SAND2020-2003C

The Use of Synthetic 3D Images to Drastically Reduce Real-World Training Data for Object Detection Models

Zoe N. Gastelum, Timothy M. Shead, Michael Higgins
zgastel@sandia.gov; (505) 845-1002

Overview: Photos from hand-held cameras that are shared via social networks, blogs, or news media are an increasingly abundant information source, and are part of an all-source approach to nuclear proliferation detection. The quantity of available information makes it unrealistic for analysts to search for potentially-relevant images without assistance. Deep learning object detection models are a promising option for this task but require significant domain-relevant training prior to be usable for this domain. Deep learning object detection models train on thousands to millions of example images to accurately locate and identify the object of interest. These models can have high levels of accuracy - sometimes better than human analysts - at recognizing objects of interest if they are well represented in the training data. Most deep learning models for object detection that are openly available today do not include content relevant for nuclear proliferation detection. Even using a pre-trained model and fine-tuning for proliferation-relevant images requires thousands of images. In the domain of nuclear proliferation detection, that quantity of images might not exist, may be too sensitive to use in model training, or may lack ground truth labels (which are time-consuming and potentially error-prone to assign).

In this project, we are developing parameterized three-dimensional computer models of objects of interest to generate unlimited numbers of synthetic two-dimensional images for training. This approach allows us to:

- Vary the object of interest within known physical limits, modifying size, shape, material, etc.
- Vary the background in which the object is situated, or eliminate the background entirely
- Determine what, if any, false cues the model is learning and generate additional images to counter-balance our training set

Our goal is to experimentally develop an approach which will allow us use synthetic images to train deep learning object detection models, while limiting or eventually excluding real-world images from training, to achieve comparable performance to those models trained on exclusively real-world images.

Background: In the international nuclear safeguards community, deep learning computer vision models are being researched for several applications, such as supporting review of surveillance imagery, incoming open source images, and existing internal archives. Open source and commercially available deep learning computer vision architectures include common objects such as animals, heavy equipment, or sporting goods, require fine-tuning prior to their use on niche domains such as international safeguards. We confirmed this hypothesis by testing three open source models on several hundred open source safeguards-relevant images in 14 classes, see results in Table 1 and Fig 1.

Table 1. Platform Performance by Nuclear-Relevance Classifications

	Platform A	Platform B	Platform C
Precision	1.00	0.81	0.62
Recall	0.10	0.31	0.29
True Positives	11	40	31
True Negatives	141	112	121
False Positives	0	9	19
False Negatives	96	87	77
TOTAL	248	248	248

Fig. 1: Example classifications from three common image classification platforms.

Expected classification: Monarch butterfly

Actual classifications:* Insect, Invertebrate, Viceroy butterfly, Monarch, Pollinator, Honeybee

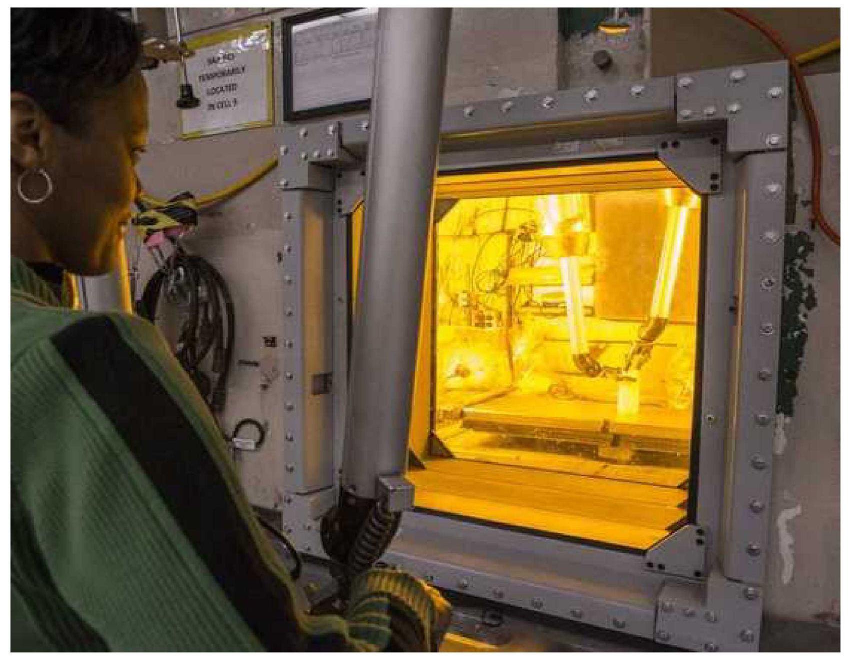

Expected classification: Hot cell

Actual classifications:* Building, Machine, Equipment, Indoors, Scientific Instrument, Apparel

** In an FY19 project funded by the National Nuclear Security Administration's Office for International Safeguards, and in collaboration with Lawrence Livermore National Laboratory, we tested a series of images across three widely known, free online platforms without any additional training. These are example classifications from that work.*

Fine-tuning these models so that they contain domain-relevant expertise requires thousands of labeled training images per class, with a potential for hundreds to thousands of domain-specific classes. In the field of international safeguards, the challenges related to collecting and labeling relevant images is compounded by a lack of availability of images, historical bias of available images, and the subject matter expertise to accurately label some images. We are developing methods to limit the number of real-word images needed to train object detection models, instead relying on 2D images that we generate using 3D computer graphics.

Use Case: Our test case is remote manipulator arms. Remote manipulators allow a user to remotely handle radioactive or other hazardous material, such as that in a hot cell that can be used for the separation of plutonium from irradiated nuclear fuel.

Fig. 2: Example remote manipulator arm configurations

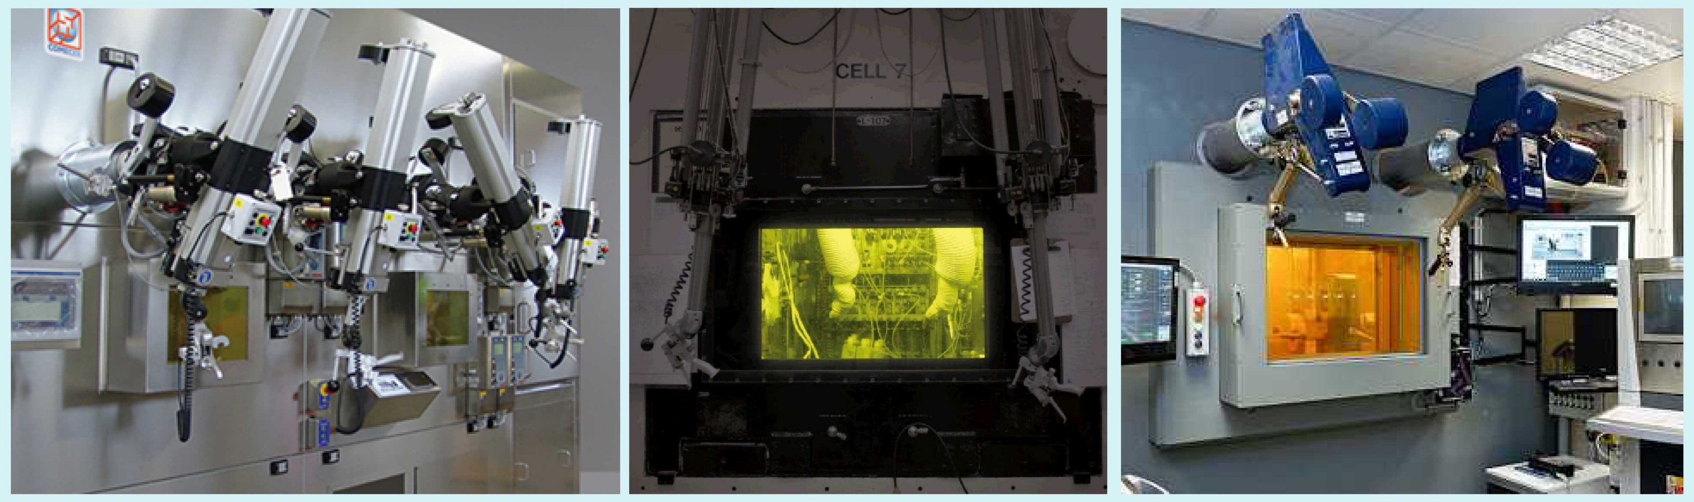

Approach: We are developing 3D models of our objects of interest using Houdini, a technical artists' tool widely used in the film and game industries. Once we develop initial technical specifications of our object in a parameterized model, we can randomly generate physics-informed instances of an object, as well as modify material, age, and condition. We can rotate the object within its environment, as well as changing lighting conditions, perspective, background, magnification, and photographic techniques such as lens or filter.

We are using these images to fine-tune state-of-the-art object detection models. We will use machine learning explanatory techniques to inform the development of our 3D models and improve our synthetic-training outcomes.

1) Fine-tune a model with synthetic images

2) Test performance on real-world images

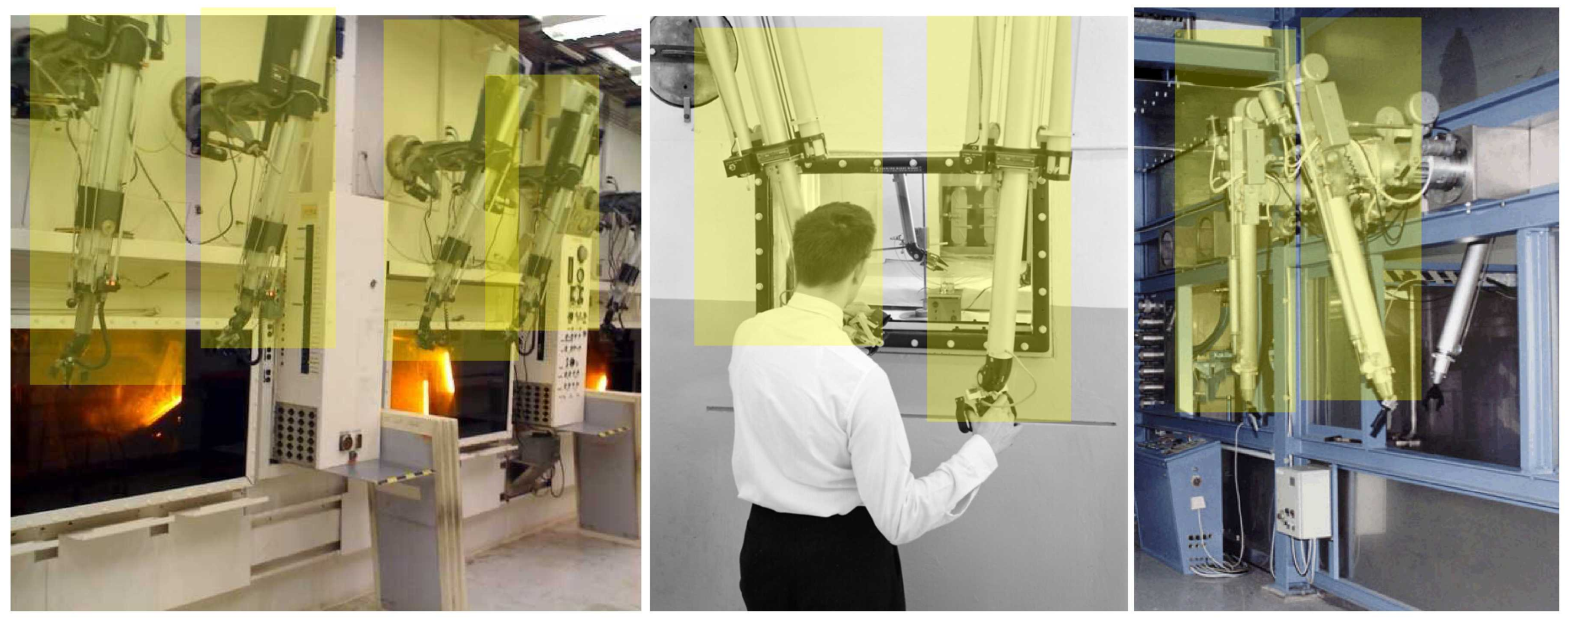

*3) Compare performance** to model fine-tuned on real-world images*

	REAL		SYNTHETIC
AP	01010010	<	01011000
AP50	01100100	<	01101001
AP75	01011011	<	01011110

***For example purposes only. These characters do not reflect actual model performance.*

Acknowledgements: This work was funded by Sandia National Laboratories' Laboratory Directed Research and Development Program.

Photo credits: Fig 1 (l-r) National Geographic, Savannah River National Laboratory; Fig 2 (l-r): Comecer, Oak Ridge National Laboratory, Research Centre Řež; Approach 2) (l-r, bounding boxes added) Argonne National Laboratory, NASA, Wälischmiller Engineering GmbH

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology &Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND No. ____________