Robust Multi-task Learning and Online Refinement for Spacecraft Pose Estimation across Domain Gap

Tae Ha Park^{*} and Simone D'Amico[†]

This work presents Spacecraft Pose Network v2 (SPNv2), a Convolutional Neural Network (CNN) for pose estimation of noncooperative spacecraft across domain gap. SPNv2 is a multi-scale, multi-task CNN which consists of a shared multi-scale feature encoder and multiple prediction heads that perform different tasks on a shared feature output. These tasks are all related to detection and pose estimation of a target spacecraft from an image, such as prediction of pre-defined satellite keypoints, direct pose regression, and binary segmentation of the satellite foreground. It is shown that by jointly training on different yet related tasks with extensive data augmentations on synthetic images only, the shared encoder learns features that are common across image domains that have fundamentally different visual characteristics compared to synthetic images. This work also introduces Online Domain Refinement (ODR) which refines the parameters of the normalization layers of SPNv2 on the target domain images online at deployment. Specifically, ODR performs self-supervised entropy minimization of the predicted satellite foreground, thereby improving the CNN's performance on the target domain images without their pose labels and with minimal computational efforts. The GitHub repository for SPNv2 is available at https://github.com/tpark94/spnv2.

keywords: Vision-only navigation, Rendezvous, Pose estimation, Computer vision, Deep learning,

1. Introduction

Monocular-based pose estimation of a noncooperative spacecraft has been a topic of research interest in recent years due to its applicability to various future mission concepts that address sustainability of the near-Earth space, such as refueling space assets³⁷ and active debris removal.¹² In such scenarios, an autonomous servicer spacecraft could generate safe and fuel-efficient rendezvous and docking trajectories based on real-time estimated pose of the target relative to the servicer. The pose information can be extracted from a sequence of images captured by the monocular camera, a low Size-Weight-Power-Cost (SWaP-C) sensor that is particularly suitable to the limited on-board capacity of small satellites such as CubeSats.

In recent years, Machine Learning (ML) techniques based on Convolutional Neural Networks (CNN) have been extensively studied for the application of spaceborne computer vision and especially pose estimation,^{2, 6, 7, 33} especially with the advent of the Spacecraft PosE Estimation Dataset (SPEED)^{39, 41} and the international Satellite Pose Estimation Competition (SPEC2019)¹⁸ co-organized by the Advanced Concepts Team (ACT) of the European Space Agency (ESA) and the Space Rendezvous Laboratory (SLAB) at Stanford University. SPEED consists of 15,000 synthetic images of the Tango spacecraft from the PRISMA mission^{9,10} and 300 Hardware-In-the-Loop (HIL) images captured from the Testbed for Rendezvous and Optical Navigation (TRON) facility at SLAB.³⁰ It allows the community to benchmark pose estimation performance and compare different methods on a unified set of test data with consistent metrics. However, spaceborne images have fundamentally different visual properties compared to computergenerated synthetic images, so CNN performance on synthetic images does not translate to the equal level of performance on spaceborne images encountered during a mission. Known as domain gap in literature,^{1,35} addressing the performance gap between training and testing data from different distributions is an active field of research in a wide range of deep learning applications. However, addressing domain gap in spaceborne applications is extremely difficult compared to terrestrial applications such as autonomous driving, since access to space is prohibitively expensive for data collection and road tests. While SPEED's HIL images are designed as surrogates of spaceborne images, they lack in quantity and diversity in pose labels and illumination conditions to comprehensively evaluate the generalization capability of the CNN models across different image domains.

To address this issue, $SPEED+^{31,32}$ has been recently introduced with two distinct HIL image domains, lightbox and sunlamp, that are captured from the upgraded TRON facility at SLAB.³⁰ The

^{*}Ph.D. Candidate, Department of Aeronautics & Astronautics, Stanford University, Stanford, CA 94305, USA, tpark94@stanford.edu

[†]Associate Professor, Department of Aeronautics & Astronautics, Stanford University, Stanford, CA 94305, USA, damicos@stanford.edu

lightbox domain consists of 6,740 images of the halfscale mockup model of the Tango spacecraft illuminated with albedo lightboxes to simulate diffuse light in the Earth's orbit, whereas the sunlamp domain contains 2,791 images of the same model illuminated with a metal halide arc lamp to simulate direct sunlight. With augmented realism of the visual characteristics of the HIL images, SPEED+ allows for a comprehensive analysis of the robustness of the CNN models trained exclusively on synthetic images. The SPEED+ dataset was used in the second international Satellite Pose Estimation Competition (SPEC2021)[‡] with emphasis on bridging the performance gap between the synthetic training and HIL test images.

This paper introduces Spacecraft Pose Network v2 (SPNv2), the SLAB's solution to resolving the domain gap in SPEED+ and beyond. Similar to its prede- ${\rm cessor}^{39,\,40}$ which consists of the shared ${\rm AlexNet}^{20}$ backbone and multiple heads for bounding box detection, attitude classification and residual regression, SPNv2 is a multi-scale, multi-task pose estimation CNN with a shared feature encoder and prediction heads that perform various tasks such as bounding box detection, pose regression, heatmap prediction and spacecraft foreground segmentation. These tasks are different yet rely on common information such as the target's geometry and pose. Therefore, jointly training SPNv2 on those tasks forces the shared encoder to learn such common information, which also translates across different image domains, instead of any task-specific features that may also be particular to the synthetic image domain.

Training SPNv2 consists of two stages:

- 1. offline robust training, which trains exclusively on the SPEED+ synthetic training set with extensive data augmentation to make CNN as domain-invariant as possible, and
- 2. **online domain refinement**, which refines the feature encoder of SPNv2 on unlabeled test images by minimizing the Shannon entropy³⁸ on the segmentation task.

The motivation of the Online Domain Refinement (ODR) stage is that, despite the realism of SPEED+ HIL images, there still remains an inherent gap between the lightbox, sunlamp and spaceborne domains. Naturally, while good performances on SPEED+ HIL domains suggest good generalization capability of the CNN models, it does not necessarily translate to an equal level of performance on the spaceborne images. Therefore, the second stage aims to directly incorporate the test or target domain images into the

training procedure, so that the CNN model parameters are influenced by the exact images on which it is designed to function. While this paper performs ODR on SPEED+ test domains, the same strategy can be applied to spaceborne images acquired during missions. Note that ODR is *source-free*, i.e., it does not require simultaneous access to the large-scale training data for adaptation, which complies with the memory and computational constraints of the on-board avionics.

It is shown that the offline training of SPNv2 alone provides remarkable performance on SPEED+ HIL domains compared to the baseline studies,³² with translation errors around 22cm / 23cm and orientation errors about 8.0° / 10.4° for lightbox and sunlamp, respectively, for a larger, batch-agnostic variant of SPNv2. ODR further refines the translation error by another 6 cm / 7 cm and orientation error by 2.4° / 0.6° on lightbox and sunlamp, respectively, after observing 4096 unlabeled images. Extensive analyses are performed on both offline training and ODR to articulate the factors contributing to this success.

This paper is organized as follows. Section 2 describes related state-of-the-art for different components of the proposed methods. Then, Sections 3 and 4 explain the proposed SPNv2 architecture, offline training procedure, and ODR. Section 5 examines the experimental results of both offline training and ODR, and Section 6 discusses limitations and future directions. The paper ends with concluding remarks in Section 7.

2. Related Work

Spacecraft Pose Estimation The first ML-based approach to pose estimation of a known target spacecraft is Spacecraft Pose Network (SPN),^{39,40} which performs relative attitude determination via a hybrid approach of attitude classification and regression and performs translation estimation by exploiting the perspective transformation and geometric constraints. Afterwards, many top-performing entries of SPEC2019 have shown diverse CNN architectures and pose estimation strategies, such as probablistic orientation estimation via soft classification³⁶ or estimation of a set of designated keypoints on the spacecraft surface via direction regression³³ or heatmaps.⁷ Specifically, the estimated 2D keypoints can then be used in Perspective-*n*-Point $(PnP)^{22}$ to recover full 6D pose. Many CNN models that followed the competition also adopt similar strategies as well.^{2,6} Due to the limitation of SPEED and other proposed datasets for spacecraft pose estimation, these CNN models are designed to drive performance on synthetic images that

[‡]https://kelvins.esa.int/pose-estimation-2021/



Fig. 1: The pose estimation CNN architecture based on the EfficientNet⁴⁵ backbone, the Bi-directional Feature Pyramid Network (BiFPN), and multi-task head networks shared by the features at all scales.

have vastly different visual characteristics compared to spaceborne images.

Domain Gap in Space Domain $gap^{1,13,35}$ is a ubiquitous problem for any real-life applications of ML models. However, addressing domain gap in spaceborne vision applications is extremely challenging due to the inaccessibility of space for even a small-scale data collection. This barrier prevents a comprehensive evaluation of the trained model on the spaceborne images on which the model should demonstrate robust performance *prior* to deployment. Therefore, existing works have performed evaluation on a handful of spaceborne images from the previous missions^{2, 33, 36, 40} or captured in a high-fidelity simulation environment with a physical model of the target satellite.^{18,32,34} In these cases, the CNN models are trained on abundant synthetic images with extensive data augmentation or domain randomization 4^{47} to render them as domain-agnostic as possible. However, their evaluation on spaceborne or other physical image domains lacks comprehensiveness due to the limited sample quantity, and they often lack in-depth analyses on the factors that attribute to successfully bridging the domain gap.

Source-Free Domain Adaptation Another consequence of inaccessibility of space is that the solution precludes many conventional domain adaptation algorithms^{13, 43, 48} that assume simultaneous availability of both labeled source and unlabeled target domain data. In reality, the spaceborne images are only available during missions in space, thus performing domain adaptation on-board the spacecraft with loaded source data is simply not feasible given the computational and memory limitations of satellite avionics. Existing source-free domain adaptation methods would leverage generative models during offline training,^{21,23} pseudo-labeling and information maximization.²⁵ entropy minimization,⁴⁹ and so on. Test-Time Training $(TTT)^{27,44}$ jointly optimizes the main task (e.g., classification) and a Self-Supervised Learning (SSL) task from a shared feature encoder. Then, at test time, the encoder parameters are updated via self-supervised learning on the target domain samples. While it is a new paradigm to source-free domain adaptation, TTT requires an additional head for SSL that is often hand-designed by users (e.g., rotation prediction,¹⁵ jigsaw puzzle²⁹) or requires a large batch of negative sample pairs as is the case for contrastive learning.⁸ Instead, this work minimizes entropy of the predicted satellite foreground and localizes the gradient update to the affine transformation parameters of the batch normalization layers,^{24,49} which does not require an additional self-supervised proxy task and minimizes the computation incurred during backpropagation.

3. Offline Robust Training

This section describes the offline robust training of SPNv2 on the synthetic dataset, which is achieved by combination of a multi-scale, multi-task CNN architecture design, extensive data augmentation and domain randomization.

3.1 Multi-Scale, Multi-Task Architecture

The main SPNv2 architecture visualized in Figure 1 closely follows EfficientPose³ based on the Efficient-Det⁴⁶ feature encoder, which comprises the Efficient-Net⁴⁵ backbone and Bi-directional Feature Pyramid Network (BiFPN) to fuse features from different scales. Then, the output of the shared feature encoder is fed into different prediction heads. Following the original design, the EfficientPose $(h_{\rm E})$ head consists of the subnets for binary classification of the object presence, bounding box prediction, target rotation and translation regression. In this work, two additional prediction heads are added: the Heatmap $(h_{\rm H})$ head which predicts the 2D heatmaps associated with each pre-designated keypoints on the spacecraft surface, and the Segmentation $(h_{\rm S})$ head for pixel-wise binary segmentation of the spacecraft foreground. The motivation for the proposed architecture is that the multi-task learning⁵ of related tasks could improve the generalization capability by suppressing the shared encoder from learning any task-specific features. By jointly training on related yet different tasks, the encoder would instead learn features that are common to those tasks, such as the spacecraft geometry and pose. Note that these features are also common across different image domains; therefore, learning them would also improve generalization from synthetic training to spaceborne test images. Henceforth, SPNv2 with $h_{\rm H}$, $h_{\rm E}$, and $h_{\rm S}$ is referred to be in full configuration.

Similar to EfficientPose,³ the convolutional parameters in the prediction heads are shared across different scales. The layouts of both the BiFPN feature fusion network and individual heads depend on the EfficientNet scaling parameter ϕ .⁴⁵ The rotation and translation networks in $h_{\rm E}$ also include the iterative refinement subnets whose layouts depend on ϕ as designed in the original work.³ The inputs to $h_{\rm E}$ are the level 3-7 feature outputs from BiFPN, where level n feature has the resolution downscaled from the input image by a factor of 2^n . The subnets of $h_{\rm H}$ and $h_{\rm S}$ mimic the structure of the bounding box subnet except the output dimension. However, in contrast to $h_{\rm E}$, these heads instead use high-resolution level 2 feature output for predictions, following the practices of the original EfficientDet⁴⁶ for semantic segmentation. Therefore, BiFPN is extended to always fuse



Fig. 2: Visualization of a 4D feature map tensor across N batches, C channels, $H \times W$ spatial resolution in BN (*left*) and GN (*right*) layers. The pixels in blue are normalized by the same mean and variance. Diagram from Wu & He.⁵⁰

the level 2-7 features, and the intermediate layers of these heads have the dimension $\min\{2W_{\mathrm{bifpn}}(\phi), 256\}$, where W_{bifpn} is the width of the BiFPN outputs at each level. The readers are referred to Tan et al.⁴⁶ and Bukschat & Vetter³ for details on different feature levels and the scaling parameter ϕ .

3.2 Normalization Layers

The original implementation of EfficientPose³ utilizes Group Normalization (GN) layers⁵⁰ throughout the entire network, which is a batch-agnostic alternative to the commonly employed Batch Normalization (BN) layers.¹⁶ As shown in Figure 2, BN normalizes the features across N images with the approximate mean and variance of the features across all images in the image domain, which are estimated using running average of the batch-wise feature statistics during training. On the other hand, GN normalizes the feature across a group of G channels per each image using the statistics of the incoming features. Since the normalization procedure does not stretch across the batch of images, GN is a completely batch-agnostic layer. Wu & He have demonstrated that a CNN with GN layers can show equivalent or better performance on a number of benchmark computer vision tasks than a CNN with BN layer.⁵⁰ The independence from batch-wise information means a CNN with GN layers can be trained on single images instead of batches in a training environment with limited computational capabilities. Moreover, as explained in Section 4 and shown in Section 5 later, the batch-agnostic nature of GN layers favors the online refinement procedure against the new image domain.

In this work, the prediction heads of SPNv2 use GN layers with 16 channels per group following the implementation of EfficientPose.³ However, for the EfficientNet backbone and the BiFPN layers, both



Table 1: List of data augmentations and equivalent commands in Albumentations.⁴ Each augmentations are activated with 50% probability.

| | Augmentation | Com | | |
|--------------------|------------------------|--------------------------------|----------------------|----------------|
| Ī | Brightness & Contrast | RandomBrigh | _ | |
| | Random $Erase^{52}$ | Coars | eDropout | |
| | Sun Flare | Random | SunFlare | |
| | Blur | <pre>OneOf(MotionBlur, M</pre> | ledianBlur, GlassBlu | r) |
| | Noise | ${\tt OneOf}({\tt GaussN})$ | loise, ISONoise) | |
| | | tur I | | |
| original | RandomBrightnessContra | st MedianBlur | MotionBlur | GlassBlur |
| style augmentation | GaussNoise | ISONoise | CoarseDropout | RandomSunFlare |

Fig. 3: Example image, its stylized version, and visualization of different augmentation techniques.

BN layers and GN layers with the group size 8 are implemented and evaluated for its contribution to resolving the domain gap.

3.3 Training Losses

Training $h_{\rm E}$ partially follows the original implementations of EfficientPose³ with a total of 9 anchor boxes of aspect ratios {0.5, 1, 2} at scales {2^{1/3}, 2^{2/3}, 2}. The classification subnet employs the focal loss²⁶ with $\alpha = 0.25$ and $\gamma = 2.0$. Unlike EfficientPose, the bounding box subnet training minimizes the complete Intersection-over-Union (IoU) loss,⁵¹ and the rotation and translation prediction subnets jointly minimize the SPEED score ($E_{\rm pose}$)¹⁸ defined as

$$E_{\text{pose}} = E_{\text{R}}(\boldsymbol{R}, \boldsymbol{R}) + E_{\text{T}}(\tilde{\boldsymbol{t}}, \boldsymbol{t}) / \|\boldsymbol{t}\|, \qquad [1]$$

where, given the predicted and ground-truth rotation matrices and translation vectors $(\tilde{\boldsymbol{R}}, \tilde{\boldsymbol{t}})$ and $(\boldsymbol{R}, \boldsymbol{t})$, the rotation error $E_{\rm R}$ and the translation error $E_{\rm T}$ are respectively defined as

$$E_{\rm R}(\tilde{\boldsymbol{R}}, \boldsymbol{R}) = \arccos \frac{\operatorname{tr}(\boldsymbol{R}^{\top} \tilde{\boldsymbol{R}}) - 1}{2},$$
 [2]

$$E_{\mathrm{T}}(\boldsymbol{t},\boldsymbol{t}) = \|\boldsymbol{t} - \boldsymbol{t}\|.$$
 [3]

The SPEED loss is then the sum of the rotation error in radians and translation error normalized by the ground-truth distance. Since SPEED loss is the official performance metric of the Satellite Pose Estimation Competition (SPEC), it is a natural choice of loss for the pose regression tasks. Note that while the EfficientPose implementation regresses the axisangle representation of the orientation, the rotation subnet in this work instead regresses the 6D representation studied by Zhou et al.,⁵³ whose continuous property has shown to outperform other parametrization methods in rotation regression tasks. With the 6D representation as a pair of 3D vectors $(\mathbf{r}_1, \mathbf{r}_2)$, the rotation matrix $\mathbf{R} = [\mathbf{R}_1 \mid \mathbf{R}_2 \mid \mathbf{R}_3]$ can be recovered column-wise via

$$\begin{cases} \boldsymbol{R}_1 &= N(\boldsymbol{r}_1) \\ \boldsymbol{R}_2 &= N(\boldsymbol{r}_2 - (\boldsymbol{r}_1 \cdot \boldsymbol{r}_2)\boldsymbol{r}_1) \\ \boldsymbol{R}_3 &= \boldsymbol{R}_1 \times \boldsymbol{R}_2 \end{cases}$$
[4]

where $N(\cdot)$ is a normalization operator.

Finally, $h_{\rm H}$ minimizes the pixel-wise mean squared error loss against the ground-truth heatmaps centered around each visible keypoints. $h_{\rm S}$ minimizes the pixelwise binary cross entropy loss. All losses are optimized simultaneously with equal weights.



Fig. 4: Visualization of complete workflow of offline training, pose inference, and ODR. The lock diagram indicates that the parameters of the corresponding component are fixed.

3.4 Data Augmentation

SPNv2 is trained with extensive data augmenta $tion^{42}$ to mitigate the overfitting to the synthetic images. The augmentations are implemented with the Albumentations library,⁴ and the list of employed augmentations and their equivalent Albumentations commands are provided in Table 1. Each of the augmentations are activated with probability 0.5. Note that when applying blur, either motion blur, median blur, or Gaussian blur is applied with equal probability, and likewise for applying noise as well. Moreover, random $erase^{52}$ and sun flare augentations, which simulate the occlusion of satellite parts due to extreme shadowing or the sun lamp's direct sunlight, are implemented after modification to localize the effect within the bounding box of the target satellite instead of the entire image frame.

In addition to the standard data augmentation techniques, style augmentation¹⁷ is employed to randomize the style and inherent texture of the synthetic spacecraft via neural style transfer,¹⁴ as done in the SPEED+ baseline studies.³² An example training image, its style-augmented version, and employed augmentation techniques are visualized in Figure 3.

4. Online Domain Refinement

While the offline robust training on synthetic images could improve the generalizability across different spacecraft image domains, its performance during real missions lacks fundamental guarantee as the spaceborne images have never been integrated into the offline training procedure. Therefore, as the spaceborne images of the target become available during the rendezvous phase, SPNv2 can be refined on the target domain images on-board the spacecraft avionics. Considering the on-board hardware limitations, there exists a number of constraints that must be met for a fully mission-compliant ODR method; specifically, it must be

- 1. *source-free*, i.e., the satellite only has access to a model trained offline on synthetic data;
- 2. *sequential*, i.e., it must not wait to collect a large batch of samples to perform refinement while withholding predictions;
- 3. *computationally efficient*, i.e., any learning components must be minimized or restricted to a small subset of the entire CNN.

This paper adopts and modifies TENT^{49} to minimize Shannon entropy³⁸ on the pixel-wise binary classification of the segmentation head. Specifically, the

gradient updates are restricted to the affine transformation parameters of the normalization layers of the shared feature encoder of SPNv2. If the encoder contains the BN layers, the ODR must re-estimate the batch-wise feature normalization statistics as well. Unlike TENT which normalizes the features using the input batch statistics, the proposed method continuously updates the running statistics of the BN layers after every B images in a manner similar to AdaBN.²⁴ If the encoder instead consists of the GN layers, then the process of estimating the batch-wise statistics can be omitted altogether. The proposed ODR is also similar to $TTT^{27,44}$ in that both works are essentially multi-task learning, and the network parameters are updated only in the shared feature encoder. However, unlike TTT which updates the entire encoder parameters, ODR only updates the affine transformation parameters of the normalization layers that are only a small fraction of the entire set of learnable parameters. Moreover, TTT introduces additional layers for the unsupervised task which must be trained from scratch and could take a while for the training to stabilize, whereas ODR simply refines the already trained parameters.

The complete workflow of ODR is visualized in Figure 4 along with the offline training and pose inference stages for comparison. The offline training updates the parameters of the entire network, including the shared feature encoder (f) and the prediction heads $(h_{\rm E}, h_{\rm H}, h_{\rm S})$, by minimizing the sum of losses from all heads, i.e., $\ell_E + \ell_H + \ell_S$. Then, at inference, only the outputs of $h_{\rm H}$ and $h_{\rm E}$ are used to predict the pose. Meanwhile, ODR does not involve $h_{\rm H}$ and $h_{\rm E}$, but instead minimizes the entropy loss $\ell_{\rm ent}$ only through $h_{\rm S}$ on the target domain images. During ODR, the parameteres of the prediction heads are locked, and only the affine transformation parameters of the BN or GN layers in f are updated.

4.1 Entropy Minimization

ODR minimizes Shannon entropy on h_S . In classification tasks, entropy minimization encourages the predicted probability distribution to have distinct peaks, effectively enhancing the confidence in predicted classes. Mathematically, let θ_f denote the set of all learnable parameters in the feature encoder, and θ_S denote the same for the segmentation head. Let $\theta_f^{\text{norm}} \subset \theta_f$ denote the set of parameters associated with the normalization layers of the feature encoder. Then, the refinement procedure amounts to solving

$$\min_{\boldsymbol{\theta}_{f}^{\text{norm}}} \frac{1}{n} \sum_{i=1}^{n} \ell_{\text{ent}}(\boldsymbol{x}_{i}; \boldsymbol{\theta}_{f}, \boldsymbol{\theta}_{\text{S}}),$$
 [5]

where $\{x_i\}$ are the unlabeled target domain images, and the entropy loss is given as

$$\ell_{\rm ent}(\boldsymbol{x}_i;\boldsymbol{\theta}_f,\boldsymbol{\theta}_{\rm S}) = -\sum_p \sigma(\hat{\boldsymbol{y}}_{i,p}) \log \sigma(\hat{\boldsymbol{y}}_{i,p}), \quad [6]$$

where $\hat{y}_i = h_{\rm S}(f(x_i))$ is the output logit of the predicted segmentation map, $\hat{y}_{i,p}$ is the *p*-th pixel of \hat{y}_i , and $\sigma(\cdot)$ is a pixel-wise sigmoid function.

4.2 Normalization Layer Parameter Updates

As explained, only the affine transformation parameters of the normalization layers of the feature encoder are updated through backpropagation during ODR. Formally, let $(\boldsymbol{z}_{j}^{\text{in}}, \boldsymbol{z}_{j}^{\text{out}})$ denote the input and output features of the *j*-th normalization layer in the feature encoder *f*. A normalization layer modulates and updates the feature according to

$$\boldsymbol{z}_{j,c}^{\text{out}} = \text{Norm}(\boldsymbol{z}_{j,c}^{\text{in}}) = \gamma_{j,c} \left(\frac{\boldsymbol{z}_{j,c}^{\text{in}} - \mu}{\sigma}\right) + \beta_{j,c}, \quad [7]$$

where $\boldsymbol{z}_{j,c}^{\text{in}}$ denotes the input feature vector at *c*-th channel, (μ, σ^2) are the feature mean and variance applied across the batch of size *N* for BN or across some *G* channels for GN, and $(\gamma_{j,c}, \beta_{j,c})$ are the affine transformation parameters of the *j*-th normalization layer associated with *c*-th channel of the input feature. Then, the set of refinement parameters, $\boldsymbol{\theta}_{f}^{\text{norm}}$, consists of the affine parameters (γ, β) of all normalization layers in *f*, i.e., $\boldsymbol{\theta}_{f}^{\text{norm}} = \{\gamma_{j,c}, \beta_{j,c}\}$.

In this work, if the BN layers are used to construct the encoder, then the feature mean and variance at each BN layer, which should approximate those of the features across the entire target domain, are updated online every B images as running averages according to

$$\mu_{(i+1)} = m\mu_{(i)} + (1-m)\mu_B,$$

$$\sigma_{(i+1)}^2 = m\sigma_{(i)}^2 + (1-m)\sigma_B^2,$$
[8]

where $m \in [0, 1]$ is a momentum hyperparameter, $(\mu_{(i)}, \sigma_{(i)})$ are the current running averages of feature statistics, and (μ_B, σ_B^2) are the new observed feature statistics over the most recent *B* images. The key difference against the conventional BN is that the samples are not processed in batches due to memory limits; they are instead processed sequentially, and the batch mean and variance are updated online¹⁹ as done for AdaBN.²⁴

5. Experiments

This section describes the experimental procedures and results for both offline robust training and online domain refinement.



Fig. 5: Example images from different domains of SPEED+. Figure from Park et al.³²

5.1 Dataset

SPNv2 is trained and evaluated on SPEED $+^{31,32}$ which comprises data from three distinct domains: synthetic, lightbox and sunlamp. The synthetic domain consists of 59,960 computer-generated images of the Tango spacecraft from the PRISMA mission.¹⁰ On the other hand, lightbox and sunlamp respectively contain 6,740 and 2,791 Hardware-In-the-Loop (HIL) images of the mockup model of the same spacecraft captured in the high-fidelity robotic simulation environment of the Testbed for Rendezvous and Optical Navigation (TRON) facility at the Space Rendezvous Laboratory (SLAB) of Stanford University. Specifically, the lightbox images are simulated with a set of calibrated lightboxes emulating the diffuse light in Earth's orbits, whereas the sunlamp images are illuminated with a metal halide arc lamp to mimic the direct sunlight commonly encountered in space. Figure 5 provides visualization of a few samples from each domain. For more information on SPEED+, the readers are referred to Park et al.³² In addition to SPEED+, the binary masks for the synthetic images are created in order to train SPNv2 for the segmentation task. SPNv2 is also evaluated on prisma25 which contains 25 labeled spaceborne images of Tango acquired during the rendezvous phase of the PRISMA mission.9,10

In this work, the synthetic domain's training and validation sets are used exclusively for offline robust training of SPNv2. Then, each lightbox and sunlamp unlabeled images are used for ODR. This practice resembles the real-life mission constraint, where images from the target domain are used only when that domain becomes operationally available. The prisma25 images are not used for ODR and reserved only for evaluation of offline training since its quantity is severely limited for ODR.

5.2 Metrics

For bounding boxes predicted by $h_{\rm E}$, the standard IoU metric is used. For pose solutions regressed by $h_{\rm E}$ and computed via Perspective-*n*-Point (PnP)²²

from the keypoints predicted from $h_{\rm H}$, rotation error $(E_{\rm R})$ in degrees (Eq. 2), translation error $(E_{\rm T})$ in meters (Eq. 3), and SPEED score (Eq. 1) are reported. Note that when evaluated on SPEED+ lightbox and sunlamp samples, the modified SPEED score $(E_{\rm pose}^*)^{32}$ is used to zero out the errors smaller than the thresholds based on the TRON calibration, i.e., for individual sample,

$$E_{\text{pose}}^* = \begin{cases} 0 & \text{if } E_{\text{R}} < 0.169^{\circ} \text{ and} \\ & E_{\text{T}} / \|\boldsymbol{t}\| < 2.173 \text{mm/m}, \quad [9] \\ E_{\text{pose}} & \text{otherwise} \end{cases}$$

where $\|\boldsymbol{t}\|$ is the ground-truth distance to the target.

5.3 Implementation Details

In all offline experiments on the synthetic domain, SPNv2 is trained with the $AdamW^{28}$ optimizer with the initial learning rate of 1×10^{-3} which decays by the factor of 0.1 after 15 and 18 epochs. The training lasts 20 epochs (i.e., trains on the entire training set 20 times) for all but a configuration with only $h_{\rm H}$, which is trained for only 10 epochs since heatmap prediction is easier to learn. The input image is resized to 768 \times 512, which minimizes the loss of information due to excessive downscaling and has an approximately matching aspect ratio compared to the original 1920 \times 1200 images. Note that unlike existing methods for spacecraft pose estimation,^{7,33,39} no separate object detection network is used to detect and crop out the region-of-interest around the spacecraft in advance before feeding it to the pose estimation network. The reason is that, unlike SPEED whose farthest image is nearly 40 m away, SPEED+ data are limited to maximum 10 m separation, so the shape of the farthest spacecraft can still be recognized at high enough image resolution.

Unless noted otherwise, the EfficientNet scaling parameter is set to $\phi = 3$, which amounts to roughly 12M parameters for the EfficientNet-B3 backbone and BiFPN layers, and BN layers are used to construct the feature encoder. All training is run on 4 NVIDIA



| Table 2: Bounding box and pose predi | ctions of SPNv2 different | head configurations. | The performances from |
|--|---------------------------|----------------------|--------------------------|
| $h_{\rm H}$ (H) and $h_{\rm E}$ (E) on the SPEED | 0+ lightbox and sunlam | p domains are report | ted wherever applicable. |
| Bold numbers indicate best perform | mances. | | |

| Heads | Source | | ligł | ntbox | sun | ılamp | | | |
|---|--------|------------|-----------------------|-----------------------|-------------------------|------------|-----------------------|-------------------------|-------------------------|
| | Source | IoU [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] | IoU [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] |
| Е | Е | 0.878 | 0.400 | 21.921 | 0.446 | 0.901 | 0.293 | 32.066 | 0.608 |
| H | Η | - | 0.368 | 13.577 | 0.298 | - | 0.424 | 19.739 | 0.414 |
| E + H | E H | 0.915 | $0.181 \\ 0.272$ | $8.757 \\ 6.924$ | $0.183 \\ 0.165$ | 0.911 | $0.251 \\ 0.314$ | $14.175 \\ 12.323$ | $0.288 \\ 0.268$ |
| $\overline{\mathbf{E} + \mathbf{H} + \mathbf{S}}$ | E H | 0.918 - | 0.175 0.271 | 8.004 6.479 | 0.169 0.158 | 0.919 - | 0.225 0.307 | 12.433 11.065 | 0.254 0.245 |

Table 3: Bounding box and pose predictions of full configuration SPNv2 and different data augmentation configurations. Starting with the baseline set of augmentations including Brightness & Contrast, Blur and Noise, different augmentation techniques are progressively added to the training. Bold numbers indicate best performances.

| Config | Source | | ligh | ntbox | | sunlamp | | | | |
|----------------|--------|---------|-----------------|-----------------|-------------------------|---------|-----------------|-----------------|-------------------------|--|
| e onng. | bource | IoU [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] | IoU [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] | |
| Deceline | Е | 0.853 | 0.518 | 24.678 | 0.509 | 0.867 | 0.641 | 47.893 | 0.937 | |
| Basenne | Η | - | 0.506 | 21.994 | 0.465 | - | 0.735 | 47.546 | 0.955 | |
| + Random Erase | Е | 0.811 | 0.756 | 24.168 | 0.534 | 0.510 | 2.766 | 79.232 | 1.771 | |
| | H | - | 0.665 | 22.544 | 0.494 | - | 2.295 | 80.778 | 1.744 | |
| Sup Flore | Ε | 0.892 | 0.314 | 11.670 | 0.252 | 0.825 | 0.875 | 33.239 | 0.709 | |
| + Sun Plate | Η | - | 0.347 | 10.018 | 0.230 | - | 0.722 | 31.504 | 0.661 | |
| + Style Aug. | Е | 0.918 | 0.175 | 8.004 | 0.169 | 0.919 | 0.225 | 12.433 | 0.254 | |
| | Η | - | 0.271 | 6.479 | 0.158 | - | 0.307 | 11.065 | 0.245 | |

V100 GPUs with per-GPU batch size 4 except for $\phi = 6$ which is trained with per-GPU batch size 1. When trained on multiple GPUs with BN layers, its statistics are synchronized across devices.

For ODR, the test images are provided one at a time with no data augmentation. The running mean and variance of each BN layers are updated after every B images with m = 0.9, and the experiment terminates after SPNv2 has seen N random samples from the test domain. ODR experiments are conducted on a single GPU.

5.4 Results: Offline Robust Training

First, SPNv2 is trained offline with a full set of data augmentations listed in Table 1 but with different configurations of prediction heads. Table 2 shows that for a full set of data augmentations applied, having more prediction heads in the architecture improves the CNN robustness on the SPEED+ HIL domains. The performance culminates with the full configuration visualized in Figure 1, which hints that jointly training a shared feature extractor for different tasks prevents it from learning task-specific features. On the other hand, Table 3 fixes the SPNv2 architeture to full configuration and instead studies the effects of three data augmentation techniques: Random Erase,⁵² Sun Flare, and Style Augmentation.¹⁷ When random erase is applied alone, the pose error actually increases, almost doubling for the sunlamp domain, suggesting the occlusion introduced by the random erase makes the learning much harder. However, as the other two augmentations are progressively added to the training, the pose error also decreases. Interestingly, adding sun flare augmentation significantly improves the pose errors on the lightbox domain compared to the baseline, whereas adding style augmentation does so on the sunlamp domain.

An interesting observation is that as the pose er-



Table 4: Comparison of SPNv2 models with different scaling parameter ϕ and normalization layers in the feature encoder. The number of parameters counts those of the feature encoder. Bold numbers indicate best performances. BN: batch normalization. GN: group normalization.

| ϕ | Norm. | n. synthetic | | | | lightbo | x | sunlamp | | |
|---------------------------------|-------|-----------------|-----------------|-----------------------|-----------------|-----------------|-------------------------|-----------------|-----------------|-------------------------|
| (Num. of Param.) | Layer | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose} [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] |
| $\phi = 0 \; (3.8 \mathrm{M})$ | BN | 0.075 | 1.149 | 0.033 | 0.355 | 12.108 | 0.268 | 0.349 | 18.734 | 0.385 |
| $\phi = 3 \; (12.0 \mathrm{M})$ | BN | 0.054 | 0.987 | 0.027 | 0.175 | 6.479 | 0.142 | 0.225 | 11.065 | 0.230 |
| $\phi = 3 \; (12.0 \mathrm{M})$ | GN | 0.056 | 1.224 | 0.031 | 0.247 | 12.919 | 0.267 | 0.342 | 28.041 | 0.545 |
| $\phi = 6 \ (52.5 \mathrm{M})$ | GN | 0.031 | 0.885 | 0.021 | 0.216 | 7.984 | 0.173 | 0.230 | 10.369 | 0.219 |

ror decreases, the translation error is generally lower when regressed from $h_{\rm E}$ than when computed via PnP from $h_{\rm H}$. This can be explained by the fact that any small deviation in the predicted keypoint locations can cause significant difference in translation, especially in depth along the camera boresight. This suggests that when deployed, predictions from both heads can be used jointly to ensure more stable predictions of the target's pose throughout the mission. Inspired from this observation, the pose predictions of the full configuration SPNv2 onward report translation error from $h_{\rm E}$ and orientation error from $h_{\rm H}$ unless otherwise noted.

Finally, the EfficientNet scaling parameter ϕ , which controls the number of parameters associated with the feature encoder, and the normalization layers (BN or GN) are varied to gauge their effect on the CNN's generalization capability across domain gaps. The SPNv2 is trained with same training parameters for different size and normalization layer configurations. The results shown in Table 4 indicate that, while there is no noticeable difference in performance on the synthetic validation set, the performances on the SPEED+ HIL domains differ considerably depending on the network size and the type of normalization layers used to build the feature encoder. Specifically, compared to the generalizability accomplished by the default architecture ($\phi = 3$, BN), reducing the network size to $\phi = 0$ or replacing the BN layers with GN layers significantly degrades the network's performance on the HIL test sets. In order to match the performance of the default architecture with a completely batchagnostic one, the network size must be increased to $\phi = 6$, which more than quadruples the size of the feature encoder. The trend observed in Table 4 is a direct opposite of the desired characteristics of CNN models for spaceborne computer vision applications elaborated in Section 4, which aim for a smaller, batchagnostic neural network architecture suitable for the on-board satellite avionics. The search for such an architecture is left as a future work.

5.5 Results: Online Domain Refinement

The ODR experiments in this section build upon the full configuration SPNv2 trained offline with the full set of augmentations, which is considered as the baseline. Based on the results in Table 4, two SPNv2 architectures are considered: small but batch-dependent ($\phi = 3$, BN) and large but batch-agnostic ($\phi = 6$, GN). Recall that ODR has two hyperparameters:

- *B*, the image batch size which controls the frequency of updating the running average of the BN layer statistics according to Eq. 8, and
- N, the total number of target domain images the neural network observes during ODR.

The hyperparameter B is irrelevant if the network only consists of the GN layers. In this section, the effect of the ODR batch size B is first tested on the batch-dependent architecture of SPNv2. The results are shown in Figure 6, where N = 1024 is fixed. It shows that there is a benefit of increasing the batch size to B = 4 or B = 8 for lightbox, whereas there is no benefit of updating the BN statistics after more than B = 1 image for sunlamp. From this observation, the batch size is fixed to B = 4, i.e., the BN layer means and variances are updated after every 4 images, as it seems to strike the balance between both HIL domains.

Then, the total number of observed samples (N) is varied for both batch-dependent and batch-agnostic architectures of SPNv2, as shown in Figure 7. As expected, as the unsupervised entropy minimization continues on, the SPEED score improves for both HIL domains. However, for the batch-dependent architecture, the improvement not only halts after N= 1024 images for lightbox and N = 2048 samples for sunlamp, but it also becomes worse than the baseline performance after nearly 8000 images for the lightbox domain. On the other hand, for the batchagnostic architecture, such failure is not observed for both HIL domains. In fact, the level of performance



Fig. 6: ODR on SPEED+ test domains with varying batch size (B). Total number of samples seen is fixed to N = 1024. Error bars represent standard deviations over 5 runs of ODR with different random seeds.



Fig. 7: ODR on SPEED+ test domains with varying number of observed samples (N). The batch size is fixed to B = 4. Error bars represent standard deviations over 5 runs of ODR with different random seeds.



Fig. 8: Distribution of change in SPEED score (E_{pose}^*) of the SPNv2 ($\phi = 6$, GN) after ODR with N = 4096 on SPEED+ HIL domains. Number of images in y-axis is shown in log scale.





Fig. 9: Visualization of examples on which the pose predictions of SPNv2 significantly improve after ODR on lightbox (*left*) and sunlamp (*right*) domains.



Table 5: Comparison of final poses predicted by SPNv2 and KRN³³ from the SPEED+ baseline studies.³² KRN's oracle performance denotes mean error from 5-fold cross validation on the respective SPEED+ HIL domains. ODR is not run on prisma25 due to limited available samples. Bold numbers indicate best performances excluding oracle.

| CNN Architecture | lightbox | | | sunlamp | | | prisma25 | | |
|---|-----------------|-----------------|-------------------------|-----------------|-----------------|-------------------------|----------------------------|-----------------|-----------------------|
| | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] | $E_{\rm T}$ [m] | $E_{\rm R}$ [°] | E_{pose}^* [-] | $\overline{E_{\rm T}}$ [m] | $E_{\rm R}$ [°] | E_{pose} [-] |
| KRN ³³ | 2.25 | 44.53 | 1.12 | 14.64 | 80.95 | 3.73 | 2.64 | 86.04 | 1.76 |
| + Style augmentation | 1.06 | 36.14 | 0.81 | 1.32 | 62.85 | 1.32 | 4.06 | 20.57 | 0.71 |
| Oracle | 0.24 | 6.15 | 0.15 | 0.19 | 5.33 | 0.13 | - | - | - |
| SPNv2 (Offline) | | | | | | | | | |
| $-\phi = 3$, BN | 0.175 | 6.479 | 0.142 | 0.225 | 11.065 | 0.230 | 1.572 | 5.202 | 0.216 |
| $-\phi = 6$, GN | 0.216 | 7.984 | 0.173 | 0.230 | 10.369 | 0.219 | 1.570 | 9.014 | 0.283 |
| SPNv2 (ODR) | | | | | | | | | |
| $-\phi = 3$, BN ($B = 4$, $N = 1024$) | 0.142 | 5.624 | 0.122 | 0.182 | 9.601 | 0.198 | - | - | - |
| $-\phi = 6, \text{GN} (N = 4096)$ | 0.150 | 5.577 | 0.122 | 0.161 | 9.788 | 0.197 | - | - | - |

improvement is much larger for the batch-agnostic architecture on lightbox. The most likely reason is that the ODR for the batch-agnostic architecture need not estimate the batch-wise feature normalization statistics which approximate those of the entire image domain. Therefore, the process is subject to significantly less uncertainty compared to that of the batch-dependent architecture.

The final scores of SPNv2 on SPEED+ and prisma25 are compared with ${\rm KRN^{33}}$ from the SPEED+ baseline studies³² in Table 5. Based on Figure 7, the ODR is performed with B = 4, N = 1024 for the batchdependent architecture and N = 4096 for the batchagnostic one. The pose predictions for prisma25 are from $h_{\rm H}$ only since its images are captured with a camera different from that of SPEED+. It shows that not only SPNv2 with just offline training is more accurate than KRN trained with and without style augmentation, but also its performance on lightbox is even better than the oracle performance of KRN, i.e., when trained directly on the lightbox images and labels. In other words, SPNv2 can generalize better to unseen lightbox images than KRN can after training on those exact images. It emphasizes a significant improvement of SPNv2 over KRN especially in its pose estimation architecture. SPNv2 also shows better performance on 25 spaceborne images of prisma25 compared to KRN for both variant of architectures. Finally, ODR for both variants of SPNv2 architecture reports around 15cm / 17cm translation errors and 5.6° / 9.7° orientation errors for lightbox and sunlamp, respectively. While there is no noticeable difference in the reported values, the ODR on the batch-agnostic SPNv2 closes a much bigger per-

formance gap on lightbox with no observable trend of failure of prediction after many iterations.

Finally, for the batch-agnostic SPNv2, Figure 8 shows the distribution of the changes in SPEED score after applying ODR, where a negative number suggests the score improved, and vice versa. For lightbox, a vast majority of images enjoys improvement in SPEED score after ODR, with the largest score drop surpassing 3. For sunlamp, the improvement is difficult to notice from the histogram, as significant number of images actually become more difficult to predict after ODR. Based on this result, the pose predictions of 4 samples that experience the largest improvement in SPEED score are visualized in Figure 9.

6. Discussion

Table 6: Number of tuned vs. entire learnable parameters in the feature encoder of SPNv2.

| Architecture | Total | Tuned | Percentage (%) |
|------------------------|-------------------|------------------|----------------|
| $\phi = 3$, BN | 12.0M | 106K | 0.88 |
| $\phi = 6, \text{GN}$ | $52.5 \mathrm{M}$ | $287 \mathrm{K}$ | 0.55 |

The proposed ODR is source-free, online and computationally efficient. Specifically, as shown in Table 6, the number of parameters tuned in the process of ODR is less than 1% of all learnable parameters in either variant of SPNv2. However, there are some limitations that need to be overcome in future work for the proposed method to be fully mission-compliant.

First, the study of offline training of SPNv2 reveals that the good generalizability across domain



gaps is best achieved by a larger network with batchdependent normalization layers such as BN. Note that the smaller variant of SPNv2 ($\phi = 3$, BN) already takes on average 0.63 seconds for inference on an $Intel(\hat{R})$ CoreTM i9-9900K CPU at 3.60GHz and 1.05 seconds for an iteration of ODR, which consists of inference, gradient backpropagation and online update step. Therefore, a bigger variant becomes prohibitively large for frequent inference and ODR updates on-board the satellite hardware. This is a problem, since the study on ODR clearly indicates that the batch-agnostic architecture performs better in the absence of the need to estimate the domain-wise feature normalization statistics. Therefore, a future work must aim to search for a smaller and more efficient CNN architecture that is batch-agnostic and can generalize its performance across domain gaps at least as well as SPNv2 can.

The second limitation is that, while ODR improves the average performance on the HIL domain images as a whole, its predictions still worsen significantly on many images as shown in Figure 8. An ideal refinement procedure would instead only introduce a non-negative change in the network's prediction accuracy on all images. This would require modifying the objective of the unsupervised learning or regularizing the loss function in such a way that prevents significant degradation of performance on individual images.

7. Conclusions

This paper presents Spacecraft Pose Network v2 (SPNv2), a multi-scale, multi-task convolutional neural network for pose estimation of noncooperative spacecraft across domain gaps. The SPNv2 architecture and its training mechanism are designed to comply with the realistic and operational constraints of machine learning in spaceborne applications. Specifically, SPNv2 is first trained offline on synthetic images of the SPEED+ dataset with extensive data augmentation and domain randomization. Then, at deployment, it refines the affine parameters of its feature extractor on unlabeled SPEED+ lightbox and sunlamp domains by self-supervised entropy minimization. It is shown that the combination of offline training and Online Domain Refinement (ODR) renders SPNv2 capable of bridging the domain gap, as evidenced by 0.15 m and 5.58° errors on lightbox and 0.16 m and 9.79° errors on sunlamp for a larger, batch-agnostic variant of SPNv2. Moreover, the proposed ODR is source-free, online, and computationally efficient, as it only refines less than 1% of the learnable parameters in the feature extractor of SPNv2. The

paper discusses a number of limitations of SPNv2 and ODR that must be addressed to build a fully missioncompliant model, which will be fully batch-agnostic, requires less computation, and does not deteriorate the prediction on the target domain samples throughout ODR. Overall, the methodology presented in this paper can easily extend to any spaceborne images of a known, noncooperative target collected in future missions. In the future, SPNv2 and ODR will be integrated into a navigation filter (e.g., unscented Kalman filter) to assess its capacity as an image processing unit of the filter in a number of close-range rendezvous scenarios. The uncertainty of the SPNv2's predictions and its impact on the filter performance will also be thoroughly evaluated.

Acknowledgement

This work is supported by the Air Force Office of Scientific Research (AFOSR) via Centauri under the project titled Modular State-Adaptive Landmark Tracking. We would like to thank OHB Sweden for the 3D model of the Tango spacecraft used to create the images used in this article and for the flight images in the prisma25 dataset. We acknowledge the Advanced Concepts Team (ACT) of the European Space Agency (ESA) for their contribution to putting together the SPEED+ dataset. We would also like to thank the Stanford Research Computing Center for providing computational resources that contributed to these research results.

References

- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2007.
- [2] Kevin Black, Shrivu Shankar, Daniel Fonseka, Jacob Deutsch, Abhimanyu Dhir, and Maruthi Akella. Real-time, flight-ready, non-cooperative spacecraft pose estimation using monocular imagery. In 31st AAS/AIAA Space Flight Mechanics Meeting, February 2021.
- [3] Yannick Bukschat and Marcus Vetter. Efficientpose: An efficient, accurate and scalable end-toend 6d multi object pose estimation approach, 2020.
- [4] Alexander Buslaev, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A. Kalinin. Albumentations: Fast

and flexible image augmentations. Information, 11(2), 2020.

- [5] Rich Caruana. Multitask learning. Mach. Learn., 28(1):41-75, jul 1997.
- [6] Lorenzo Pasqualetto Cassinis, Robert Fonod, Eberhard Gill, Ingo Ahrns, and Jesus Gil Fernandez. Cnn-based pose estimation system for close-proximity operations around uncooperative spacecraft. In AIAA Scitech 2020 Forum, 2020.
- [7] Bo Chen, Jiewei Cao, Álvaro Parra Bustos, and Tat-Jun Chin. Satellite pose estimation with deep landmark regression and nonlinear pose refinement. 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pages 2816–2824, 2019.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1597– 1607, 2020.
- [9] Simone D'Amico, Mathias Benn, and John L. Jørgensen. Pose estimation of an uncooperative spacecraft from actual space imagery. *International Journal of Space Science and Engineering*, 2(2):171, 2014.
- [10] Simone D'Amico, Per Bodin, Michel Delpech, and Ron Noteborn. PRISMA. In Marco D'Errico, editor, Distributed Space Missions for Earth System Monitoring Space Technology Library, volume 31, chapter 21, pages 599–637. 2013.
- [11] Jun Deng, Wei Dong, Richard Socher, Li-Jia Li, Kuntai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 248–255, 2009.
- [12] Jason L. Forshaw, Guglielmo S. Aglietti, Nimal Navarathinam, Haval Kadhem, Thierry Salmon, Aurélien Pisseloup, Eric Joffre, Thomas Chabot, Ingo Retat, Robert Axthelm, and et al. RemoveDEBRIS: An in-orbit active debris removal demonstration mission. Acta Astronautica, 127:448–463, 2016.
- [13] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In Francis Bach and David Blei, editors, Proceedings of the 32nd International Conference on Machine Learning, volume 37 of Proceedings

of Machine Learning Research, pages 1180–1189, Lille, France, 07–09 Jul 2015. PMLR.

- [14] Golnaz Ghiasi, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network. In *British Machine* Vision Conference, 2017.
- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, Proceedings of the 32nd International Conference on Machine Learning, volume 37 of Proceedings of Machine Learning Research, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
- [17] Philip T Jackson, Amir Atapour-Abarghouei, Stephen Bonner, Toby P Breckon, and Boguslaw Obara. Style augmentation: Data augmentation via style randomization. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pages 83–92, 2019.
- [18] Mate Kisantal, Sumant Sharma, Tae Ha Park, Dario Izzo, Marcus Märtens, and Simone D'Amico. Satellite pose estimation challenge: Dataset, competition design and results. *IEEE Transactions on Aerospace and Electronic Systems*, 56(5):4083–4098, 2020.
- [19] Donald E. Knuth. The Art of Computer Programming, Vol. 1: Fundamental Algorithms. Addison-Wesley, Reading, Mass., third edition, 1997.
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 25. Curran Associates, Inc., 2012.
- [21] Jogendra Nath Kundu, Naveen Venkat, Rahul M V, and R. Venkatesh Babu. Universal sourcefree domain adaptation. June 2020.
- [22] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. EPnP: An accurate O(n) solution to the PnP problem. *International Journal of Computer Vision*, 81(2):155–166, 2008.

- [23] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.
- [24] Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation, 2016.
- [25] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 6028–6039, 2020.
- [26] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In 2017 IEEE International Conference on Computer Vision (ICCV), pages 2999–3007, 2017.
- [27] Yuejiang Liu, Parth Kothari, Bastien Germain van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. TTT++: When does selfsupervised test-time training fail or thrive? In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [28] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. CoRR, abs/1711.05101, 2017.
- [29] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In European Conference on Computer Vision (ECCV), 2016.
- [30] Tae Ha Park, Juergen Bosse, and Simone D'Amico. Robotic testbed for rendezvous and optical navigation: Multi-source calibration and machine learning use cases. In 2021 AAS/AIAA Astrodynamics Specialist Conference, Big Sky, Vitrual, August 9-11 2021.
- [31] Tae Ha Park, Marcus Märtens, Gurvan Lecuyer, Dario Izzo, and Simone D'Amico. Next generation spacecraft pose estimation dataset (SPEED+). Stanford Digital Repository, 2021. Available at https://purl.stanford. edu/wv398fc4383.
- [32] Tae Ha Park, Marcus Märtens, Gurvan Lecuyer, Dario Izzo, and Simone D'Amico. SPEED+: Next-generation dataset for spacecraft pose estimation across domain gap. arXiv preprint arXiv:2110.03101, 2021.

- [33] Tae Ha Park, Sumant Sharma, and Simone D'Amico. Towards robust learning-based pose estimation of noncooperative spacecraft. In 2019 AAS/AIAA Astrodynamics Specialist Conference, Portland, Maine, August 11-15 2019.
- [34] Lorenzo Pasqualetto Cassinis, Alessandra Menicucci, Eberhard Gill, Ingo Ahrns, and Jesus Gil Fernandez. On-ground validation of a cnn-based monocular pose estimation system for uncooperative spacecraft. In 8th European Conference on Space Debris, volume 8, 2021.
- [35] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. VisDA: The visual domain adaptation challenge, 2017.
- [36] Pedro F Proenca and Yang Gao. Deep learning for spacecraft pose estimation from photorealistic rendering. arXiv preprint arXiv:1907.04298, 2019.
- [37] Benjamin B. Reed, Robert C. Smith, Bo J. Naasz, Joseph F. Pellegrino, and Charles E. Bacon. The Restore-L servicing mission. *Aiaa Space 2016*, 2016.
- [38] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948.
- [39] Sumant Sharma and Simone D'Amico. Pose estimation for non-cooperative spacecraft rendezvous using neural networks. In 2019 AAS/AIAA Space Flight Mechanics Meeting, Ka'anapali, Maui, HI, January 13-17 2019.
- [40] Sumant Sharma and Simone D'Amico. Neural network-based pose estimation for noncooperative spacecraft rendezvous. *IEEE Transactions* on Aerospace and Electronic Systems, 56(6):4638– 4658, 2020.
- [41] Sumant Sharma, Tae Ha Park, and Sumant D'Amico. Spacecraft pose estimation dataset (speed). Stanford Digital Repository. Available at: https://doi.org/10.25740/dz692fn7184, 2019.
- [42] Connor Shorten and Taghi M. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6, 2019.
- [43] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In ECCV 2016 Workshops, 2016.



- [44] Yu Sun, Xiaolong Wang, Liu Zhuang, John Miller, Moritz Hardt, and Alexei A. Efros. Test-time training with self-supervision for generalization under distribution shifts. In *ICML*, 2020.
- [45] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 09–15 Jun 2019.
- [46] Mingxing Tan, Ruoming Pang, and Quoc V. Le. EfficientDet: Scalable and efficient object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.
- [47] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 23–30, 2017.
- [48] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2962–2971, 2017.
- [49] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In International Conference on Learning Representations, 2021.
- [50] Yuxin Wu and Kaiming He. Group normalization. In Proceedings of the European Conference on Computer Vision (ECCV), September 2018.
- [51] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *The AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [52] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), 2020.
- [53] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the Continuity of Rotation

Representations in Neural Networks. In *Proceedings of CVPR*, page 9, Long Beach, CA, June 2019. IEEE.