

DexFruit: Dexterous Manipulation and Gaussian Splatting Inspection of Fruit

Aiden Swann^{1*}, Alex Qiu^{1*}, Matthew Strong², Angelina Zhang¹, Samuel Morstein¹,
Kai Rayle¹, Monroe Kennedy III^{1,2}

Abstract—Dexfruit is a robotic manipulation framework that enables gentle, autonomous handling of fragile fruit and precise evaluation of damage. Soft fruits have long faced an issue of produce loss in both the harvesting and post-harvesting processes due to their extreme fragility and susceptibility to bruising, making them one of the hardest produce type to manipulate with automation. In this work, we demonstrate by using optical tactile sensing, autonomous manipulation of fruit with minimal damage can be achieved. We show that our tactile informed diffusion policies outperform baselines in both reduced bruising and pick-and-place success rate across three fruits: strawberries, tomatoes, and blackberries. In addition, we introduce FruitSplat, a novel technique to represent and quantify visual damage in a high-resolution 3D representation via 3D Gaussian Splatting (3DGS). Existing metrics for measuring damage lack quantitative rigor or require expensive equipment. With FruitSplat, we distill a 2D fruit mask as well as a 2D bruise segmentation mask into the 3DGS representation from just a web-cam video. Furthermore, this representation is modular and general, compatible with any relevant 2D model. Overall, we demonstrate a 92% grasping policy success rate, up to a 15% reduction in visual bruising, and up to a 31% improvement in grasp success rate on challenging fruit compared to our baselines across our three tested fruits. We rigorously evaluate this result with over 630 trials. Please checkout our website, which contains our code and datasets at <https://dex-fruit.github.io/>.

Index Terms—Force and Tactile Sensing, Agricultural Automation, Dexterous Manipulation

I. INTRODUCTION

DEMOGRAPHIC-DRIVEN labor shortages and a growing population present significant challenges for the global food supply [1–3]. To address these impending issues, the agricultural industry has taken many strides into increased applications of machinery and automation [4, 5]. While some agricultural production is highly mechanized, such as grain and cereal, automation for fruits and vegetables is comparatively limited, often requiring workers to manually harvest produce [1, 6]. Fruit production, in particular, requires a high level of labor-intensive operations to be performed, primarily due to the delicate handling required to preserve quality [6]. Fruits, such as strawberries, demand significant human labor because

Manuscript received: August 8th, 2025; Revised: November 2nd, 2025; Accepted: December 3rd, 2025.

This paper was recommended for publication by Editor Júlia Borràs Sol upon evaluation of the Associate Editor and Reviewers’ comments.

A. Swann and M. Strong are supported by NSF GRFP Fellowship No. DGE-2146755. This work was also supported in part by NSF Grant No. 2142773 and 2220867. *Both authors contributed equally to this work.

[¹] are with the Department of Mechanical Engineering, [²] Department of Computer Science, Stanford University, Stanford CA, USA. Emails: {swann, aqiu34, mastro1, azhang11, morstein, kairayle, monroek}@stanford.edu

©2026 IEEE




Fig. 1: Our framework, DexFruit, enables safe manipulation of fragile fruit using optical tactile feedback. We further present FruitSplat, a method for accurate 3D reconstruction, automated segmentation, and bruise localization in strawberries.

of their fragile structure and short post-harvest shelf life; they are typically hand-picked to minimize bruising and decay [6]. Although there are some mechanized solutions for fruit harvesting, they often do not meet the quality standards of the fresh market, damage fruit more than their human counterparts [7], and still involve significant human labor in sorting. Soft fruits in particular are especially challenging for robotic systems, as they require extra care and precision in harvesting and manipulation not required for other produce [8–10], presenting a clear need for dexterous robotic manipulation.

Beyond harvesting, the industry has also turned to robotics and automation for the post-harvest process, which includes sorting, quality control, packing, and shipping. These processes incur substantial losses to farmers due to post-harvest damage, accounting for almost a third of all of the world’s produce [11]. This significant loss of food compounded with high process inefficiencies and labor costs have specifically driven the industry to integrate robots into this post-harvesting downstream pipeline, applying soft-grippers and various computer vision techniques for manipulation and damage assessment

[12]. With automation introduced into the post-harvest process, pick-and-place operations have become the most repetitive robotic manipulation task [13]. In order for these pick-and-place tasks to be successful, the robot must maintain a stable but soft grasp on the target produce to ensure minimal damage. Recent advances in imitation learning have allowed for tasks like pick-and-place to be executed with ease, yet most of these methods employ rigid grippers and target rigid objects, neglecting damage mitigation. Moreover, these methods do not address the question of post-manipulation damage, a crucial step that must be tightly coupled with automated manipulation. As of now, there is a lack of low-cost, automated methods for assessing damage quantitatively for quality control. Practical damage metrics are a pre-requisite to real-world robot policy benchmarking and will facilitate continued development of soft fruit manipulation.

To address these gaps, we introduce **DexFruit**, a novel framework that integrates tactile-sensing into imitation learning for careful manipulation of soft fruit and present a stand-alone post-manipulation damage analysis method **FruitSplat**, which offers a quantitative, 3D solution for bruise detection. Our method employs DenseTact sensors [14] to manipulate strawberries, cherry tomatoes, and blackberries with minimal bruising or damage. We implement Diffusion Policy with tactile feedback from DenseTact sensors, which are mounted on simple general-purpose parallel jaw-grippers. FruitSplat is a 3D Gaussian Splatting method that integrates 2D bruising information from a fine-tuned YOLO [15] segmentation model into a 3D representation, providing both a qualitative and quantitative measurement of post-manipulation damage. Using only a simple webcam, FruitSplat bridges the gap between high-fidelity scanning techniques and traditional deep learning methods. Our framework aims to address the crucial need for dexterous manipulation with soft fruit using state of the art imitation learning methods infused with tactile feedback, as well as providing a simple but effective approach to damage assessment for post-harvest quality control.

Our key contributions are as follows:

- We demonstrate the integration of optical tactile sensing into diffusion-based imitation learning policies for soft body manipulation.
- We demonstrate quantitative and qualitative reductions in bruising and improvements in success rate compared to our baselines while using tactile sensing. We show fruit remains damage free after up to 50 repeated grasps.
- We introduce FruitSplat, a novel method for qualitative and quantitative damage analysis for strawberries.

II. RELATED WORK

A. Robotic Fruit Manipulation

Robotic fruit manipulation is still relatively underexplored in literature, with most applications targeting harvesting as opposed to the post-harvest processes. Previous works include robotic fruit picking of apples, strawberries, blackberries, and

peppers [16–20]. These works typically involve the construction of a novel soft gripper, which is unique to a specific fruit or berry. In [16] they focus on apples, and [17, 19] on blackberries. While these grippers generally perform well, they are limited to a single type of fruit. They typically prevent damage through modality specific designs [18, 19, 21] and low dimensional sensing [20] or by not touching the fruit at all [22]. In this work, we focus on a general gripper capable of manipulating arbitrary fruit. Rather than relying solely on soft robotics techniques or complex designs, we utilize optical tactile sensing.

Some recent studies have explored the utility of optical tactile sensing in agricultural applications, particularly for fruit manipulation and evaluation. For instance, [23] employs optical tactile sensing to facilitate fruit manipulation; however, their approach is limited to hard fruits such as apples and lemons, which restricts its applicability across a diverse range of produce. Furthermore, they do not quantify the damage over repeated grasping. In [24] the authors utilize a vision-based tactile sensor, to determine the stiffness of cherry tomatoes. While our current work does not incorporate direct stiffness measurements, the underlying principles of this approach could be integrated into our pipeline. Furthermore, we focus on the manipulation task while preventing damage.

B. Bruise and Metrics

Current approaches to bruise detection largely rely on qualitative assessments, with little emphasis on quantitative metrics. For example, [20] employs a visual classification system for raspberries, classifying them as undamaged, mildly damaged, or damaged, which suffers from subjectivity and limited resolution in damage assessment. As quality control of fruit on the market is commonly conducted through visual inspection, allowable fruit quality varies based on inspector [25]. This system is inherently subjective and is in need of a standardized method to quantify damage.

There are a number of quantitative methods for damage detection for agricultural products. [26] provides a comprehensive survey of high-precision methods currently in use. Techniques such as spectroscopy, optical coherence tomography (OCT), magnetic resonance imaging (MRI), and x-ray imaging offer detailed insights into 3D tissue damage, yet their reliance on expensive and complex setups restricts their practical use in robotics. While these techniques offer high resolution and rich information about the quality of a fruit, this is often unnecessary to determine if a fruit should be brought to market. Multi-spectral imaging has also been used [26], however this can also require specialized hardware. A broad body of work has quantified fruit damage using stiffness-based metrics using direct contact force measurements [27, 28]. While informative, these metrics generally fall short of providing fine-grained analyses.

Deep learning has also been explored to overcome some of these limitations. For example, [29] demonstrated bruise detection on red apples using convolutional models, while [30] addressed strawberry bruise detection aided by UV light. Although these methods enhance detection capabilities compared

to purely visual grading, they do not fully leverage the three-dimensional metric information of the fruit.

C. Imitation Learning and Tactile Sensing

Recent advances have significantly improved the performance of imitation learning through the use of Denoising Diffusion Probabilistic Models (DDPM) [31]. These approaches learn observation conditioned, high-dimensional policies with relatively few demonstrations. In [32] the authors survey the landscape of imitation learning for agriculture. While there are many examples of agricultural tasks like peeling bananas and identifying fruit, none rely on tactile sensing or use Diffusion Policies. Simple tactile sensing through force conditioning is used in [33] to pick up fragile grapes and placing them on a plate. However, DexFruit not only safely grasps diverse fragile fruit using expressive tactile sensing, but also *quantitatively* measures the damage done to the fruit across hundreds of trials providing a rigorous and necessary analysis of the fruit after manipulation.

Recently, several studies have explored incorporating tactile sensing into imitation learning policies, utilizing diverse modalities such as audio [34], low-resolution capacitive sensing [33, 35], and optical tactile sensing [36, 37]. Reactivate Diffusion Policy (RDP) [37] delivers impressive dexterity (e.g., cucumber peeling) via a two-tier tactile scheme, yet reduces tactile input to an optical deformation field. DexSkin [35] completes blueberry manipulation tasks with low-resolution capacitive sensing and reports no rigorous quantitative damage analysis; however, their qualitative examples show visible damage in over 40% of fruit after a single move, whereas in our work we manipulate the same piece of fruit up to 50 times without significant damage. Although these studies demonstrate impressive performance, our work is the first to rigorously validate tactile-feedback-based manipulation of delicate, soft fruits over hundreds of experimental trials.

III. PRELIMINARIES

A. Diffusion Policies

Given an observation \mathcal{O} , which corresponds to external camera images and robot state, and an action A in the space of robot positions, we seek to learn a policy $\pi : \mathcal{O} \rightarrow \mathcal{P}(A)$. Instead of learning this policy directly, we model our policy as a Denoising Diffusion Implicit Model (DDIM) and instead learn a noise predictor network following [31]:

$$\hat{\epsilon}_k = \epsilon_\theta(A_t^k, \mathcal{O}_t, k) \quad (1)$$

Where A_t^k is a sequence of noisy actions, \mathcal{O}_t is the observation and denoising iteration k . Given a dataset $\mathcal{D} = \{\{O_t, A_t\}\}_{t=1}^n$ we train this denoising function by applying noise ϵ^k to data A_t and minimizing the following loss:

$$\mathcal{L} = \text{MSE}\left(\epsilon^k, \epsilon_\theta(A_t^k, \mathcal{O}_t, k)\right) \quad (2)$$

During training we sample random actions A_t^K and perform denoising using a DDIM scheduler [38].

B. 3D Gaussian Splatting

3D Gaussian Splatting (3DGS) is an explicit 3D visual representation for robotics. Given a series of 100-200 images, 3DGS constructs a visually accurate representation that can be rendered at high FPS on a GPU. 3DGS utilizes a set of 3D Gaussian primitives g_i with learnable parameters, including 3D mean position $\mu \in \mathbb{R}^3$, 3D covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, opacity $\alpha \in \mathbb{R}$, and spherical harmonic (SH) coefficients to represent color. Each Gaussian can be defined as:

$$G(x) = e^{-\frac{1}{2}(x)^T \Sigma^{-1}(x)} \quad (3)$$

The 3D scene gets initialized typically from a sparse point cloud, which is computed via Structure-from-Motion (SfM) [39]. Images are rendered by projecting the 3D Gaussians onto the image plane using an affine approximation of projective transformation. The final pixel color C is computed as follows:

$$C = \sum_{i \in N} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) \quad (4)$$

where α_i is given by evaluating a 2D projection of a Gaussian with covariance Σ . The loss used to supervise the Gaussian Splat is a weighted combination of the L1 and Structural Dissimilarity Index Measure (D-SSIM) loss between each ground truth (GT) image and volumetric rendered image.

$$\mathcal{L} = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{D-SSIM} \quad (5)$$

IV. METHOD

DexFruit combines a robust manipulation policy and advanced 3D vision techniques to address the challenges of handling delicate fruits in agricultural robotics. We utilize Diffusion Policy trained on expert demonstrations with optical tactile sensing. To quantify fruit damage, we develop FruitSplat, a novel 3D Gaussian Splatting-based analysis tool that accurately reconstructs and compares fruit models before and after manipulation to quantify damage caused during manipulation.

A. Fruit Manipulation Policy

Safely handling fruit in agricultural robotic manipulation requires policies that balance two critical aspects: robustness, to reliably pick and deposit fruits, and sensitivity, to prevent damage during handling. These factors inherently conflict, as robustness encourages a firm grip, whereas damage prevention demands minimal contact. To meet these conflicting demands, we leverage a Diffusion Policy (see Section III), known for smooth and safe execution in complex real-world scenarios. While many existing approaches achieve gentleness through soft gripper hardware, our standard parallel-jaw gripper equipped with DenseTact implicitly achieves softness via the learned policy, successfully minimizing fruit damage.

For this task, our action space consists of the end-effector pose (3D position + 6D rotation [40]) and the width of the gripper normalized between [0,1]. The observation space consists of one RGB camera mounted on the robot arm wrist and RGB images from two DenseTacts. We perform data

collection and inference at approximately 10hz. To accentuate the changes in the gripper state during contact, the average channel image difference is computed for each sensor. In addition, we include the current robot state as an observation in the same format as the action. The RGB images are down-sampled to (224x224) and encoded into 512-dimensional latent vectors using ResNet-34 [41], while the DenseTact images are processed with ResNet-18 [41]. Both ResNets are trained from scratch and implemented from PyTorch Image Models (TIMM) [42]. We train the model for 1000 epochs using a DDIM noise scheduler [38] with 50 train steps and 16 inference steps. We train with a batch size of 128 and a learning rate of $3e-4$ across 4 L40 GPUs. Expert demonstrations are collected via teleoperation, varying the robot’s starting position and individual fruit to ensure the policy conditions on the fruit itself. To ensure consistency across methods and enable a fair comparison with a purely tactile sensing baseline, we fix the initial location of the strawberry. Since diffusion policies have already shown strong capabilities in identifying and grasping objects from diverse positions, we focus exclusively on gentle grasping rather than localization in this work.

During testing, we observed that raw DenseTact images alone provided insufficient signal strength for the policy to effectively leverage tactile information. This was due to minimal variation in the raw images upon contact, making learning challenging. We compute the image difference relative to the undeformed sensor and find that the resulting grayscale image contains sufficient information. Additionally, we noticed that straightforwardly combining vision and tactile images led to worse performance than using either modality individually, likely caused by the model overfitting to the visual inputs. Inspired by human tactile exploration, we resolved this issue by adaptively switching off visual input when tactile contact occurs, which proved essential for high-performing policies. Contact detection is implemented by computing the sum of squared pixel differences (after converting to grayscale), and triggering the visual–tactile switch when this energy exceeds a predefined threshold τ .

B. FruitSplat

For this work, we chose to implement FruitSplat only for strawberries and discuss the reasons for this decision in Section VII. To quantify damage made to strawberries, we leverage 3D Gaussian Splatting to accurately construct strawberries models before and after manipulation. FruitSplat has the following pipeline:

- 1) *Data Pre-Processing*: Videos of the entire target strawberry are first collected using a web-cam mounted on a turntable. To ensure good coverage of the entire surface the strawberry is suspended on two pieces of fishing line as shown in Fig. 3. These images are then fed directly through COLMAP [43] to obtain sparse points and camera poses of each frame. We then pass each frame into our custom fine-tuned YOLOv12 [44] bruising segmentation model for predicted bruise masks. We fine-tuned the model on two thousand hand-labeled images of store-bought strawberries, with bruise masks manually annotated using open-source data labeling tools. The final




Fig. 2: Here we show the FruitSplat pipeline. Prior to training, the 3D Gaussian Splat the images are processed through 2 models. YOLO segments the bruising, while SAM2 segments the strawberry. These two additional masks are used to supervise additional Gaussian parameters.

segmentation model utilized in experiments achieved 94.2% precision, 89.5% recall, and 84.7% Intersection Over Union (IoU) for our evaluation dataset of 162 images. We also apply SAM 2 [45] to retrieve segmented strawberry masks.

- 2) *3D Gaussian Splatting*: Our custom method directly inherits from Nerfstudio’s Splatfacto [46], however we augment the Gaussian representation by adding additional parameters s_i and b_i to each Gaussian primitive [47]. The parameters $s_i \in [0, 1]$ and $b_i \in [0, 1]$ are probabilities that a Gaussian g_i is part of a strawberry and in a bruised region. Strawberry and bruise mask ground-truth maps can be provided as 2D images for supervision. b_i and s_i are rasterized similarly to color (4) and are supervised using a binary cross entropy loss. Our dataloader utilizes the pre-processed YOLO and SAM2 masks as pseudo-ground truth for the supervised strawberry and bruise mask learning. During training, bruise and strawberry loss terms are simply added to the main loss. Additionally, we gate the bruise losses on the strawberry mask as bruising should only occur on the strawberry and nowhere else in the scene. In order to improve training stability we use a stop-grad in the rasterization of bruise and strawberry masks, which only allows gradients with respect to the s_i and b_i parameters of each Gaussian and does not allow gradients to flow to position and shape parameters of the Gaussians.

- 3) *Post-Processing and Damage Analysis*: Once the Gaussian Splat is fully trained after 10,000 steps, FruitSplat exports the Gaussian means along with the learned Gaussian strawberry and Gaussian bruise parameters. It then self-filters the point cloud based on the strawberry parameter using a fixed probability threshold, leaving a high-resolution point cloud




Fig. 3: On the left (a) is the experimental setup. The two sensing modalities used are highlighted in red: the DenseTact optical tactile sensor and IntelRealsense D435i. On the right (b) is the force measuring apparatus which is used to quantify internal damage in the fruit. (c) Shows that apparatus we utilize for scanning the strawberries.

model of the strawberry. Each point is also exported with an encoded bruise score learned from the previously trained 3DGS, which represents the probability of being a bruise. FruitSplat then utilizes the per-point encoded bruise score to calculate the percentage of bruised points within this filtered strawberry point cloud that pass a preset confidence threshold, and subsequently compares the difference in bruising between pre-manipulation and post-manipulation strawberry. Due to the fact that the point cloud points already represent only the surface of the strawberry, this method can effectively represent visual surface bruising information. It should be noted that damage analysis requires a pre-manipulation and post-manipulation Gaussian Splat of the strawberry as an input; thus, steps 1 and 2 for FruitSplat must be run independently twice. However, both steps only require the use of a camera, with no extra hardware necessary.

C. Internal Stiffness Metric

As discussed in Section II fruit stiffness can be correlated with post harvest damage [27, 28]. Furthermore during testing, we found that there were often instances where the internal structure of the fruit was significantly compromised but did not show observable visual changes. In order to measure this internal damage, we assess fruit damage before and after experimentation by measuring the fruits ability to resist an external force. We developed a system that measures the spring constant of the fruit surface under small deformations. In practice, we find that this accurately measures internal damage to the fruit. We calculate this metric by deforming the fruit a known distance and measuring the associated force required for the deformation. This setup can be seen in Fig. 3. We select equal spaced points along the circumference of each fruit and position the fruit so that the point is normal to a 3D-printed circular indenter. We measure the force in the z-direction after a 1.25mm displacement of the z-axis linear guide. The resulting force is measured using an ATI F/T Sensor Nano17. By averaging the spring constant of the four points on each fruit, we compare the effectiveness of the gripper with the DenseTact and its baselines before and after gripping. This metric, *% Stiffness*, is the percent

of pre-interaction stiffness, which is maintained after the experiment. 100% is the best possible value. Blackberries are excluded from this metric due to their high susceptibility to damage.

V. RESULTS

We demonstrate significant improvements over our baselines in grasping success rate and damage reduction over 630 trials across strawberries, tomatoes and blackberries. Strawberries are fragile, vary greatly in size, and show clear bruising on their surface after damage, making them an ideal challenge fruit for evaluating our method. Cherry tomatoes were selected due to their firm yet delicate skin, which makes them particularly sensitive to improper handling. Blackberries present additional complexity with their soft, easily ruptured structure, necessitating precise manipulation to avoid damage. Together, these fruits provide diverse tactile challenges critical for demonstrating the versatility of our approach.

A. Experimental Setup

The experimental setup is shown in Fig. 3. From a starting position, our policy is able to move to the strawberry and a grasp is attempted. The strawberry is then lifted and placed on the green square (goal). Any strawberry that is moved by the policy to the goal is labeled a success, otherwise it is labeled a failure. Failure modalities include missing a grasp of the strawberry completely or dropping the strawberry before reaching the target zone. Following best practices, the strawberries which are used for expert demonstrations are discarded and new strawberries are used for testing. We used the same two DenseTact sensors for the entire course of the project, across thousands of demos and experimental roll-outs.

To clearly highlight differences between the evaluated methods, we conduct a series of pick-and-place trials on strawberries, tomatoes, and blackberries. Specifically, we perform 25 trials per each strawberry, 50 per each tomato, and 20 per each blackberry. By conducting such a large number of trials, we both demonstrate the consistency of our method and create a key point of distinction from prior work. For each fruit type, we further vary the size categories: strawberries are tested across three size ranges (small: <4 cm, medium: 4–6 cm, large: >6 cm), tomatoes across two sizes (small: <3 cm, medium: \geq 3 cm), and blackberries across three sizes (small: <2.5 cm, medium: 2.5–3.5 cm, large: >3.5 cm). These size categories are based on measured fruit length. For each method and fruit size we perform all trials on the *same* fruit. This allows subtle damage to build up and become apparent, while showcasing the true delicacy of our method. To ensure consistency, all fruit used in the experiment were sourced on the same day from a single lot, ensuring similar ripeness and quality. The fruit starts on a small 3D-printed pedestal positioned in the center of the white start square. The gripper starts 20 cm away from the strawberry and moves on average 20 cm to the goal region. We employ an Intel Realsense D435i camera positioned on the robot wrist, angled toward the gripper. We utilize the Ufactory xArm7 operating in position control mode for the experiments.

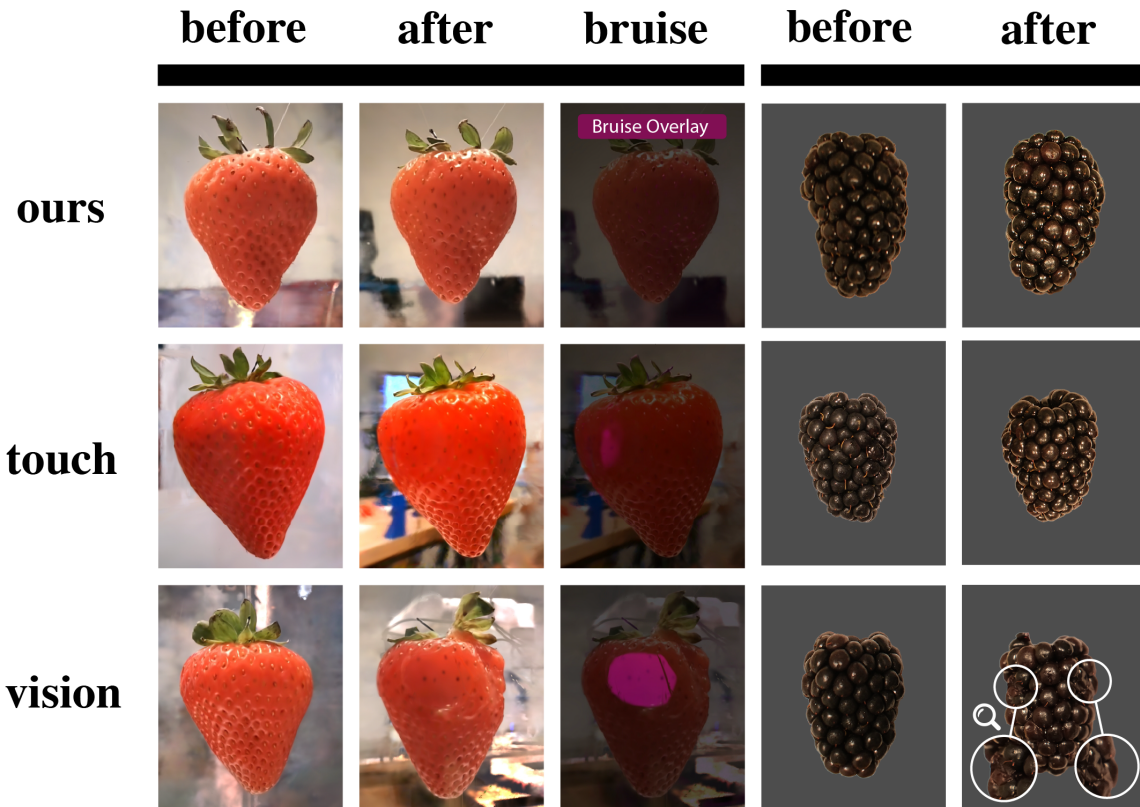


Fig. 4: We show the qualitative results of our method compared to the baselines. On the left, we show the same strawberries before and after the experiments along with the corresponding bruise mask. All the strawberry images are rendered from FruitSplat. Blackberries are shown on the right while tomatoes are not shown as they do not show visible damage.

Method	Strawberry			Tomato		Blackberry
	Success \uparrow	Stiffness \uparrow	Bruise \downarrow	Success \uparrow	Stiffness \uparrow	Success \uparrow
Ours	95%	80%	0.00%	100%	95%	80%
Touch	91%	78%	-1.75%	100%	93%	55%
Vision	67%	29%	14.61%	69%	77%	58%

TABLE I: Experimental results. Bold indicates best with margin.

Data collection is performed via a 3D Connexion space mouse and a mouse-based slider to control gripper position. Additionally, for our full method with tactile sensing, the RGB images of the DenseTact sensors are displayed to the teleoperator. We collect between 100 and 200 demos for each fruit type, ensuring a wide range of variability in our dataset. As our baselines are ablations of the full method, we use the same set of demos for each. We take 360-degree videos of each target fruit before and after the manipulations trials. These videos are then fed into our FruitSplat pipeline to produce quantitative damage results.

B. Baselines

In order to effectively evaluate our method, we compare our results to a number of baselines and ablations. *Vision* is our full gripper setup with the DenseTact sensors except the model does not incorporate touch information and instead relies only on the wrist-mounted camera. Likewise *Touch* relies only on

the tactile input. Finally we also test a naive combination of vision and touch where both inputs are combined without any natural switching labeled *Vision+Touch**.

C. Experimental Results

We perform a rigorous experimental testing program to effectively compare our method to the baselines of vision-only and touch-only. The quantitative results are shown in Tab. I and Fig. 5, while the qualitative images are shown in Fig. 4. To emphasize the differences between methods we run the same fruit with multiple grasps, 25 for strawberries, 50 for cherry tomatoes and 20 for blackberries. Examining the results in Tab. I, we see that our method outperforms or matches the baselines in success rate across all three of the tested fruits, while also achieving the least damage across our measured metrics. Success is defined as moving the fruit from the start to the goal and is independent from damage. Specifically for strawberries, we found that post-manipulation bruise difference was almost

zero, compared to the worst performing baseline of vision only with 14.6% bruising. It should also be noted that the negative percentage in strawberry bruising for touch only is likely caused by slight inconsistencies in pseudo-ground truth bruise mask detection by our YOLO model. We refer the readers to Fig. 5 which includes our estimated error bars based on the variance across trials.

Fig. 5 provides more detail showing the success rate across the different fruit sizes. We can see that the vision only methods struggle the most on small and medium fruit. We hypothesize that the vision only failure is due to an overfitting on size where the model fits to the average fruit size. While this leads to good grasping success rate on large fruit, it causes grasps which are far too loose on small and medium fruit. Furthermore, this often leads to over-squeezing of larger sized fruit causing more damage. This increased damage is clearly visible on the bottom of Fig. 5. Here we show the % change in stiffness relative to before the experiment. For comparison, we also include both *human* where a human performs the same action as the robot for each trial, and *teleop* where a human teleoperator controls the robot. While *ours* and *touch* perform similar to the theoretical optimal of *teleop* and *human* the *vision* policy causes a much greater amount of damage. Furthermore, we see that the vision only policy leads to much more external bruising.

We can also see in Fig. 5 an additional ablation labeled *vision+touch** which combines vision and touch without any intelligent switching. As discussed in Section V-B this policy performs poorly. We speculate this is due to the policy ignoring the latent tactile vector and overfitting on visual information even after contact has occurred. We can see that unlike the vision or touch only policies, this policy is worse across all three sizes of fruit indicating that it suffers from a different mode of failure.

Finally Fig. 4 shows the qualitative results of our experiments. We can see clear differences in bruising across our method and baselines. The high quality visual accuracy of the FruitSplat reconstructions of the strawberries are also visible. Tomatoes are not shown as they exhibit nearly no visible bruising when damaged.

VI. CONCLUSION

This work presents DexFruit, an optical tactile sensing embedded diffusion policy for soft fruit manipulation. We demonstrate that an intelligent switching between visual and tactile inputs greatly improves soft handling of fragile fruits such as strawberries, tomatoes, and blackberries compared to utilizing pure vision or touch sensing. We also present FruitSplat, a novel technique that harnesses the geometric accuracy and photorealism of 3D Gaussian Splatting to qualitatively and quantitatively represent fruit damage post-robot manipulation. Experimentation across various policy ablations and fruit types demonstrate DexFruit’s best performance in both policy grasping success rate and post-manipulation damage minimization, demonstrating an average 92% grasping policy success rate, up to a 14.6% reduction in visual bruising.

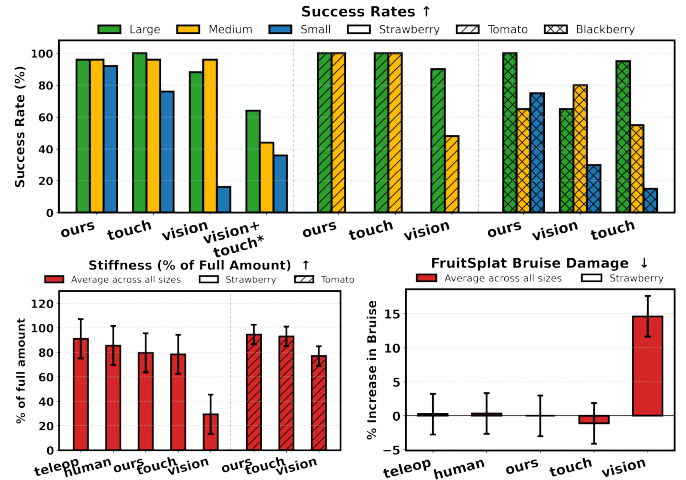


Fig. 5: Above we show the success rates across methods and fruit size. Below on the left we quantify the internal fruit damage measured as a percent change in stiffness, while on the right we show the results from FruitSplat representing external damage.

VII. LIMITATIONS AND FUTURE WORK

While DexFruit presents improvements in dexterous manipulation of soft fruits, some limitations and several opportunities for future work exist. With regards to our fruit selection for FruitSplat, we chose to restrict our initial scope of work on strawberries chiefly due to the data-collection and training time associated with implementing more fruits, which is an inherent limitation of our method. In principle, we expect our method can extend to any fruit that displays visible bruising such as blackberries. This is also the reason we did not implement tomatoes, due to the absence of visible bruising. In our experience with manipulating tomatoes for the baseline methods, we found that the best way to quantify any damage was through our stiffness metric. Since the body of the tomatoes are substantially less stiff than their outer layer of skin, excessive gripping of tomatoes only resulted in internal damage and was visually undetectable. For future work, we hope to extend FruitSplat to other fruits through curating an extensive database of labeled fruit damage. Additionally, integrating multi-spectral data [48] into FruitSplat could further increase its utility and robustness. Our experiments were conducted in controlled laboratory settings rather than actual field conditions. Further method evaluation should be done for both the harvesting and post-harvest processes out in the field and in produce warehouses. We believe that integrated ripeness sensing could be achieved by leveraging soft optical tactile sensing output. This system could simultaneously improve yield by reducing damage and improving harvest ripeness. On the policy side, the method for switching off vision and turning on tactile in our method is conditioned on a contact heuristic. While this is sufficient for this work, this could be extended to a learned mode switch as the policy chooses actively which input is more useful.

REFERENCES

- [1] L. Calvin, P. Martin, and S. Simmitt, "Adjusting to higher labor costs in selected u.s. fresh fruit and vegetable industries," U.S. Department of Agriculture, Economic Research Service, Tech. Rep. EIB-235, July 2022.
- [2] T. Luo and C. L. Escalante, "Us farm workers: What drives their job retention and work time allocation decisions?" *The Economic and Labour Relations Review*, vol. 28, no. 2, pp. 270–293, 2017.
- [3] H. C. J. Godfray, J. R. Beddington, I. R. Crute *et al.*, "Food security: the challenge of feeding 9 billion people," *science*, vol. 327, no. 5967, pp. 812–818, 2010.
- [4] Y. Edan, S. Han, and N. Kondo, "Automation in agriculture," *Springer handbook of automation*, pp. 1095–1128, 2009.
- [5] M. S. A. Mahmud, M. S. Z. Abidin, A. A. Emmanuel, and H. S. Hasan, "Robotics and automation in agriculture: present and future applications," *Applications of Modelling and Simulation*, vol. 4, pp. 130–140, 2020.
- [6] A. Klodd, E. Tepe, and E. Hoover, "Commercial fruit growing guides: Harvesting strawberries," 2021. [Online]. Available: <https://extension.umn.edu/strawberry-farming/harvesting-strawberries>
- [7] H. Zhou, Y. Zhao, P. Jin *et al.*, "Intelligent robots for fruit harvesting: recent developments and future challenges," *Precision Agriculture*, vol. 24, no. 2, p. 420–455, 2023.
- [8] C. Lenz, R. Menon, M. Schreiber *et al.*, "Hortibot: An adaptive multi-arm system for robotic horticulture of sweet peppers," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 31–38.
- [9] A. L. Gunderman, J. A. Collins, A. L. Myers *et al.*, "Tendon-driven soft robotic gripper for blackberry harvesting," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2652–2659, 2022.
- [10] A. Myers, A. Gunderman, R. Threlfall, and Y. Chen, "Determining hand-harvest parameters and postharvest marketability impacts of fresh-market blackberries to develop a soft-robotic gripper for robotic harvesting," *HortScience*, vol. 57, pp. 592–594, 05 2022.
- [11] B. Das, A. Hoque, S. Roy *et al.*, "Post-harvest technologies and automation: AI-driven innovations in food processing and supply chains," *Int. J. Sci. Res. Sci. Technol.*, vol. 12, no. 1, pp. 183–205, 2025.
- [12] A. Singh, G. Vaidya, V. Jagota *et al.*, "Recent advancement in postharvest loss mitigation and quality management of fruits and vegetables using machine learning frameworks," *Journal of Food Quality*, vol. 2022, no. 1, p. 6447282, 2022.
- [13] A. Chauhan, B. Brouwer, and E. Westra, "Robotics for a quality-driven post-harvest supply chain," *Current Robotics Reports*, vol. 3, 06 2022.
- [14] W. K. Do, B. Jurewicz, and M. Kennedy, "Densetact 2.0: Optical tactile sensor," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 12 549–12 555.
- [15] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE CVPR*, 2016, pp. 779–788.
- [16] J. Lv, Y. Wang, L. Xu *et al.*, "A method to obtain the near-large fruit from apple image in orchard for single-arm apple harvesting robot," *Scientia horticulturae*, vol. 257, p. 108758, 2019.
- [17] A. Qiu, C. Young, A. L. Gunderman *et al.*, "Tendon-driven soft robotic gripper with integrated ripeness sensing for blackberry harvesting," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 831–11 837.
- [18] C. J. Hohimer, H. Wang, S. Bhusal *et al.*, "Design and field evaluation of a robotic apple harvesting system with a 3d-printed soft-robotic end-effector," *Transactions of the ASABE*, vol. 62, no. 2, pp. 405–414, 2019.
- [19] A. De, D. Kumar, I. Kwuan *et al.*, "Compliant robotic gripper with integrated ripeness sensing for blackberry harvesting," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 056–11 062.
- [20] K. Junge, C. Pires, and J. Hughes, "Lab2field transfer of a robotic raspberry harvester enabled by a soft sensorized physical twin," *Communications Engineering*, vol. 2, no. 1, p. 40, 2023.
- [21] A. Porichis, M. Inglezou, N. Kegkeroglou *et al.*, "Imitation learning from a single demonstration leveraging vector quantization for robotic harvesting," *Robotics*, vol. 13, no. 7, p. 98, 2024.
- [22] C. H. Kim, A. Silwal, and G. Kantor, "Autonomous robotic pepper harvesting: Imitation learning in unstructured agricultural environments," *IEEE Robotics and Automation Letters*, 2025.
- [23] Y. Han, K. Yu, R. Batra *et al.*, "Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer," *IEEE/ASME Transactions on Mechatronics*, 2024.
- [24] L. He, L. Tao, Z. Ma *et al.*, "Cherry tomato firmness detection and prediction using a vision-based tactile sensor," *Journal of Food Measurement and Characterization*, vol. 18, no. 2, pp. 1053–1064, 2024.
- [25] U. L. Opara and P. B. Pathare, "Bruise damage measurement and analysis of fresh horticultural produce—a review," *Postharvest Biology and Technology*, vol. 91, pp. 9–24, May 2014.
- [26] Y. He, Q. Xiao, X. Bai *et al.*, "Recent progress of nondestructive techniques for fruits damage inspection: a review," *Critical Reviews in Food Science and Nutrition*, vol. 62, no. 20, pp. 5476–5494, 2022.
- [27] C. Moggia, J. Graell, I. Lara *et al.*, "Firmness at harvest impacts postharvest fruit softening and internal browning development in mechanically damaged and non-damaged highbush blueberries (*vaccinium corymbosum* L.)," *Frontiers in plant science*, vol. 8, p. 535, 2017.
- [28] W. Yang, F. Takeda, M. Zhang *et al.*, "Internal bruise damage in machine-harvested blueberries," in *XXXI International Horticultural Congress (IHC2022): International Symposium on Advances in Berry Crops 1381*, 2022, pp. 392–400.
- [29] Z. Ünal, T. Kızıldeniz, M. Özden *et al.*, "Detection of bruises on red apples using deep learning models," *Scientia Horticulturae*, vol. 329, p. 113021, 2024.
- [30] X. Zhou, Y. Ampatzidis, W. S. Lee *et al.*, "Deep learning-based postharvest strawberry bruise detection under uv and incandescent light," *Computers and Electronics in Agriculture*, vol. 202, p. 107389, 2022.
- [31] C. Chi, Z. Xu, S. Feng *et al.*, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [32] S. Mahmoudi, A. Davar, P. Sohrabipour *et al.*, "Leveraging imitation learning in agricultural robotics: a comprehensive survey and comparative analysis," *Frontiers in Robotics and AI*, vol. 11, p. 1441312, 2024.
- [33] B. Huang, Y. Wang, X. Yang *et al.*, "3d vitac: learning fine-grained manipulation with visuo-tactile sensing," in *Proceedings of Robotics: Conference on Robot Learning (CoRL)*, 2024.
- [34] Z. Liu, C. Chi, E. Cousineau *et al.*, "Maniwav: Learning robot manipulation from in-the-wild audio-visual data," *arXiv preprint arXiv:2406.19464*, 2024.
- [35] S. Wistreich, B. Shi, S. Tian *et al.*, "Dexskin: High-coverage conformable robotic skin for learning contact-rich manipulation," *arXiv preprint arXiv:2509.18830*, 2025.
- [36] S. Jiang, S. Zhao, Y. Fan, and P. Yin, "Gelfusion: Enhancing robotic manipulation under visual constraints via visuotactile fusion," 2025. [Online]. Available: <https://arxiv.org/abs/2505.07455>
- [37] H. Xue, J. Ren, W. Chen *et al.*, "Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation," *arXiv preprint arXiv:2503.02881*, 2025.
- [38] J. Song, C. Meng, and S. Ermon, "Denosing diffusion implicit models," in *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.
- [39] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 4104–4113.
- [40] Y. Zhou, C. Barnes, J. Lu *et al.*, "On the continuity of rotation representations in neural networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5745–5753.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [42] R. Wightman, "Pytorch image models," <https://github.com/rwightman/pytorch-image-models>, 2019.
- [43] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, "Pixel-wise view selection for unstructured multi-view stereo," in *European Conference on Computer Vision (ECCV)*, 2016.
- [44] Y. Tian, Q. Ye, and D. Doermann, "Yolov12: Attention-centric real-time object detectors," 2025. [Online]. Available: <https://arxiv.org/abs/2502.12524>
- [45] N. Ravi, V. Gabeur, Y.-T. Hu *et al.*, "Sam 2: Segment anything in images and videos," *arXiv preprint arXiv:2408.00714*, 2024. [Online]. Available: <https://arxiv.org/abs/2408.00714>
- [46] M. Tancik, E. Weber, E. Ng *et al.*, "Nerfstudio: A modular framework for neural radiance field development," in *ACM SIGGRAPH 2023 Conference Proceedings*, 2023.
- [47] O. Shorinwa, J. Sun, and M. Schwager, "Fast-splat: Fast, ambiguity-free semantics transfer in gaussian splatting," *arXiv preprint arXiv:2411.13753*, 2024.
- [48] S. N. Sinha, H. Graf, and M. Weinmann, "Spectralgaussians: Semantic, spectral 3d gaussian splatting for multi-spectral scene representation, visualization and analysis," *arXiv preprint arXiv:2408.06975*, 2024.