

Making the Flow Glow – Robot Perception under Severe Lighting Conditions using Normalizing Flow Gradients

Simon Kristoffersson Lind^{1,*}

Rudolph Triebel²

Volker Krüger¹

Abstract—Modern robotic perception is highly dependent on neural networks. It is well known that neural network-based perception can be unreliable in real-world deployment, especially in difficult imaging conditions. Out-of-distribution detection is commonly proposed as a solution for ensuring reliability in real-world deployment. Previous work has shown that normalizing flow models can be used for out-of-distribution detection to improve reliability of robotic perception tasks. Specifically, camera parameters can be optimized with respect to the likelihood output from a normalizing flow, which allows a perception system to adapt to difficult vision scenarios. With this work we propose to use the absolute gradient values from a normalizing flow, which allows the perception system to optimize local regions rather than the whole image. By setting up a table top picking experiment with exceptionally difficult lighting conditions, we show that our method achieves a 60% higher success rate for an object detection task compared to previous methods.

I. INTRODUCTION

In robotics, neural networks have emerged as an indispensable tool for perception and computer vision tasks. However, they operate under the assumption that all inputs come from the same distribution as their training data. In other words, when confronted with unfamiliar inputs, their outputs can be unpredictable [1]. This unpredictability poses a significant risk for robots operating in dynamic real-world settings. Moreover, the accuracy of computer vision hinges on the camera’s ability to produce clear images for the neural network. A case in point is the 2015 Tesla crash, where the autopilot system failed to detect a truck’s side against a bright sky [2].

One common way to prevent unpredictable behaviour is out-of-distribution (OOD) detection [1]. Most OOD methods work by directly or indirectly approximating the probability density function (PDF) of the training data. The approximate PDF is then used as a score function to determine whether an input is in- or out-of-distribution. Using such a score, an autonomous system can react by performing a safety behaviour whenever an input is deemed too far from the training distribution. Most commonly, this safety behaviour simply discards any OOD input while waiting for better inputs.

While discarding inputs may be feasible for many applications, it can be problematic for many autonomous systems.

¹Lund University LTH, Lund, Sweden
simon.kristoffersson_lind@cs.lth.se,
volker.krueger@cs.lth.se

²German Aerospace Center, Oberpfaffenhofer, Germany,
and Karlsruhe Institute of Technology, Karlsruhe, Germany
rudolph.triebel@dlr.de

*Corresponding Author



Fig. 1. In the top image, taken with default camera parameters and auto-exposure, only the cup can be detected by YOLOv4. In the bottom image, where camera parameters have been optimized using our proposed method, the cup, spoon, fork, scissors, and book are successfully detected.

Consider an autonomous car that encounters OOD data. Should it keep driving while discarding inputs, and risk hitting someone? Or should it stop while waiting for better inputs, and risk an accident by blocking traffic? We argue that it is necessary for autonomous systems to detect and adapt to OOD data.

In our earlier work [3], we used a Normalizing Flow (NF) as an OOD detector, and we created an adaptable vision system by allowing our robot to change parameters in its camera on-the-fly. We chose to implement this behaviour as an optimization procedure, where the parameters *backlight compensation*, *brightness*, *exposure*, *gain*, *saturation*, and *sharpness* were optimized with respect to the OOD score




Fig. 2. Left: Images from the COCO validation set. right: Absolute gradient images from YOLOv4 + Normalizing Flow.

from our NF. Effectively, our robot learns the best camera parameters for any given scene, by minimizing the OOD score. We were able to show that this type of adaptive behaviour improves the performance of a YOLOv4 object detector in difficult visual scenarios.

A major weakness of our earlier method is the fact that the NF evaluates an entire image as either in- or out-of-distribution. Consider the example in Fig. 1. Due to the bright light, it is not possible to find any single camera parameter setting that makes the entire image visible. Despite this, the objects around the bright light can be made almost perfectly visible by changing camera parameters. Therefore, we see a need to express OOD detection on regions rather than the full image.

When inspecting the pixel-wise gradient with regard to the OOD score from our NF model, it seems that OOD regions correspond to large absolute gradient values, and in-distribution regions correspond to small absolute gradient values. This effect can be seen in Fig. 2.

In this work, we propose to use the absolute pixel-wise NF gradient values as a local OOD score. We conduct two experiments to provide empirical evidence that the absolute NF gradient values carry OOD information at the pixel-level. Finally, we show that our proposed absolute NF gradient outperforms our previous method [3] on a real-world robotic picking experiment in a difficult visual scenario. To summarize, we make the following contributions:

- We propose to use the absolute gradient from an NF to provide pixel-level OOD information.
- We empirically show the efficacy of our proposed NF gradient for OOD detection.
- We show, across two common object detection architectures (YOLOv4[4], Faster-RCNN[5]) that object

detectors tend to perform better in regions with small absolute NF gradient values.

- We experimentally establish the superior performance of our method over [3] in adapting camera parameters for challenging imaging scenarios.

II. RELATED WORK

A substantial body of research focuses on OOD detection, such as [6], [7], [8], [9], [10], [11]. However, these studies predominantly validate their findings using benchmarks. The benchmarks employed generally fail to capture the diverse scenarios a robot might face in real-world operations. This limitation arises primarily because these works ensure that in-distribution and out-of-distribution inputs come from distinct sets.

A more niche segment of research delves into the application of OOD detection in robotics, for example [12], [13]. These studies, while applying OOD detection to robotics, typically resort to aborting operations upon encountering OOD data.

Another research avenue explores OOD techniques for visual anomaly detection, with studies like [14], [15], [16]. These investigations typically employ autoencoder methods to identify OOD elements within images, such as in medical imaging contexts. However, the OOD strategies in these studies are usually tailored to their specific applications, limiting their broader applicability.

Lastly, [3] leverages an NF to adapt to challenging imaging situations by fine-tuning camera parameters. Our contribution enhances this by using the NF’s gradient to incorporating pixel-level OOD information.

III. BACKGROUND

Given a training distribution $p_{\text{train}}(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^D$, the goal of OOD detection is to identify whether or not a new sample $\hat{\mathbf{x}}$ comes from $p_{\text{train}}(\mathbf{x})$. We assume that $p_{\text{train}}(\mathbf{x})$ is unknown or intractable to evaluate – otherwise OOD detection would be trivial. NFs enable OOD detection by approximating p_{train} [17].

Formally, an NF is a *diffeomorphism*, an invertible transformation $T(\mathbf{x})$ where both $T(\mathbf{x})$ and $T^{-1}(\mathbf{x})$ are differentiable [18]. In training an NF, the goal is to learn this diffeomorphism such that $\mathbf{u} = T(\mathbf{x})$, $\mathbf{u} \in \mathbb{R}^D$, follows a different, tractable distribution $p_u(\mathbf{u})$. The distribution $p_u(\mathbf{u})$ is typically chosen to be a unit Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{1})$, since a Gaussian is easy to evaluate and sample. T being a diffeomorphism allows $p_{\text{train}}(\mathbf{x})$ to be approximated using the following formula [18]:

$$\begin{aligned} p_{\text{train}}(\mathbf{x}) &= p_u(\mathbf{u}) |\det J_{T^{-1}}(\mathbf{u})|^{-1} \\ &= p_u(T(\mathbf{x})) |\det J_T(\mathbf{x})|, \end{aligned}$$

where J_T denotes the Jacobian of T . As such, the NF allows us to approximate p_{train} by transforming it into a tractable distribution. In the standard OOD detection setting, an OOD threshold is set with respect to $p_{\text{train}}(\mathbf{x})$, typically based on a fixed false positive rate [19].

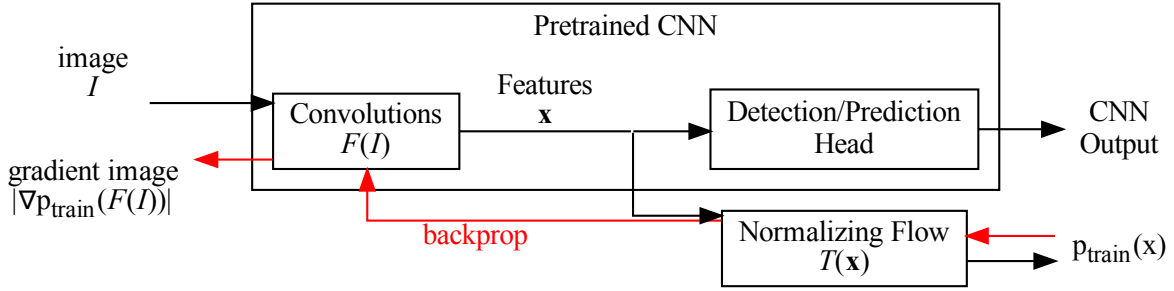


Fig. 3. Illustration of how we apply NFs alongside a pretrained CNN for vision tasks.

The primary difficulty in constructing an NF is ensuring invertibility, along with a tractable $|\det J_T(\mathbf{x})|$. One common way to solve these problems is *affine coupling layers*, used in for example RealNVP[20] and GLOW[21]. Coupling layers divide their input \mathbf{x} into two non-overlapping parts $\mathbf{x}_1, \mathbf{x}_2$. An affine transformation is then performed on one part, while leaving the other unchanged:

$$\mathbf{x}'_1 = \boldsymbol{\alpha} \cdot \mathbf{x}_1 + \boldsymbol{\beta},$$

where $\boldsymbol{\alpha} = \exp(F_1(\mathbf{x}_2))$, $\boldsymbol{\beta} = F_2(\mathbf{x}_2)$.

Here we use \cdot to denote element-wise multiplication. F_1 and F_2 denote arbitrary functions, for example neural networks. The output of the coupling layer is then computed by simply reconstructing \mathbf{x} using the new \mathbf{x}'_1 and the unchanged \mathbf{x}_2 . Commonly, \mathbf{x}_1 and \mathbf{x}_2 are constructed by simply splitting \mathbf{x} in half, in which case the output is constructed by simply concatenating: $\mathbf{x}' = [\mathbf{x}'_1 | \mathbf{x}_2]$ [18]. In this setting, the layers typically operate on alternating halves of \mathbf{x} , otherwise one half would be left permanently unchanged.

Kirichenko *et al.* [17] observe that applying NFs directly to images tends to work poorly for OOD detection. They show that NFs seem to learn pixel-to-pixel correspondences, rather than extracting any semantic information about the image distribution. As a solution, they propose to apply an NF to features extracted from another pretrained neural network.

In our previous work [3], we adopted the approach from [17], and applied it to a robotic perception setting. Instead of using a pretrained network only for feature extraction, we applied an NF alongside a pretrained network for a vision task. This approach allows a single network to perform a vision task, while also being used as a feature extractor for the NF as illustrated in Fig. 3. In turn, this allows the NF to provide an OOD score for each image with negligible runtime overhead. Using this NF construction, we applied an evolutionary optimization procedure [22] to minimize $p_{\text{train}}(\mathbf{x})$ by adjusting camera parameters (backlight compensation, brightness, contrast, exposure, gain, saturation, and sharpness). We implemented this optimization procedure in our robot, using YOLOv4[4] as the pretrained CNN. By deploying our robot in a number of difficult visual scenarios, we were able to show an increase in object detection performance using our optimization procedure (average increase of 3.86 percentage units in F1 score).

IV. OUR APPROACH

Instead of computing a global OOD score, we extend the NF approach from [3] and [17] by computing gradients to extract pixel-level OOD information. Formally, given an input image

$$I = \begin{bmatrix} i_{11} & i_{12} & \dots & i_{1W} \\ i_{21} & i_{22} & \dots & i_{2W} \\ \vdots & \vdots & \ddots & \vdots \\ i_{H1} & i_{H2} & \dots & i_{HW} \end{bmatrix}, \quad I \in \mathbb{R}^{H \times W},$$

the OOD detection method in [3], [17] computes:

$$\mathbf{x} = F(I), \quad F: \mathbb{R}^{H \times W} \mapsto \mathbb{R}^D$$

$$p_{\text{train}}(\mathbf{x}) = p_u(T(\mathbf{x})) |\det J_T(\mathbf{x})|, \quad (1)$$

where F denotes the feature extraction layers from a pretrained neural network, and T denotes the learned NF.

In order to gain pixel-level OOD information, we propose to use the gradient of Eq. (1). Specifically, we propose to use the absolute value of each element in the gradient. With some notational abuse, we write

$$|\nabla p_{\text{train}}(F(I))| \in \mathbb{R}^{H \times W}$$

$$|\nabla p_{\text{train}}(F(I))|_{yx} = \left| \frac{\delta p_{\text{train}}}{\delta i_{yx}} \right|. \quad (2)$$

In other words, we simply backpropagate through p_{train} and F , as indicated by the red edges in Fig. 3. Here we use 2-dimensional grayscale images purely for notational convenience. Our method extends directly to RGB images by averaging across the channel dimension.

Through initial testing we have observed that image regions that closely match the distribution p_{train} tend to have small gradient values compared to regions that are OOD. Intuitively, we reason that in-distribution data will likely lie close to a maximum of $p_{\text{train}}(\mathbf{x})$. Due to the continuity and differentiability of $p_{\text{train}}(\mathbf{x})$, these data points should exhibit small gradient values. This intuition also serves as our rationale for using the absolute value of gradients, as the direction of change is inconsequential.

Just as before, our goal is to construct an optimization procedure over camera parameters, in order to allow our robot to adaptively overcome difficult visual scenarios. In our robotics work, we mostly perform object detection tasks, for example object picking. All object detection tasks involve

some type of region proposal, whether it be bounding-boxes or image segmentations. For such tasks, we propose to optimize:

$$\min_{\theta} \frac{1}{|\text{ROI}|} \sum_{x,y \in \text{ROI}} |\nabla p_{\text{train}}(F(I(\theta)))|_{xy} ,$$

where ROI denotes the region of interest proposed by the object detector, and θ denotes camera parameters. We use $|\text{ROI}|$ to represent the number of pixels inside the region. Here we also write $I(\theta)$ to highlight the dependence of the image on the camera parameters. When the object detector fails to propose a suitable region of interest, we fall back to computing the average over the entire absolute NF gradient image $|\nabla p_{\text{train}}(F(I(\theta)))|$. We expect this fallback behaviour to perform similarly to optimizing p_{train} , as in [3]. Additionally, for multi-object detectors, we apply the sum over the union of all proposed regions:

$$\text{ROI} = \bigcup_{\text{ROI}_i} \text{ROI}_i .$$

For other types of visual tasks, it may be necessary to construct other ways to select suitable regions to optimize over.

V. EXPERIMENTS AND RESULTS

In this section, we begin by performing two separate experiments to empirically establish evidence supporting our proposed optimization procedure.

First, we perform a conventional OOD detection experiment using the average absolute NF gradient as an OOD score:

$$\frac{1}{HW} \sum_{y=1}^H \sum_{x=1}^W \left| \frac{\delta p_{\text{train}}}{\delta i_{yx}} \right| . \quad (3)$$

Showing that the absolute NF gradient can be used successfully for OOD detection directly implies that it is useful for constructing an adaptive optimization procedure as in [3]. Additionally, showing success in OOD detection using the average absolute gradient over the entire image supports our decision to use this as a fallback when our object detection fails to propose a region of interest.

Second, we investigate the relationship between object detector performance, and the absolute NF gradient. Specifically, we look at the predicted bounding-boxes from two common object detectors (YOLOv4[4], Faster-RCNN[5]), and we show that bounding-boxes that are predicted correctly tend to contain smaller absolute NF gradient values compared to incorrectly predicted bounding-boxes. This experiment serves directly to support our decision to optimize camera parameters over the predicted regions of interest from our object detector.

Finally, we conduct a real-world experiment using our robot, where we set up an object picking task in a cluttered scene. We manually construct an exceedingly difficult visual scenario by making our lab as dark as possible, while aiming a bright light at our robot. Our proposed absolute NF gradient optimization procedure is then compared to [3], and simple auto-exposure and auto-whitebalance.

A. NF Training

In all of our experiments we use an NF consisting of 10 affine coupling layers, operating on alternating halves of the input. Each coupling layer consists of two shared ReLU-activated linear layers followed by two separate prediction heads, one for α and one for β . Our NF is trained for 200 epochs, using the Adam[23] optimizer with a constant learning rate of 10^{-4} .

B. OOD Experiment

We run the same conventional OOD detection experiment as in [17], in order to compare performance of our proposed average absolute NF gradient to that of using p_{train} directly. For feature extraction, we use an EfficientNet [24] model pre-trained on ImageNet [25]. We perform OOD detection on the CIFAR-10 [26], SVHN [27], and CelebA [28] datasets. An NF model is trained for each of these datasets, which results in the other two datasets being considered OOD. First, we compute OOD detection performance using $p_{\text{train}}(F(I))$ as an OOD score directly, exactly like [17]. Then, we compute OOD detection performance using the average absolute NF gradient (Eq. 3). We calculate an OOD threshold with a fixed 5% false-positive rate. The resulting true-positive rates are shown in Table I and II respectively.

From Table II, it is evident that the NF gradient works well for OOD detection. While we do observe a minor decrease in overall OOD detection performance, our NF gradient is not far behind the plain NF approach, and even achieves improved performance on CIFAR10/SVHN. We would especially like to stress that we do not aim to achieve state-of-the-art OOD detection performance, but only highlight the usefulness of accessing this type of information at a pixel-level.

C. Correlation with Model Performance

In order to further establish our proposed method for adapting camera parameters, we wish to investigate whether the absolute NF gradient correlates with object detection performance. We perform an experiment using YOLOv4[4] and Faster-RCNN[5], both trained on the COCO[29] dataset.

TABLE I
OOD DETECTION TRUE-POSITIVE RATES USING $p_{\text{TRAIN}}(F(I))$.

		OOD		
		Cifar10	SVHN	CelebA
Train	Cifar10	-	0.4148	0.9953
	SVHN	0.999	-	1.0
	CelebA	1.0	1.0	-

TABLE II
OOD DETECTION TRUE-POSITIVE RATES USING AVERAGE ABSOLUTE NF GRADIENT.

		OOD		
		Cifar10	SVHN	CelebA
Train	Cifar10	-	0.6688	0.9205
	SVHN	0.972	-	0.9999
	CelebA	0.9464	0.9747	-



Fig. 5. Example images captured during experiments. Top left: default parameters. Top right: Default parameters with auto-exposure and auto-whitebalance. Bottom left: parameters optimized according to [3]. Bottom right: parameters optimized using our NF gradient image.

TABLE IV

NUMBER OF CORRECT OBJECT DETECTIONS. DEFAULT: CAMERA'S DEFAULT PARAMETERS. AUTO: CAMERA'S DEFAULT PARAMETERS WITH AUTO-EXPOSURE AND AUTO-WHITEBALANCE. NF: PARAMETERS OPTIMIZED W.R.T. $p_{\text{TRAIN}}(\mathbf{x})$ AS IN [3]. NF GRADIENT: PARAMETERS OPTIMIZED W.R.T. OUR PROPOSED NF GRADIENT IMAGE.

	Default	Auto	NF	NF Gradient
apple	0	0	0	0
banana	1	8	6	16
book	5	2	6	11
cup	9	12	12	18
fork	24	19	34	45
mouse	0	0	0	0
scissors	2	0	4	10
spoon	4	1	9	14
Total:	45	42	71	114

apple and *mouse*. Since they were placed close to the light source, they were typically more difficult to discern. This further highlights the importance of local OOD information compared to a single global OOD score. By computing a local OOD score, our NF gradient method allows the perception pipeline to focus only on the regions that matter, and make sure those are visible.

We would like to note that it is likely possible to further improve the object detection performance. In our experiment, camera parameters were jointly optimized for all detected objects. However, we believe that individual optimization for each object region might yield even better performance.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we introduce a technique that employs gradients from a Normalizing Flow (NF) to obtain pixel-level Out-of-Distribution (OOD) information for an image. Experimental results validate the efficacy of our proposed NF gradient for general OOD detection. Moreover, this approach significantly surpasses our prior methods for adapting to challenging imaging scenarios. Consequently, our robotic perception framework demonstrates enhanced object detection capabilities, even in the face of adverse lighting conditions.

We will make our code publicly available at <https://github.com/SimonKLind/NF-Gradients>. Additionally, we will compile the images collected from our experiment into a dataset, which we will also make publicly available for future research.

A notable limitation of our current method is its dependency on the camera to sample parameter configurations

during the optimization procedure. The process of acquiring multiple image samples is inherently time-consuming, resulting in a runtime that is slower than desired, which may prohibit its use in some robot applications. To broaden the applicability of our techniques across diverse robotic tasks, we hope to improve the runtime by removing the dependency on the camera for sampling.

ACKNOWLEDGMENTS

This research is funded by the Excellence Center at Linköping-Lund in Information Technology (ELLIIT), and the Wallenberg AI, Autonomous Systems and Software Program (WASP). Computations for this publication were enabled by the supercomputing resource Berzelius provided by the National Supercomputer Centre at Linköping University and the Knut and Alice Wallenberg foundation.

REFERENCES

- [1] J. Gawlikowski, C. Rovile Njeutcheu Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, M. Shahzad, W. Yang, R. Bamler, and X. X. Zhu, "A survey of uncertainty in deep neural networks," *arXiv e-prints*, Jul 2021, aDS Bibcode: 2021arXiv210703342G.
- [2] Tesla, "A tragic loss," <https://www.tesla.com/blog/tragic-loss>, accessed 2023-09-15.
- [3] S. Kristofferson Lind, R. Triebel, L. Nardi, and V. Krueger, "Out-of-distribution detection for adaptive computer vision," in *Image Analysis*, ser. Lecture Notes in Computer Science, R. Gade, M. Felsberg, and J.-K. Kämäräinen, Eds. Cham: Springer Nature Switzerland, 2023, pp. 311–325.
- [4] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "Yolov4: Optimal speed and accuracy of object detection," Apr 2020, arXiv:2004.10934 [cs, eess].
- [5] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," Jan 2016, arXiv:1506.01497 [cs].
- [6] L. Deecke, R. Vandermeulen, L. Ruff, S. Mandt, and M. Kloft, "Image anomaly detection with generative adversarial networks," in *Machine Learning and Knowledge Discovery in Databases*, M. Berlingerio, F. Bonchi, T. Gärtner, N. Hurley, and G. Ifrim, Eds. Cham: Springer International Publishing, 2019, pp. 3–17.
- [7] Y.-C. Hsu, Y. Shen, H. Jin, and Z. Kira, "Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [8] S. Liang, Y. Li, and R. Srikant, "Principled detection of out-of-distribution examples in neural networks," *CoRR*, vol. abs/1706.02690, 2017.
- [9] J. Winkens, R. Bunel, A. G. Roy, R. Stanforth, V. Natarajan, J. R. Ledsam, P. MacWilliams, P. Kohli, A. Karthikesalingam, S. Kohl, A. T. Cemgil, S. M. A. Eslami, and O. Ronneberger, "Contrastive training for improved out-of-distribution detection," *CoRR*, vol. abs/2007.05566, 2020.
- [10] J. Ren, P. J. Liu, E. Fertig, J. Snoek, R. Poplin, M. Depristo, J. Dillon, and B. Lakshminarayanan, "Likelihood ratios for out-of-distribution detection," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019.
- [11] S. Mohseni, M. Pitale, J. Yadawa, and Z. Wang, "Self-supervised learning for generalizable out-of-distribution detection," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, pp. 5216–5223, April 2020.
- [12] M. Yuhas, Y. Feng, D. J. X. Ng, Z. Rahiminasab, and A. Easwaran, "Embedded out-of-distribution detection on an autonomous robot platform," in *Proceedings of the Workshop on Design Automation for CPS and IoT*, ser. Destion '21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 13–18.
- [13] R. McAllister, G. Kahn, J. Clune, and S. Levine, "Robustness to out-of-distribution inputs via task-aware generative uncertainty," in *International Conference on Robotics and Automation (ICRA)*, 2019, pp. 2083–2089.
- [14] Y. Du, "An anomaly detection method using deep convolution neural network for vision image of robot," *Multimedia Tools and Applications*, vol. 79, no. 13, pp. 9629–9642, Apr 2020.
- [15] D. J. Samuel and F. Cuzzolin, "Unsupervised anomaly detection for a smart autonomous robotic assistant surgeon (saras) using a deep residual autoencoder," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7256–7261, Oct 2021.
- [16] S. Salimpour, J. P. Queralta, and T. Westerlund, "Self-calibrating anomaly and change detection for autonomous inspection robots," in *2022 Sixth IEEE International Conference on Robotic Computing (IRC)*, Dec 2022, pp. 207–214.
- [17] P. Kirichenko, P. Izmailov, and A. G. Wilson, "Why Normalizing Flows Fail to Detect Out-of-Distribution Data," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 20 578–20 589.
- [18] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, "Normalizing Flows for Probabilistic Modeling and Inference," *Journal of Machine Learning Research*, vol. 22, no. 57, pp. 1–64, 2021.
- [19] J. Yang, K. Zhou, Y. Li, and Z. Liu, "Generalized out-of-distribution detection: A survey," Aug 2022, arXiv:2110.11334 [cs].
- [20] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using real nvp," Feb 2017, arXiv:1605.08803 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1605.08803>
- [21] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1x1 convolutions," Jul 2018, arXiv:1807.03039 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1807.03039>
- [22] C. Audet and W. Hare, *Derivative-Free and Blackbox Optimization*. Springer, 2017.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [24] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," Sep 2020, arXiv:1905.11946 [cs, stat].
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [26] A. Krizhevsky, G. Hinton, *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [27] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," in *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- [28] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [29] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár, "Microsoft coco: Common objects in context," Feb 2015, arXiv:1405.0312 [cs].