

# Playing with Food: Learning Food Item Representations through Interactive Exploration

Amrita Sawhney\*\*, Steven Lee\*\*, Kevin Zhang\*\*,  
Manuela Veloso, and Oliver Kroemer

Carnegie Mellon University, Pittsburgh PA 15232, USA  
{amritasa, stevenl3, klz1, mmv, okroemer}@cs.cmu.edu

**Abstract.** A key challenge in robotic food manipulation is modeling the material properties of diverse and deformable food items. We propose using a multimodal sensory approach to interact and play with food that facilitates the ability to distinguish these properties across food items. First, we use a robotic arm and an array of sensors, which are synchronized using ROS, to collect a diverse dataset consisting of 21 unique food items with varying slices and properties. Afterwards, we learn visual embedding networks that utilize a combination of proprioceptive, audio, and visual data to encode similarities among food items using a triplet loss formulation. Our evaluations show that embeddings learned through interactions can successfully increase performance in a wide range of material and shape classification tasks. We envision that these learned embeddings can be utilized as a basis for planning and selecting optimal parameters for more material-aware robotic food manipulation skills. Furthermore, we hope to stimulate further innovations in the field of food robotics by sharing this food playing dataset with the research community.

## 1 Introduction

Knowledge of an object’s material properties is important for robots learning to perform tasks that involve physical interactions. However, obtaining material properties for deformable objects can be difficult and time consuming. Food items, in particular, vary widely between and within food types depending on factors such as how they were grown, how they were stored, and whether they have been cooked [10]. As a result of our prior knowledge, humans typically can ascertain the basic properties of many different food items using only vision. For properties that are not always clearly conveyed through vision, humans will touch or interact with objects in order to disambiguate its internal material properties, such as knocking on a watermelon to determine its ripeness. Analogously, we believe that the ability to distinguish properties between food items can be learned using multimodal sensor data from interactive robot exploration [15].

In this paper, we present a unique multimodal food interaction dataset consisting of vision, audio, proprioceptive, and force data acquired autonomously through robot interactions with a variety of food items. Additionally, we use this dataset to explore a self-supervised method for learning embeddings that encode various food material

---

\*\* Equal Contribution

This work was supported by Sony AI.

properties and use visual data as input. The network used to learn these embeddings is trained using a triplet loss formulation, which groups similar and dissimilar samples based on the different types of interactive sensor data. In this manner, the robot can learn complex representations of these items autonomously without the need for subjective and time-consuming human labels.

To demonstrate the utility of the learned representations, we subsequently use the learned embeddings as input features to train regressors and classifiers for different tasks and compare their performances with baseline vision-only and audio-only approaches. Our experiments show that regressors and classifiers that use our learned embeddings outperform similar baseline networks on a variety of tasks, indicating that the visual embedding network encodes additional material property information from the different modalities, without requiring the robot to interact with the object at test time. Our project website is located here<sup>1</sup> along with a link to our dataset here<sup>2</sup>.

## 2 Related Work

Many recent works have utilized simulations in order to learn dynamics models and material properties of deformable objects [20,29]. For example, Matl et al. [21] collect data on granular materials and compare the visual depth information with simulation results in order to infer their properties. Yan et al. [32] learn latent dynamics models and visual representations of deformable objects by manipulating them in simulation and using contrastive estimation, which is similar to our approach. In comparison, we worked with real-world robots and objects to collect data since representations of food that are learned through a simulation environment may not accurately transfer to real world due to variable behaviors during complex tasks, such as large plastic deformations during cutting. There are simulators that can simulate large elasto-plastic deformation [4], but they are computationally expensive, unavailable to the public, and have not yet shown their efficacy in this particular domain.

Other works also use a variety of multimodal sensors to inform a robot of deformable object properties in order to better manipulate them [3,13,17]. Erickson et al. [8] use a highly specialized near-infrared spectrometer and texture imaging to classify the materials of objects. Meanwhile, Feng et al. [9] use a visual network and forces to determine the best location at which to skewer a variety of food items for bite acquisition. Finally, Zhang et al. [35] use forces and contact microphones to classify the hardness of ingredients in order to adjust the cutting parameters for slicing actions. In contrast to these works, we only use overhead images as input to our embedding network during inference time, while still incorporating multi-modal data during training.

Numerous works have focused on using interactive perception to learn about objects in the environment [2]. Katz and Brock [14] used interactive perception to determine the location and type of articulation on a variety of random objects. Sinapov et al. [26,27,28] had a robot interact with objects using vision, proprioception, and audio in order to categorize them. Chu et al. [5] utilized 2 Biotac sensors on a PR2 along with 5 exploratory actions in order to learn human labeled adjectives from haptic feedback. In

<sup>1</sup> <https://sites.google.com/view/playing-with-food>

<sup>2</sup> <https://tinyurl.com/playing-with-food-dataset>




Fig. 1: Our cutting experimental setup.




Fig. 2: Images from the 4 cameras spread around our cutting experimental setup.

our work, we focus our exploratory data collection on deformable food objects instead of rigid objects. Unlike the Epic Kitchens dataset [6], which captures a first person view of humans cooking, we capture proprioceptive information using a robot and repeat consistent interactive actions across food items.

Finally, researchers have been interested in learning deep embeddings for objects in order to create simpler representations for a variety of tasks [1,30]; however, performing the same task for food items has not been studied extensively. Sharma et al. [25] learn a semantic embedding network to represent the size of food slices in order to plan a sequence of cuts. On the other hand, Isola et al. [12] used human labeled adjectives to generalize the transformation of object states, such as the ripeness in fruit, over time. Although the above works learn embeddings for food objects, both of them focus on using solely visual inputs during both train and test time, while we incorporate additional synchronized multi-modal sensory information during the training of our vision-based embedding networks.

### 3 Robotic Food Manipulation Dataset

#### 3.1 Experimental Setup

Our data collection pipeline involves two different experimental setups: one for robotic food cutting and another for food playing, wherein the robot interacts with the food slices that have already been cut. We collected multi-modal sensor data during the cutting and playing data collection processes using our Franka robot control framework [34]. The sections below detail the setups for both cutting and playing in more detail.




Fig. 3: Our playing experimental setup.




Fig. 4: Examples of Kinect, Realsense, and FingerVision images of a cut tomato.

**Robot Cutting:** Our experimental setup for cutting data collection consists of two Franka Emika Panda Arms mounted on a Vention frame with a cutting board in the center, as shown in Fig. 1. One arm is grasping a custom knife attachment while the other has a set of 8” kitchen tongs mounted to its fingers. There are four cameras attached to the setup: an overhead Microsoft Azure Kinect Camera, a side-view Intel Realsense D435, another D435 mounted above the wrist of the robot holding the knife, and a third D435 mounted on the wrist of the robot holding the tongs. Sample images from each camera are shown in Fig. 2. In addition, there are 3 contact microphones: one mounted underneath the center of the cutting board, another mounted on the knife attachment, and the last mounted on the tongs.

**Robot Playing:** For our playing setup, we have a single Franka Emika Panda Arm mounted on a Vention frame with a cutting board in the center, as shown in Fig. 3. We mounted an overhead Microsoft Azure Kinect Camera and a frontal Intel Realsense D435 that faces the robot. We attach a fisheye 1080P USB Camera<sup>3</sup> to each fingertip as in FingerVision [31], except we use a laser-cut clear acrylic plate cover instead of a soft gel-based cover over the camera. This acrylic plate allows us to observe the compression of the object being grasped relative to fingertips. We also added a white LED to better illuminate the object while it is being grasped. Images from all of the cameras are shown in Fig. 4. We have 2 contact microphones on the setup: one mounted underneath the center of the cutting board and the other mounted on the back of the Franka Panda

<sup>3</sup> <https://www.amazon.com/180degree-Fisheye-Camera-usb-Android-Windows/dp/B00LQ854AG/>




Fig. 5: Food item slices. From left to right, top to bottom: apple, banana, bell pepper, boiled apple, boiled bell pepper, boiled carrot, boiled celery, boiled jalapeno, boiled pear, boiled potato, bread, carrot, celery, cheddar, cooked steak, cucumber, jalapeno, kiwi, lemon, mozzarella, onion, orange, pear, potato, raw steak, spam, strawberry, and tomato.

hand. The Piezo contact microphones<sup>4</sup> from both setups capture vibro-tactile feedback through the cutting board, fingers, and tools. The audio from the contact microphones of both the cutting and playing setup are captured using a Behringer UMC404HD Audio Interface<sup>5</sup> and synchronized with ROS using `sounddevice_ros` [33].

### 3.2 Data Collection

Using our robot cutting setup mentioned in Section 3.1, we taught the robot simple cutting skills using Dynamic Movement Primitives (DMPs) [16,23]. Through ridge regression, we fitted DMP parameters to trajectories collected using kinesthetic human

<sup>4</sup> <https://www.amazon.com/Agile-Shop-Contact-Microphone-Pickup-Guitar/dp/B07HVFTGTH/>

<sup>5</sup> <https://www.amazon.com/BEHRINGER-Audio-Interface-4-Channel-UMC404HD/dp/B00QHURLHM>

demonstrations as in [35]. Afterwards, we chained DMPs into multiple slicing motions until the specified food item was cut completely through.

In total, we cut 10 slices each from 21 different food types which include: apples, bananas, bell peppers, bread, carrots, celery, cheddar, cooked steak, cucumbers, jalapenos, kiwis, lemons, mozzarella, onions, oranges, pears, potatoes, raw steak, spam, strawberries, and tomatoes. The slices are enumerated from 1 to 14 and were created using similar skill parameters across the food types. The skill parameters vary in slice thickness from 3mm to 50mm, angles from  $\pm 30$  degrees, and slice orientation where we had normal vs. in-hand cuts when the knife robot cut between the tongs. There are more than 10 slice types because food items are shaped differently and not all cuts can be executed on every food item, especially the angled cuts. While the robot is cutting the food items, we collect audio, image, force, and proprioceptive data. The resulting slices from various food items are shown in Fig. 5.

After the 10 different types of slices have been created, we transfer them to the playing robot setup and begin the data collection process. We run 5 trials on each of the 10 slices in order to capture variations in the objects' behaviors due to differing initial orientations, positions, and changes over time. In each trial, we capture a RGBD image from the overhead Kinect and specify the center of the object. Afterwards, the robot closes its fingers and pushes down on the object until 10N of force is measured. We record the robot's position and forces during this action. Next, the robot resets to a known position, grasps the object, and releases it from a height of 15cm. During the push, grasp, and release actions, we record videos from the Realsense and audio from the two contact microphones as mentioned in Section 3.1. Videos are only recorded from the FingerVision cameras during the grasp and release actions. Additionally, we record the gripper width when the grasp action has finished and save RGBD images from the overhead Kinect before and after each trial.

Our full dataset is available for download here<sup>6</sup>. The data is located in the appropriately named folders and are sorted first by food type, then slice type, and finally trial number. Additionally, we provide food segmentation masks in the silhouette data folder where we used Deep Extreme Cut (DEXTR) [19] to obtain hand labeled masks of the objects in the overhead and side view images. Then we fine-tuned a PSPNet [36] pre-trained on the Ade20k [37] dataset with our manually labeled masks to generate additional neural network labeled segmentation masks. Finally, we have additional playing data in the old playing data folder, but those slices were all hand cut, and the data was collected in a different environment.

### 3.3 Data Processing

To train the embedding networks, we first extract features from the data. We transform the raw audio data from both the cutting data and playing data (during the release action, push-down action, and grasp action) into Mel-frequency cepstrum coefficient (MFCC) features [18] using Librosa [22]. These features have been shown to effectively differentiate between materials and contact events [35]. Subsequently, we use

<sup>6</sup> <https://tinyurl.com/playing-with-food-dataset>




Fig. 6: An overview of our approach. The different features (modalities) defined in Section 3.3 are used to form triplets to learn embeddings in an unsupervised manner, which are used for supervised classification and regression tasks (blue text).

PCA to extract a lower-dimensional representation of the cutting ( $\mathcal{A}_{cut}$ ) and playing ( $\mathcal{A}_{play}$ ) audio features.

To form proprioceptive features ( $\mathcal{P}$ ), we use the robot poses and forces from the push down and grasp actions. More specifically, for the push down action we extract the final  $z$  position ( $z_f$ ) of the robot’s end-effector once 10N of force has been reached. Using this value, we also find the change in  $z$  position between the point of first contact and  $z_f$  as  $\Delta z$ , which is an indication of the object’s stiffness. We retrieve the final gripper width ( $w_g$ ) during the grasping action once 60N of force has been reached. These three values are combined to form  $\mathcal{P}$ . Labels for each data sample, such as food class label ( $\mathcal{F}$ ) and slice type label ( $\mathcal{S}$ ), are also created according to the food type and slice type performed during cutting.

## 4 Learning Food Embeddings

We train convolutional neural networks, in an unsupervised manner, to output embeddings from overhead images. Our architecture is comprised of ResNet34 [11], which is pretrained on ImageNet [7] and has the last fully connected layer removed. We add an additional three hidden layers, with ReLU activation for the first two, to reduce the dimensionality of the embeddings. We use a triplet loss [24] to train the network, so similarities across food types are captured in the embeddings.

The different modalities of data mentioned in Section 3.3 are used as metrics to form the triplets. More specifically, the food class labels ( $\mathcal{F}$ ), slice type labels ( $\mathcal{S}$ ), playing audio features ( $\mathcal{A}_{play}$ ), cutting audio features ( $\mathcal{A}_{cut}$ ), proprioceptive features ( $\mathcal{P}$ ), and combined audio and proprioceptive features ( $\mathcal{A}_{play} + \mathcal{P}$ ) are used as metrics. For each sample in the training set, we define the  $n$  nearest samples in the PCA feature space (using the L2 norm) as the possible positive samples in a triplet and all other samples as possible negative samples, where  $n$  is a hyperparameter ( $n = 10$  was used here). At training time, triplets are randomly formed using these positive/negative identifiers. Fig. 6 shows an overview of our approach.

To evaluate the usefulness of these learned embeddings, we train multiple 3-layer multilayer perceptron classifiers and regressors for a variety of tasks, using the learned embeddings as inputs. These results are presented in Section 5. When combining multiple modalities, we concatenate the embeddings output from each separate network.

## 5 Experiments

Embeddings	Food Type Accuracy - 21 classes (%)	Hardness Accuracy - 3 classes (%)	Juiciness Accuracy - 3 classes (%)	Slice Type Accuracy - 14 classes (%)	Slice Width RMSE (mm)
$\mathcal{F}$	92.0	40.7	36.6	12.9	10.9
$\mathcal{S}$	17.1	37.0	34.9	<b>40.5</b>	11.8
$\mathcal{A}_{play}$	85.7	35.0	<b>46.0</b>	17.1	9.9
$\mathcal{A}_{cut}$	93.5	33.5	45.6	16.8	11.3
$\mathcal{P}$	49.5	<b>47.1</b>	37.0	20.0	<b>7.9</b>
$\mathcal{A}_{play}+\mathcal{P}$	83.8	36.4	40.2	21.4	9.5
ResNet	<b>98.9</b>	34.9	36.5	30.0	13.9
Classifier w/ $\mathcal{A}_{play}$ as input	84.4	40.8	34.0	30.1	34.4

Table 1: Baseline and multi-layer perceptron results on 5 evaluation tasks using different learned embedding networks that were trained on our full dataset.

As mentioned in the previous section, embedding networks were trained using the food class label ( $\mathcal{F}$ ), slice type label ( $\mathcal{S}$ ), playing audio features ( $\mathcal{A}_{play}$ ), cutting audio features ( $\mathcal{A}_{cut}$ ), proprioceptive features ( $\mathcal{P}$ ), and combined audio and proprioceptive features ( $\mathcal{A}_{play}+\mathcal{P}$ ) for creating triplets. These embeddings were then used to train the multi-layer perceptrons, mentioned in Section 4, to predict the labels or values for five different tasks: classifying food type (21 classes), predicting slice width (the width of the gripper after grasping), classifying the hardness (3 human-labeled classes - hard, medium, and soft), classifying juiciness (3 human-labeled classes - juicy, medium, dry), and classifying slice type (14 different classes based on the type of cuts the cutting robot performed to generate the slice). For our two baselines, we trained convolutional neural networks, with a ResNet34 architecture, that use only visual data and another set of 3-layer multi-layer perceptrons that use only  $\mathcal{A}_{play}$  data as input to generate predictions for each of the tasks.

To assess the generalizability of our approach, we evaluated the hardness and juiciness classification tasks based on leave-one-out classification, where we left an entire food class out of the training set and evaluated the trained classifiers on this class at test time. We then averaged the results across the 21 leave-one-out classification trials. The performance of all the trained networks on each of the tasks are shown in Table 1.

As illustrated in Table 1, the purely visual baseline outperforms our embeddings in the food type classification task due to the vast number of labeled images in the ImageNet dataset that ResNet was pre-trained on. However, ResNet performs worse on the other 4 tasks as ImageNet did not contain prior relevant information on the physical properties that are important for these tasks. Additionally, ResNet was trained to differentiate object classes instead of finding similarities between classes, so when an entire food category was left out of the training dataset, it most likely had no way of extrapolating the correct answer from previous data. Meanwhile, our embeddings contained auxiliary information that encoded the similarity of slices through various multimodal features, without ever being given explicit human labels. This indicates that our inter-



Embeddings	Hardness Accuracy (%)	Juiciness Accuracy (%)	Cooked Accuracy (%)
$\mathcal{A}_{play}$	98.0	62.9	98.9
$\mathcal{P}$	63.0	68.4	60.6
$\mathcal{A}_{play}+\mathcal{P}$	<b>99.7</b>	<b>70.6</b>	<b>99.1</b>
ResNet	90.5	66.1	90.4
Classifier w/ $\mathcal{A}_{play}$ as input	82.1	67.4	88.8

Table 2: Results on 3 evaluation tasks for different learned embedding networks that were trained using the auxiliary cooked vs. uncooked dataset.

active multimodal embeddings provided the neural networks with a greater ability to generalize to unseen data as compared to the supervised, non-interactive baselines.

Additionally, the results show that the audio embeddings provide some implicit information that can help the robot distinguish vegetable types. On the other hand, it makes sense that the proprioceptive embeddings are more useful at predicting hardness and slice width as their triplets were generated using similar information. However, absolute labels were never provided when training the embedding networks, so the learned embeddings encoded this relative information themselves. It should also be noted that in the hardness and juiciness leave-one-out classification tasks, some food types, such as tomatoes, were more difficult to classify when left out of the training dataset than others, such as carrots. This may be due to the small size of our diverse dataset, which has few items with similar properties.

Finally, with respect to the slice type prediction task, there were poor performances across the board due to the inherent difficulty of the task. Due to the variability of shapes between food items, the resulting slices generated by the cutting robot, while executing the same actions, differed greatly at times. Thus, it was to be expected that only the embeddings trained using slice type labels performed relatively well on this classification task. Overall, the results of the evaluations above show that certain embeddings performed better on relevant tasks, which supports the hypothesis that they are encoding information on different material properties and can be applied in different use cases.

**Auxiliary Study with Additional Cooked vs. Uncooked Food Data:** As an addendum to the 21-food class dataset described in Section 3.2, we collected an additional dataset of boiled food classes to further explore and evaluate our method’s ability to detect whether a food item is cooked or not through interactive learning. The additional boiled food classes collected were: apples, bell peppers, carrots, celery, jalapenos, pears, and potatoes. Each item was boiled for 10 minutes. Note that for these additional cooked food classes, we did not have the robot cut the food slices due to difficulties grasping the objects. We combined these boiled classes with their uncooked counterparts from the full dataset to form a 14-class dataset and conducted a subset of evaluations on the embeddings learned from this dataset, shown in Table 2.

The auxiliary study with the boiled food dataset shows that the playing audio data is effective at autonomously distinguishing between cooked and uncooked food items without being provided any human-provided labels. This is likely because cooking food




Fig. 7: Fig. 7a visualizes playing audio features using PCA. Fig. 7b visualizes the embeddings learned using the playing audio features also in PCA space.

significantly changes its material properties. The high performance of the  $(\mathcal{A}_{play} + \mathcal{P})$  embeddings on the hardness, juiciness, and cooked leave-one-out classification tasks on this smaller dataset demonstrates the ability of our approach to generalize to new data when a food class with similar properties was present during training (apples and pears, potatoes and carrots, bell peppers and jalapenos).

Fig. 7a shows the playing audio features ( $\mathcal{A}_{play}$ ) of the different cooked and uncooked food classes in this dataset using the top 3 principal components. Fig. 7b visualizes the learned embeddings based on  $\mathcal{A}_{play}$ , also in PCA space. As shown in the plots, there is a distinct separation between the boiled and raw foods in the audio feature space and also in the learned embedding space. Within the cooked and uncooked groupings, there are certain food types that tend to cluster together. For example, the uncooked pear and apple cluster close together, which makes sense given the similarity of these two fruits. Interestingly, they remain clustered close to one another even after they are cooked, even though there is a shift in the feature space between cooked and uncooked.

## 6 Conclusions and Future Work

In this work, we have presented a novel dataset consisting of autonomously collected audio, proprioceptive, force, and visual data that was recorded while a robot played with a variety of slices from 21 unique food types that have different shapes and properties. In addition, we learned visual embedding networks that utilized our multi-modal dataset to encode properties of food items using a triplet loss formulation. These learned embeddings were shown to encode similarities between food types without explicit human labeling and outperformed normal visual-only and audio-only baselines on a variety of tasks. We hope others can utilize our publicly available dataset to explore novel ways of deciphering and encoding the material properties of food items in order to further propel research on food robotics forward.

For example, an interesting extension to this work would be to apply state of the art computer vision tracking algorithms in order to observe the deformations and movements of the slices using the videos captured from the RealSense and FingerVision cameras. This could serve as an additional metric for creating triplets or possibly be used as a basis for creating dynamics models for different types of food. In addition,

since we recorded videos when cutting the slices, it may be possible to predict the shape of resulting slices given an action or vice versa.

For us, the next challenges we hope to tackle using this dataset include: monitoring the progression of food as it is being cooked in order to inform the robot when to intervene, and using the learned embeddings to better execute manipulation tasks such as cutting, flipping, mixing, picking, and placing food items. If we collect data that is complementary to this dataset then we will append the additional data to the original dataset. We also welcome others to contribute as well. Overall, we believe that this work is an exciting step towards autonomously learning about deformable food items through robotic interactions and play.

## References

1. Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
2. J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, and G. S. Sukhatme. Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017.
3. P. Boonvisut and M. C. Çavuşoğlu. Estimation of soft tissue mechanical parameters from robotic manipulation data. *IEEE/ASME Transactions on Mechatronics*, 18(5):1602–1611, 2012.
4. N. Chentanez, M. Müller, and M. Macklin. Real-time simulation of large elasto-plastic deformation with shape matching. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 159–167, 2016.
5. V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. M. Perez-Tejada, M. Arrigo, N. Fitter, J. C. Nappo, T. Darrell, et al. Using robotic exploratory procedures to learn the meaning of haptic adjectives. In *2013 IEEE International Conference on Robotics and Automation*, pages 3048–3055. IEEE, 2013.
6. D. Damen, H. Doughty, G. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, et al. The epic-kitchens dataset: collection, challenges and baselines. *IEEE Computer Architecture Letters*, (01):1–1, 2020.
7. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009.
8. Z. Erickson, E. Xing, B. Srirangam, S. Chernova, and C. C. Kemp. Multimodal material classification for robots using spectroscopy and high resolution texture imaging. *arXiv:2004.01160*, 2020.
9. R. Feng, Y. Kim, G. Lee, E. K. Gordon, M. Schmittle, S. Kumar, T. Bhattacharjee, and S. S. Srinivasa. Robot-assisted feeding: Generalizing skewering strategies across food items on a realistic plate. *arXiv preprint arXiv:1906.02350*, 2019.
10. L. Figura and A. A. Teixeira. *Food physics: physical properties-measurement and applications*. Springer Science & Business Media, 2007.
11. K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
12. P. Isola, J. J. Lim, and E. H. Adelson. Discovering states and transformations in image collections. In *CVPR*, 2015.
13. B. Jia, Z. Hu, J. Pan, and D. Manocha. Manipulating highly deformable materials using a visual feedback dictionary. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 239–246. IEEE, 2018.

14. D. Katz and O. Brock. Manipulating articulated objects with interactive perception. In *2008 IEEE International Conference on Robotics and Automation*, pages 272–277. IEEE, 2008.
15. O. Kroemer, S. Niekum, and G. Konidaris. A review of robot learning for manipulation: Challenges, representations, and algorithms. *arXiv preprint arXiv:1907.03146*, 2019.
16. O. Kroemer and G. Sukhatme. Meta-level priors for learning manipulation skills with sparse features. In *International Symposium on Experimental Robotics*, pages 211–222. Springer, 2016.
17. Q. Li, O. Kroemer, Z. Su, F. F. Veiga, M. Kaboli, and H. J. Ritter. A review of tactile information: Perception and action through touch. *IEEE Transactions on Robotics*, 2020.
18. B. Logan et al. Mel frequency cepstral coefficients for music modeling. In *Ismir*, volume 270, pages 1–11, 2000.
19. K. Maninis, S. Caelles, J. Pont-Tuset, and L. Van Gool. Deep extreme cut: From extreme points to object segmentation. In *CVPR*, 2018.
20. J. Matas, S. James, and A. J. Davison. Sim-to-real reinforcement learning for deformable object manipulation, 2018.
21. C. Matl, Y. Narang, R. Bajcsy, F. Ramos, and D. Fox. Inferring the material properties of granular media for robotic tasks. *arXiv:2003.08032*, 2020.
22. B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto. librosa: Audio and music signal analysis in python. 2015.
23. S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert. Learning movement primitives. In *Robotics research. the eleventh international symposium*, pages 561–572. Springer, 2005.
24. F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015.
25. M. Sharma, K. Zhang, and O. Kroemer. Learning semantic embedding spaces for slicing vegetables. *arXiv:1904.00303*, 2019.
26. J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, and A. Stoytchev. Interactive object recognition using proprioceptive and auditory feedback. *The International Journal of Robotics Research*, 30(10):1250–1262, 2011.
27. J. Sinapov, C. Schenck, K. Staley, V. Sukhoy, and A. Stoytchev. Grounding semantic categories in behavioral interactions: Experiments with 100 objects. *Robotics and Autonomous Systems*, 62(5):632–645, 2014.
28. G. Tatiya and J. Sinapov. Deep multi-sensory object category recognition using interactive behavioral exploration. In *ICRA*, pages 7872–7878. IEEE, 2019.
29. Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel. Learning to manipulate deformable objects without demonstrations, 2020.
30. J. Xie, R. Girshick, and A. Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, pages 478–487, 2016.
31. A. Yamaguchi and C. G. Atkeson. Combining finger vision and optical tactile sensing: Reducing and handling errors while cutting vegetables. In *Humanoids*, pages 1045–1051. IEEE, 2016.
32. W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto. Learning predictive representations for deformable objects using contrastive estimation, 2020.
33. K. Zhang. [https://github.com/firephnix/sounddevice\\_ros](https://github.com/firephnix/sounddevice_ros).
34. K. Zhang, M. Sharma, J. Liang, and O. Kroemer. A modular robotic arm control stack for research: Franka-interface and frankapy. *arXiv preprint arXiv:2011.02398*, 2020.
35. K. Zhang, M. Sharma, M. Veloso, and O. Kroemer. Leveraging multimodal haptic sensory data for robust cutting. In *Humanoids*, pages 409–416, 2019.
36. H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *CVPR*, pages 2881–2890, 2017.
37. B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba. Scene parsing through ade20k dataset. In *CVPR*, pages 633–641, 2017.