Vehicles and Systems Factsheet. Retrieved from https://www.nato.int/topics/mcc/Factsheet-Military-Engineering-Vehicles_en.pdf

[22] Erlanger, S., & Gibbons-neff, T. (2023, January 6). New Armored Vehicles will help Ukraine take the fight to Russia. The New York Times. Retrieved from https://www.nytimes.com/2023/01/06/world/europe/ukraine-war-armored-vehicles.html

[23] Holmes, B., Harkrider, S., & McGlynn, L. (2016). Using Serious Games to Enhance Recognition of Combatants Training for Electro Optic and Infrared (EO/IR) Sensors. In STO-MP-MSG-133. NATO S&T Organization. https://www.sto.nato.int/publications/STO%20Meeting%20Proceedings/STO-MP-MSG-133/MP-MSG-133-04.pdf

[24] Marler, Timothy, Carra S. Sims, Mark Toukan, Ajay K. Kochhar, Shawn Cochran, Christine Kistler LaCoste, and Matt Strawn, Assessment of Joint All Domain Command and Control Requirements and the Use of Live, Virtual, and Constructive Capabilities for Training. Santa Monica, CA: RAND Corporation, 2023. https://www.rand.org/pubs/research_reports/RRA985-2.html

[25] US Air Force. (January 9, 2023) ABMS moves forward on cloud-based C2 Retrieved from https://www.af.mil/News/Article-Display/Article/3262645/abms-moves-forward-on-cloud-based-c2/