

Augmenting Tactile Simulators with Real-like and Zero-Shot Capabilities

Osher Azulay* Alon Mizrahi*, Nimrod Curtis* and Avishai Sintov

Abstract—Simulating tactile perception could potentially leverage the learning capabilities of robotic systems in manipulation tasks. However, the reality gap of simulators for high-resolution tactile sensors remains large. Models trained on simulated data often fail in zero-shot inference and require fine-tuning with real data. In addition, work on high-resolution sensors commonly focus on ones with flat surfaces while 3D round sensors are essential for dexterous manipulation. In this paper, we propose a bi-directional Generative Adversarial Network (GAN) termed SightGAN. SightGAN relies on the early CycleGAN while including two additional loss components aimed to accurately reconstruct background and contact patterns including small contact traces. The proposed SightGAN learns real-to-sim and sim-to-real processes over difference images. It is shown to generate real-like synthetic images while maintaining accurate contact positioning. The generated images can be used to train zero-shot models for newly fabricated sensors. Consequently, the resulted sim-to-real generator could be built on top of the tactile simulator to provide a real-world framework. Potentially, the framework can be used to train, for instance, reinforcement learning policies of manipulation tasks. The proposed model is verified in extensive experiments with test data collected from real sensors and also shown to maintain embedded force information within the tactile images.

I. INTRODUCTION

Tactile sensing is a fundamental aspect of human perception and, therefore, is a topic for extensive research in robotics [1]–[3]. Such sensing plays a crucial role in enabling robots to interact with the physical world potentially with precision and dexterity. The advancement of tactile sensor technologies has led to high-dimensional and complex data representations which, in turn, requires sufficient data in order to train accurate models and policies [4]. Hence, tactile simulations and sim-to-real approaches are gaining momentum as researchers explore new ways to bridge the gap between virtual and physical worlds [5], [6].

Tactile sensors come in various technologies including capacitive transducers [7], force sensitive resistors [8] and piezo-resistors [9]. However, they usually fit to specific applications and provide low-resolution data. Optical-based tactile sensors, on the other hand, have become increasingly common due to their ability to provide high-resolution signals [10]–[13]. In such sensors, an internal camera observes the deformation of the contact pad, typically made of a soft elastomer, during contact with an object. An image captured by the internal camera has the potential to encode essential

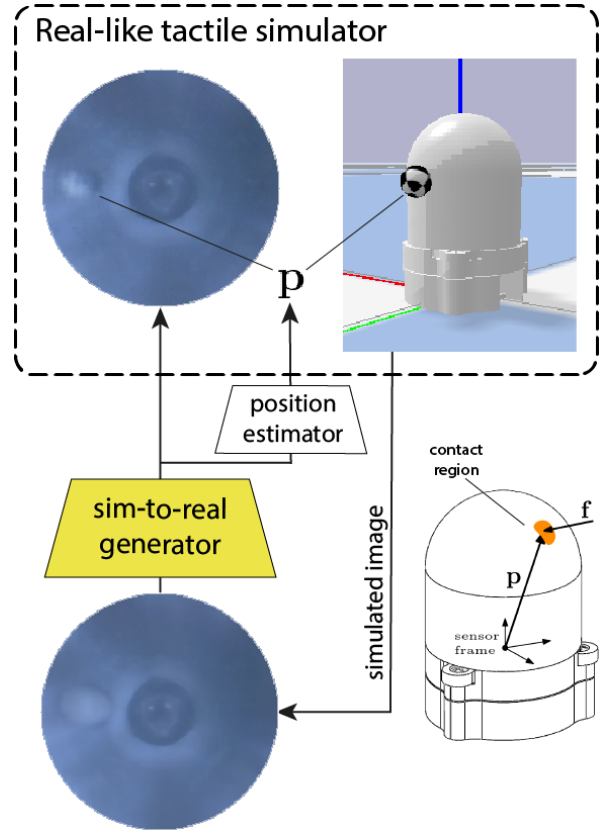


Fig. 1: The sim-to-real generator from the trained SightGAN model is used to map simulated tactile images to real-like images of a 3D round tactile sensor. Since the generated image is close to reality, various models can be trained using the simulator. In this example, a position estimator can provide accurate labeling to the image making it a fine simulator for various tasks.

information regarding the contact including its position with respect to the sensor’s coordinate frame. Nevertheless, due to the rich data in such images, training models to estimate features in the data requires an extensive amount of samples.

In order to cope with the data amount requirements, simulations of optical-based tactile sensors have been addressed and have the potential to rapidly generate large datasets of tactile images [14], [15]. However, transferring a model trained on simulated data to a real sensor, i.e., sim-to-real, may present difficulties. Primarily, real-world tactile images often exhibit substantial disparities when compared to their simulated counterparts [16], [17]. In attempt to bridge the gap between simulated tactile data and real-world tactile information, various approaches have been explored. Domain randomization was included in simulation of the TacTip sensor where some parameters are constantly randomised [18]. Similarly, a simulation was augmented with random texture

* These authors contributed equally.

O. Azulay, A. Mizrahi, N. Curtis and A. Sintov are with the School of Mechanical Engineering, Tel-Aviv University, Israel. E-mail: {osherazulay, alonmizrahi2, nimrodcurtis}@mail.tau.ac.il; sintov1@tauex.tau.ac.il.

This work was supported by the Pazy Foundation (grant No. 283-20).

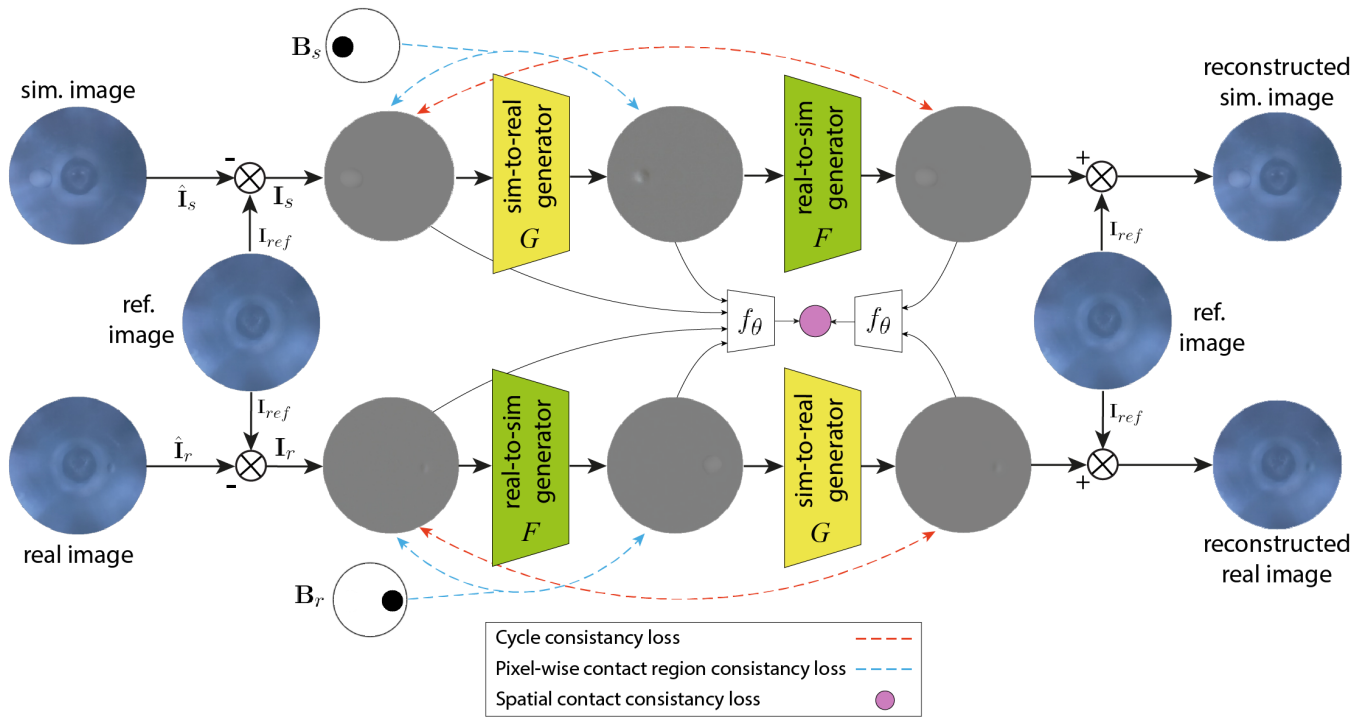


Fig. 2: Scheme of the SightGAN model. The model operates on difference images in order to enhance generability to new sensors. Top and bottom rows illustrate the sim-to-real-to-sim and real-to-sim-to-real processes, respectively. The cycle consistency loss of CycleGAN is augmented by two additional losses aimed to provide pixel-level domain adaptation of the contacts.

perturbations in order to train a GelSight sensor model [19]. Domain randomization, however, requires careful choices of the parameters to randomize and is limited in complex tasks such as in high-resolution tactile images. Some work have harnessed Finite Element Methods (FEM) in order to generate simulated deformation of the contact pad in tactile sim-to-real [20]. However, the computational complexity of such an approach limits real-time sensing [21].

A notable technique involves the use of Generative Adversarial Networks (GANs), a class of deep learning models that can generate new image samples from learned data distributions [22]. In tactile real-to-sim, a GAN was used to match real images of the BioTac sensor in order to target simulated ones [2]. Yet, the approach requires paired images which are difficult to extract. Hence, a variant of GAN termed CycleGAN [23] has gained traction for its ability to facilitate sim-to-real and real-to-sim transfer of tactile information without a paired dataset. CycleGAN is able to learn the mapping of data distributions from the simulated source domain to the real-world target domain and vice versa, without explicit correspondence between individual samples. In the first work to utilize CycleGAN in tactile sensing, a bidirectional sim-to-real approach was proposed for the GelSight sensor [24]. In a later work, an improved CycleGAN architecture was introduced with task-specific loss functions for enhanced structural fidelity of generated tactile images [25]. However, these approaches and others [26], [27] focused on a specific tactile sensor with a flat contact surface and without exhibiting zero-shot capability. In addition, prior work focused on tactile images where the contact trace seen in the image is rather large

[1]. Consequently, CycleGAN may learn to hide information in the adapted image instead of explicitly retaining the semantics when the trace is small [28], [29].

In this paper, we tackle the sim-to-real problem for high-resolution 3D round sensors while enabling zero-shot inference of accurate contact position estimation (Figure 1). Building upon CycleGAN, we propose the *SightGAN* model which augments CycleGAN with contact-specific consistency losses as illustrated in Figure 2. The losses reduce background disparities between simulated and real tactile images while minimizing contact position errors. For the latter, distillation with a trained contact position estimator is exerted to compare accuracies between generated tactile images, either in real or simulated domains. One of the key advantages of employing SightGAN is its bidirectional capability, allowing to seamlessly transfer knowledge both from real to simulated domains and vice versa. This versatility enables to train models for various sensors of different illumination and fabrication uncertainties.

SightGAN is evaluated on the novel AllSight sensor [13] whereas it can potentially be applied to any optical-based tactile sensor. AllSight is a high-resolution and all-round tactile sensor in which a model for it requires a sufficient amount of data. Hence, our approach provides a simulated environment with real-like and accurately labeled tactile images as demonstrated in Figure 1. As such, models trained on these synthetic images exhibit a zero-shot inference capability on new, real and untrained sensors. In addition, unlike prior work, SightGAN is able to withstand small contact traces. Hence, the proposed model can be used in applications with small loads in, for instance, in-hand




Fig. 3: (Left) The AllSight tactile sensor with the internal view of the camera during contact with a round indenter. (Right) Structure illustration of the AllSight sensor.

manipulation. The simulator of AllSight along with the proposed SightGAN sim-to-real framework and datasets are provided open-source¹ for the benefit of the community and to advance research in the field.

II. OPTICAL TACTILE SENSORS

A. Design

An optical-based tactile sensor typically uses an internal camera to track the deformation of a soft elastomer upon contact with an object. In this work, we consider the AllSight sensor proposed in [13] which has a full 360° round contact region clearly visible without blind spots or obscurity. Hence, the camera is covered by a tube in the shape of a cylinder with an hemispherical end as seen in Figure 3. The tube is three-layered where the inner layer is a rigid crystal-clear shell. A transparent elastomer covers the shell and is coated on its exterior by a reflective silicone paint. In such design, the camera observes the deformation of the elastomer from within upon contact with the exterior layer. The inner-surface of the shell is evenly illuminated by an annular printed circuit board (PCB) with embedded LEDs. The lighting provides better visibility and informative images. In this work we consider white lighting while any other LED setting is possible.

B. Data collection of real images

A dataset \mathcal{P}_{real} of real images with spatial position labels is collected using an automated setup. A robotic arm with a round indenter mounted on its tip repeatedly touched the surface of AllSight in various contact locations, as can be seen in Figure 4. First, a reference image $\mathbf{I}_{ref} \in \mathcal{I}_r$, where \mathcal{I}_r denotes the space of real images, is recorded for a sensor without any contact. During contact, an image $\hat{\mathbf{I}}_i$ is taken along with its position \mathbf{p}_i on the contact surface. The position is represented in the sensors




Fig. 4: Data collection system including a robotic arm and an F/T sensor with several indenters.

Fig. 4: Data collection system including a robotic arm and an F/T sensor with several indenters.

¹AllSight simulator, SightGAN sim-to-real framework and datasets: <https://github.com/osheraz/allightsim>

coordinate frame as seen in Figure 1 and calculated through the forward kinematics of the arm. Furthermore, we consider difference images in the dataset such that an image used for training is $\mathbf{I}_i = \hat{\mathbf{I}}_i - \mathbf{I}_{ref}$. Such subtraction has been shown to enhance the learning process making the model agnostic to the background and focuses only on the color gradients that occur around the deformations [5]. Consequently, the model is expected to generalize to new sensors of different backgrounds in zero-shot. The acquisition and labeling process yields dataset $\mathcal{P}_{real} = \{(\mathbf{I}_i, \mathbf{p}_i)\}_{i=1}^N$ of N labeled images.

C. Data collection in simulation

TACTO is a physics-engine simulator for optical-based tactile sensors [14]. An AllSight simulation was set in TACTO as seen in Figure 1 and calibrated by including reference images from real AllSight sensors. To enhance sim-to-real pre-training of the state estimation model, we collected different reference images from different AllSight sensors and used them for augmentation. The simulated dataset \mathcal{P}_{sim} was generated by labeling M images captured in TACTO during random contacts. Here also, a reference image from a real sensor is subtracted from the contact images. An image $\mathbf{I}_i \in \mathcal{I}_s$, where \mathcal{I}_s denotes the space of simulated images, is taken along with the contact position \mathbf{p}_i such that $\mathcal{P}_{sim} = \{(\mathbf{I}_i, \mathbf{p}_i)\}_{i=1}^M$.

D. Contact position estimation model

Using either datasets \mathcal{P}_{real} or \mathcal{P}_{sim} , we train a contact position estimation model. The model $f_\theta : \mathcal{I} \rightarrow \mathbb{R}^3$ maps a tactile image to the spatial position of the contact on the sensor $\mathbf{p} \in \mathbb{R}^3$. Vector θ is the trainable parameters of the model. The model is based on the ResNet-18 architecture [30]. The top layer is removed and the flattened output features are fed through two fully-connected layers of size 512 and 256. At each iteration, both reference \mathbf{I}_{ref} and contact \mathbf{I}_i images are down-sampled to resolution 224×224 and stacked along the channel. The stacked image is then passed through the model to get the estimated position $\hat{\mathbf{p}}_i$.

III. METHOD

The proposed SightGAN, illustrated in Figure 2, integrates the CycleGAN architecture with additional auxiliary losses designated for tactile images. The losses aim to reduce disparities in background and contact reconstruction of images in the bidirectional transfer between simulation and real domains. In this section, we briefly outline the GAN and CycleGAN losses prior to presenting the SightGAN loss.

A. Generative Adversarial Network (GAN)

In GAN, a mapping $G : \mathcal{I}_s \rightarrow \mathcal{I}_r$ is trained such that a discriminator D_r cannot distinguish between an original image $\mathbf{I}_r \in \mathcal{I}_r$ and a synthetic one $\tilde{\mathbf{I}}_r = G(\mathbf{I}_s)$ where $\mathbf{I}_s \in \mathcal{I}_s$. Model G is trained to minimize the adversarial loss

$$\mathcal{L}_{GAN}(G, D_r, \mathcal{I}_s, \mathcal{I}_r) = \mathbb{E}_{\mathbf{I}_r \sim \mu_r} [\log(D_r(\mathbf{I}_r))] + \mathbb{E}_{\mathbf{I}_s \sim \mu_s} [\log(1 - D_r(G(\mathbf{I}_s)))] \quad (1)$$

where μ_r and μ_s are the data distributions of \mathcal{I}_r and \mathcal{I}_s , respectively.

B. CycleGAN: Bidirectional Image Translation

Our approach facilitates bidirectional mapping, as in CycleGAN, between unpaired image datasets from two domains, \mathcal{I}_s and \mathcal{I}_r . This process involves the generators of $G : \mathcal{I}_s \rightarrow \mathcal{I}_r$ and $F : \mathcal{I}_r \rightarrow \mathcal{I}_s$ along with adversarial discriminators D_s and D_r . In addition to adversarial losses for both mappings, a cycle consistency loss is included in order to minimize the cyclic reconstruction of images given by

$$\mathcal{L}_{\text{cycle}}(G, D_r, \mathcal{I}_s, \mathcal{I}_r) = \mathbb{E}_{\mathbf{I}_r \sim \mu_r} \|G(F(\mathbf{I}_r)) - \mathbf{I}_r\|_1 \quad (2)$$

$$+ \mathbb{E}_{\mathbf{I}_s \sim \mu_s} \|F(G(\mathbf{I}_s)) - \mathbf{I}_s\|_1.$$

Consequently, CycleGAN is trained with the following loss:

$$\mathcal{L}_{\text{CycleGAN}}(G, F, D_s, D_r) = \mathcal{L}_{\text{GAN}}(G, D_r, \mathcal{I}_s, \mathcal{I}_r) \quad (3)$$

$$+ \mathcal{L}_{\text{GAN}}(F, D_s, \mathcal{I}_r, \mathcal{I}_s)$$

$$+ \lambda_{\text{cycle}} \mathcal{L}_{\text{cycle}}(F, G)$$

where λ_{cycle} is some pre-defined weight.

C. SightGAN Losses

Drawing inspiration from the perception consistency loss of RetinaGAN [31], we introduce two auxiliary contact losses, one in the image domain and the other in the contact space. In the context of optical tactile images, these images can be partitioned into two distinct regions: the background region representing the no-contact area of the tactile sensor and the foreground region embodying the tactile sensor's interaction with objects. For the background region, the *Pixel-wise Contact Region Consistency loss* emphasizes the constraint of color similarity in the generated images while forcing no contact traces in the region. Furthermore, the *Spatial Contact Consistency Loss* focuses on the contact estimation accuracy as minor deviations in the generated image can have great significance on the spatial positioning of the contact. This meticulous attention to foreground texture and structural nuances of the background is pivotal in optimizing sim-to-real tactile images from 3D tactile sensors.

1) *Spatial Contact Consistency loss*: The spatial contact consistency loss penalizes disparities in contact localization after transferring images across domains. We define function \mathcal{L}_{sp} to compare between contact position estimations of two images, \mathbf{I} and \mathbf{J} as

$$\mathcal{L}_{sp}(\mathbf{I}, \mathbf{J}) = \|f_{\theta}(\mathbf{I}) - f_{\theta}(\mathbf{J})\|^2. \quad (4)$$

With this function and based on the perception consistency loss of the RetinaGAN [31], we define the spatial contact consistency loss as

$$\mathcal{L}_{\text{spatial}}(\mathbf{I}_s, \mathbf{I}_r, F, G) = \mathcal{L}_{sp}(\mathbf{I}_s, G(\mathbf{I}_s)) + \quad (5)$$

$$+ \frac{1}{2} \mathcal{L}_{sp}(\mathbf{I}_s, F(G(\mathbf{I}_s))) + \frac{1}{2} \mathcal{L}_{sp}(G(\mathbf{I}_s), F(G(\mathbf{I}_s))) +$$

$$+ \mathcal{L}_{sp}(\mathbf{I}_r, F(\mathbf{I}_r)) + \frac{1}{2} \mathcal{L}_{sp}(\mathbf{I}_r, G(F(\mathbf{I}_r))) +$$

$$+ \frac{1}{2} \mathcal{L}_{sp}(F(\mathbf{I}_r), G(F(\mathbf{I}_r))).$$

The halving of losses involving the cycled images accounts for their dual comparison against the original and transferred images.

2) *Pixel-wise Contact Region Consistency loss*: In order to further augment the the accuracy of contact localization in domain transfer and enhance structural fidelity, we introduce a loss related to the contact region. Minor changes in contact pixels can lead to significant alterations in the contact domain, potentially causing errors in contact estimation with f_{θ} . For either simulated or real training images, the contact position \mathbf{p} is labeled as described in Section II. For each image \mathbf{I} , binary image \mathbf{B} is defined where a mask is placed on the contact region of the image. In practice, pixels in and out of the contact regions are marked with zeros and ones, respectively, in \mathbf{B} . The contact loss between an image and its transfer is, therefore, defined by

$$\mathcal{L}_m(\mathbf{I}, \mathbf{B}, H) = \|\mathbf{I} * \mathbf{B} - H(\mathbf{I}) * \mathbf{B}\|_1 \quad (6)$$

where $*$ denotes pixel-wise multiplication and, H is either F or G for real or simulated images, respectively. Hence, the pixel-wise contact region consistency loss is defined as

$$\mathcal{L}_{\text{mask}}(\mathbf{I}_s, \mathbf{I}_r, F, G) = \mathcal{L}_m(\mathbf{I}_s, \mathbf{B}_s, G) + \mathcal{L}_m(\mathbf{I}_r, \mathbf{B}_r, F). \quad (7)$$

Note that additional loss components in (7) similar to $\mathcal{L}_{\text{spatial}}$ do not enhance performance as resulted in preliminary analysis and, therefore, not included.

The above two auxiliary losses are applied across batches of simulated and real images with sim-to-real G and real-to-sim F generators. Hence, the overall auxiliary loss is defined by the sum of both

$$\mathcal{L}_c(\mathbf{I}_s, \mathbf{I}_r, F, G) = \lambda_{\text{spatial}} \mathcal{L}_{\text{spatial}}(\mathbf{I}_s, \mathbf{I}_r, F, G) \quad (8)$$

$$+ \lambda_{\text{mask}} \mathcal{L}_{\text{mask}}(\mathbf{I}_s, \mathbf{I}_r, F, G)$$

where λ_{spatial} and λ_{mask} are weight parameters. Combining (3) with (8), the complete SightGAN loss is

$$\mathcal{L}_{\text{SightGAN}}(G, F, D_s, D_r) = \mathcal{L}_{\text{GAN}}(G, D_r, \mathcal{I}_s, \mathcal{I}_r)$$

$$+ \mathcal{L}_{\text{GAN}}(F, D_s, \mathcal{I}_r, \mathcal{I}_s)$$

$$+ \lambda_{\text{cycle}} \mathcal{L}_{\text{cycle}}(F, G)$$

$$+ \mathcal{L}_c(\mathbf{I}_s, \mathbf{I}_r, F, G). \quad (9)$$

Once the SightGAN model has been trained, a simulator is available. In order to generate real-like images for a new sensor, one only needs to reintegrate the generated foreground image outputted from the sim-to-real generator with a real reference image of the sensor. Hence, the generated images are expected to provide zero-shot contact inference.

IV. EXPERIMENTS

Experiments are presented in this section in order to analyze the performance of SightGAN. While prior work often use one specific sensor, we collect a training dataset $\mathcal{P}_{\text{real}}$ over six different AllSight sensors. The dataset is collected in an automated process where each AllSight sensor is mounted on a fixed frame. An indenter with a round tip of radius 3 mm is mounted to a robotic arm. The arm is also equipped with a Robotiq FT-300 Force/Torque (F/T) sensor for labeling contact forces. We allocate 2% of all samples randomly to avoid contact with the surface. For each sample, the robot selects a point to press on the surface of the sensor and computes its location through its kinematics. An image taken during the press is labeled with its position. The collection yielded 1,000 labeled images for each sensor and

a total of $N = 6,000$ labeled images in \mathcal{P}_{real} . In addition, we aim to analyze the generalization abilities of the trained model for new sensors. Hence, a test set is collected over two new sensors not included in the training. The test set is comprised of 2,000 labeled images. Results are presented with white LED illumination for all sensors whereas RGB illumination was tested in preliminary experiments yielding similar results. Also, for each image, the reference image is subtracted as described in Section II-B.

For \mathcal{P}_{sim} , data is collected in TACTO as described in Section II-C. \mathcal{P}_{sim} consists of simulated tactile images and their associated contact positions, featuring a 3 mm radius spherical indenter at varying penetration depths. To fine-tune the simulation, we utilized reference images obtained from six different sensors. The collection yielded 1,000 labeled images for each sensor and a total of $M = 6,000$ labeled images in \mathcal{P}_{sim} .

A. Contact position evaluation

Table I summarizes the Root-Mean-Square-Errors (RMSE) for position estimation with model f_θ while trained with different origins of training data. All models were evaluated on distinct data from the two test sensors. The results include the lower accuracy bound of directly training with real data (different from the test data) from the two test sensors. Also, we include accuracy when training f_θ directly with \mathcal{P}_{real} . Next, model f_θ is trained with data \mathcal{P}_{sim} generated in the simulation without any GAN and while using the reference images of the six training sensors. Using only simulated data provides poor accuracy showing that the simulation, even with real reference images, is far from representing reality. Then, f_θ with data generated by the CycleGAN (trained with $\lambda_{cycle} = 10$) alone without additional losses is evaluated. The error with only CycleGAN is the highest due to its inability to focus and reconstruct the contact. This is a known issue when most of the image is not under contact and can even lead to a mode of collapse [28]. Adding either spatial contact consistency loss (5) with $\lambda_{spatial} = 0.1$ or pixel-wise contact region consistency loss (7) with $\lambda_{mask} = 30$ is shown to significantly reduce the error. The accuracy of SightGAN, which combines CycleGAN with both losses, provides the lowest error out of the generative models. Hence, SightGAN generates real-like images from simulated ones and enables zero-shot position estimation of contacts.

We now evaluate the accuracy which the sim-to-real of SightGAN provides to images generated from simulation considering diversity in the training data. Model f_θ is trained over synthetic images generated by the sim-to-real of SightGAN. Figure 5 presents the position estimation RMSE of f_θ over the test sensors with regards to the number of real train sensors used to train SightGAN. The addition of more sensors in the training set increases diversity and decreases the estimation error. Hence, more train sensors improves the zero-shot capability over new real sensors. While zero-shot provides relatively accurate predictions, real data from the target sensors can improve accuracy. Figure 6 shows the error of position estimation over the test data of the two

TABLE I: Estimation accuracy of contact positions over data from two test sensors with regards to the origin of the train data for f_θ

Origin of training data for f_θ		Position RMSE (mm)
Direct	Data from the 2 test sensors	0.66
	Data from 6 train sensors	2.16
	6 sensors from simulation	7.48
Generative	CycleGAN	13.30
	CycleGAN + spatial contact loss (5)	3.86
	CycleGAN + pixel-wise contact loss (7)	3.70
	SightGAN	3.49

new sensors with regards to the number of new samples used to fine-tune the model. The addition of a small amount of new samples for fine-tuning further improves accuracy. With 300 additional samples, the position RMSE reaches to approximately 1 mm.

Next, we evaluate the generability of a model, trained on samples with round indenters, to estimate contacts of other geometries. Test data of 2,000 labeled images was collected from one test sensor with square (edge length 6 mm) and elliptical (axis lengths 8 mm and 4 mm) indenters. Table II presents the RMSE results for position contact estimation on the test data. The first row shows a baseline for training and testing model f_θ with the new test indenters. All other results in the Table are based on training data of round indenters and evaluation on the test indenters. While the second and third rows are estimation models trained on real data, the rows below are based on estimation models trained with simulated images taken from the simulator while using similar square and elliptical indenters. The simulated images are passed through the respective sim-to-real model and used to train f_θ . Here also, SightGAN exhibits the lowest error over all sim-to-real generated image origins. Furthermore, the results show that using SightGAN generalizes well and shows similar accuracy to directly training with real data. Hence, synthetic data from the sim-to-real generator of SightGAN generalizes well and can be used to train zero-shot estimation models.




Fig. 5: Position estimation error of f_θ trained with synthetic data from SightGAN over the test sensor with regards to the number of train sensors used to train SightGAN.

B. Sim-to-real quality

In order to evaluate the sim-to-real of SightGAN for generating realistic tactile images, we consider the Fréchet Inception Distance (FID) [32] and Kernel Inception Distance (KID) [33] metrics. FID and KID are common pixel-level metrics of domain adaptation and quantify the quality and diversity of generated images. Furthermore, they do not




Fig. 6: Position estimation error with regards to the number of real images from the test sensor used to fine-tune model f_θ . Results with zero new tactile images are the zero-shot transfer errors without any fine-tuning.

TABLE II: Estimation accuracy of contact positions over data from one test sensor with indenters of different geometries

Origin of training data for f_θ		Position RMSE (mm)
Direct	Train sensors w/ square and ellipse indenters	0.94
	Train sensors	2.67
	Test sensor	3.19
	6 simulated sensors	7.74
Generative	CycleGAN	13.28
	CycleGAN + spatial contact loss (5)	3.93
	CycleGAN + pixel-wise contact loss (7)	3.97
	SightGAN	3.67

require paired images. Table III exhibits FID and KID comparative results for direct mapping between simulated and real images, for CycleGAN and for SightGAN. Direct comparison between simulated and real images is obviously poor and emphasizes the reality gap. CycleGAN provides some improvement. Nevertheless, SightGAN is shown to be superior to CycleGAN with approximately 47% and 16% quality improvements for FID and KID, respectively.

C. Contact force evaluation

Common tactile simulations provide force readings which are linear to the deformation of the contact pad [14], [34]. In reality, however, modeling the deflection of the contact pad is nearly impossible and the reaction forces are more complex [20], [35]. Hence, we further assess the potential of using SightGAN to provide accurate force estimations to the real-like simulated tactile images. A trained contact force estimation model $g_\psi : \mathcal{I} \rightarrow \mathbb{R}^3$ is evaluated to approximate the corresponding contact force $\mathbf{f} = [f_x, f_y, f_z]^T \in \mathbb{R}^3$. Vector ψ includes trainable model parameters. A real dataset was collected with the F/T sensor by labeling tactile images of the test sensors with force magnitudes during contact. The dataset includes a wide range of contact force magnitudes within ranges $f_z \in [0N, 12N]$ and $f_x, f_y \in [-2N, 2N]$ where f_z is the normal force and f_x, f_y are the forces tangential to the contact surface at \mathbf{p} .

TABLE III: FID and KID comparison over the generated tactile images.

Model		FID	KID
Direct	\mathbf{I}_r vs. \mathbf{I}_s	98.54±0.90	29.19±0.45
CycleGAN	\mathbf{I}_r vs. $G(\mathbf{I}_s)$	73.63±0.51	13.27±0.10
SightGAN	\mathbf{I}_r vs. $G(\mathbf{I}_s)$	39.12±0.30	11.24±0.12

TABLE IV: Estimation accuracy of contact forces with regards to the origin of the train data for g_ψ

Origin of train data for g_ψ	Test data ¹	Force RMSE (N)
6 train sensors	2 train sensors	0.30
2 test sensors	2 test sensors	0.78
SightGAN real-to-sim-to-real	2 train sensors	0.42
SightGAN real-to-sim-to-real	2 test sensors	0.81

¹Train and test data are distinct in all cases.

Table IV presents a comparison of force estimation accuracy over several origins of train data for g_ψ . For a baseline, model g_ψ is trained on data from the six train sensors and evaluated on different data from the same sensors. Similarly, the model was trained on data from the two test sensors and evaluated on different data from the same sensors. Obviously, having more diversity originating from more sensors provides better generability and lower errors. Next, SightGAN and g_ψ are trained on data from the six training sensors. Following, force accuracy is evaluated by passing new test data, either from the six train sensors or the two test sensors, through the real-to-sim and sim-to-real generators. The reconstructed images are then evaluated with model g_ψ .

The results with SightGAN show relatively similar accuracy to the baseline errors. Evidently, SightGAN maintains force information embedded within the original images and reconstructs it. Potentially, simulated images from the simulator upon contact can be mapped with the sim-to-real to the real domain and evaluated with g_ψ . Hence, the simulator would be able to provide a comprehensive contact state in various applications. However, such an approach requires further research left to future work.

V. CONCLUSIONS

In this paper, a novel generative model termed SightGAN was proposed for augmenting tactile simulators with real-like capabilities. The model is based on a bidirectional GAN and can be used to simulate 3D round tactile sensors. Using unpaired real and simulated data, SightGAN is trained to provide sim-to-real and real-to-sim mappings while minimizing the cycle consistency loss of CycleGAN and two additional loss components. The two loss components aim to minimize contact localization errors in the reconstruction of the image through distillation with a trained contact position estimator. Such approach augments a simulator with real-like performance and enables zero-shot capabilities for models trained on the simulated sensors and evaluated on real ones. Indeed, using real data to train a contact position estimator provides similar or better accuracy than doing so with synthetic data generated by the SightGAN. However, the proposed model provides a framework for augmenting a simulator with real-like capabilities and can leverage tedious policy learning processes. SightGAN was shown to embed contact force data within the generated images. Hence, future work could address the calibration of simulated linear forces to real-world contact forces. Future work may also include investigating the integration of SightGAN based policy learning into robotic in-hand manipulation, addressing complex real-world challenges and constraints of online learning.

REFERENCES

- [1] J. Xu, S. Kim, T. Chen, A. R. Garcia, P. Agrawal, W. Matusik, and S. Sueda, "Efficient tactile simulation with differentiability for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 1488–1498.
- [2] A. Church, J. Lloyd, N. F. Lepora *et al.*, "Tactile sim-to-real policy transfer via real-to-sim image translation," in *Conference on Robot Learning*. PMLR, 2022, pp. 1645–1654.
- [3] S. Suresh, Z. Si, S. Anderson, M. Kaess, and M. Mukadam, "Midastouch: Monte-carlo inference over distributions across sliding touch," in *Conference on Robot Learning*. PMLR, 2023, pp. 319–331.
- [4] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez, "Tactile-rl for insertion: Generalization to objects of unknown geometry," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6437–6443, 2021.
- [5] C. Higuera, B. Boots, and M. Mukadam, "Learning to read braille: Bridging the tactile reality gap with diffusion models," *arXiv preprint arXiv:2304.01182*, 2023.
- [6] D. F. Gomes, S. Luo, and P. Paoletti, "Beyond flat gelsight sensors: Simulation of optical tactile sensors of complex morphologies for sim2real learning," in *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, K. E. Bekris, K. Hauser, S. L. Herbert, and J. Yu, Eds., 2023.
- [7] A. Maslyczyk, J.-P. Roberge, V. Duchaine *et al.*, "A highly sensitive multimodal capacitive tactile sensor," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 407–412.
- [8] O. Azulay, I. Ben-David, and A. Sintov, "Learning haptic-based object pose estimation for in-hand manipulation control with underactuated robotic hands," *IEEE Transactions on Haptics*, pp. 1–12, 2022.
- [9] M.-Y. Cheng, C.-M. Tsao, Y.-T. Lai, and Y.-J. Yang, "A novel highly-twistable tactile sensing array using extendable spiral electrodes," in *IEEE Int. Conf. on Micro Electro Mech. Sys.*, 2009, pp. 92–95.
- [10] I. H. Taylor, S. Dong, and A. Rodriguez, "Gelslim 3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger," in *IEEE Int. Conf. on Rob. & Auto.*, 2022, pp. 10781–10787.
- [11] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer *et al.*, "Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [12] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.
- [13] O. Azulay, N. Curtis, R. Sokolovsky, G. Levitski, D. Slomovik, G. Lilling, and A. Sintov, "Allsight: A low-cost and high-resolution round tactile sensor with zero-shot learning capability," *arXiv preprint arXiv:2307.02928*, 2023.
- [14] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra, "Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3930–3937, 2022.
- [15] Z. Si and W. Yuan, "Taxim: An example-based simulation model for gelsight tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2361–2368, 2022.
- [16] K. Patel, S. Iba, and N. Jamali, "Deep tactile experience: Estimating tactile sensor output from depth sensor data," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9846–9853.
- [17] Z. Si, Z. Zhu, A. Agarwal, S. Anderson, and W. Yuan, "Grasp stability prediction with sim-to-real transfer from tactile sensing," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 7809–7816.
- [18] Z. Ding, N. Lepora, and E. Johns, "Sim-to-real transfer for optical tactile sensing," in *IEEE International Conference on Robotics and Automation (ICRA)*, 05 2020, pp. 1639–1645.
- [19] D. Fernandes gomes, P. Paoletti, and S. Luo, "Generation of gelsight tactile images for sim2real learning," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 03 2021.
- [20] Y. S. Narang, B. Sundaralingam, K. V. Wyk, A. Mousavian, and D. Fox, "Interpreting and predicting tactile signals for the syntouch biotac," *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1467–1487, 2021.
- [21] C. Sferrazza and R. D'Andrea, "Sim-to-real for high-resolution optical tactile sensing: From images to three-dimensional contact force distributions," *Soft Robotics*, vol. 9, no. 5, pp. 926–937, 2022.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [23] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [24] W. Chen, Y. Xu, Z. Chen, P. Zeng, R. Dang, R. Chen, and J. Xu, "Bidirectional sim-to-real transfer for gelsight tactile sensors with cyclegan," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6187–6194, 2022.
- [25] X. Jing, K. Qian, T. Jianu, and S. Luo, "Unsupervised adversarial domain adaptation for sim-to-real transfer of tactile images," *IEEE Transactions on Instrumentation and Measurement*, 2023.
- [26] Y. Zhao, X. Jing, K. Qian, D. F. Gomes, and S. Luo, "Skill generalization of tubular object manipulation with tactile sensing and sim2real learning," *Robotics and Autonomous Systems*, vol. 160, p. 104321, 2023.
- [27] W. D. Kim, S. Yang, W. Kim, J.-J. Kim, C.-H. Kim, and J. Kim, "Marker-embedded tactile image generation via generative adversarial networks," *IEEE Robotics and Automation Letters*, 2023.
- [28] K. Rao, C. Harris, A. Irpan, S. Levine, J. Ibarz, and M. Khansari, "Rl-cyclegan: Reinforcement learning aware simulation-to-real," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11157–11166.
- [29] C. Chu, A. Zhmoginov, and M. Sandler, "Cyclegan, a master of steganography," *arXiv preprint arXiv:1712.02950*, 2017.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [31] D. Ho, K. Rao, Z. Xu, E. Jang, M. Khansari, and Y. Bai, "Retinagan: An object-aware approach to sim-to-real transfer," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 10920–10926.
- [32] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," in *International Conference on Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 6629–6640.
- [33] M. Binkowski, D. J. Sutherland, M. Arbel, and A. Gretton, "Demystifying MMD gans," in *International Conference on Learning Representations*, 2018.
- [34] Z. Si and W. Yuan, "Taxim: An example-based simulation model for gelsight tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2361–2368, 2022.
- [35] P. Ruppel, Y. Jonetzko, M. Görner, N. Hendrich, and J. Zhang, "Simulation of the syntouch biotac sensor," in *Intelligent Autonomous Systems*, M. Strand, R. Dillmann, E. Menegatti, and S. Ghidoni, Eds., 2019, pp. 374–387.