

TactEx: An Explainable Multimodal Robotic Interaction Framework for Human-Like Touch and Hardness Estimation

Felix Verstraete*, Lan Wei*, Wen Fan*, and Dandan Zhang

Abstract—Accurate perception of object hardness is essential for safe and dexterous contact-rich robotic manipulation. Here, we present TactEx, an explainable multimodal robotic interaction framework that unifies vision, touch, and language for human-like hardness estimation and interactive guidance. We evaluate TactEx on fruit-ripeness assessment, a representative task that requires both tactile sensing and contextual understanding. The system fuses GelSight-Mini tactile streams with RGB observations and language prompts. A ResNet50+LSTM model estimates hardness from sequential tactile data, while a cross-modal alignment module combines visual cues with guidance from a large language model (LLM). This explainable multimodal interface allows users to distinguish ripeness levels with statistically significant class separation ($p < 0.01$ for all fruit pairs). For touch placement, we compare YOLO with Grounded-SAM (GSAM) and find GSAM to be more robust for fine-grained segmentation and contact-site selection. A lightweight LLM parses user instructions and produces grounded natural-language explanations linked to the tactile outputs. In end-to-end evaluations, TactEx attains 90% task success on simple user queries and generalises to novel tasks without large-scale tuning. These results highlight the promise of combining pretrained visual and tactile models with language grounding to advance explainable, human-like touch perception and decision-making in robotics.

I. INTRODUCTION

Humans rely on touch to infer intrinsic physical properties (e.g., hardness, friction, compliance) and to regulate contact, especially when vision is unreliable due to occlusion, poor lighting, or visually ambiguous materials [1]. In contact-rich interaction, mechanical properties govern how objects deform under applied forces and whether contact remains stable. Touch therefore enables not only stable manipulation but also interpretable comparative judgments (e.g., “this fruit is softer than that one”), which are often ambiguous from visual cues alone.

In robotics, tactile sensing provides complementary information that cannot be directly recovered from vision. High-resolution tactile imaging supports texture recognition [2] and hardness estimation [3], while high-bandwidth tactile cues enable early slip detection and precise force modulation [4]. These capabilities make tactile sensing particularly important for intrinsic property inference during physical interaction. Among the intrinsic properties accessible through touch, hardness is especially critical for manipulation. Accurate hardness estimation supports safe force regulation, adaptive grasping, and delicate object handling. However,

fine-grained hardness discrimination requires controlled contact, high-resolution tactile sensing, and statistically reliable comparison, particularly when visually similar objects differ mechanically. Fruit ripeness assessment exemplifies this challenge: subtle compliance variations must be distinguished under visually ambiguous conditions [5]. Beyond its practical relevance in agriculture and domestic robotics, ripeness assessment provides a principled benchmark for multimodal hardness estimation.

Despite recent advances in tactile hardware [6] and tactile-driven manipulation, most approaches prioritize task success [7], such as pushing, grasping, or tactile servoing, rather than explicit intrinsic property inference. Hardness estimation models remain limited in practice, particularly for fine-grained ripeness discrimination. Existing methods often require large datasets or extensive fine-tuning [8], [9], and many lack statistical validation for subtle compliance differences [10]. As a result, calibrated and reliable hardness estimation remains an open challenge.

Beyond accuracy, explainability is critical for deploying tactile perception in human-facing applications. When robots estimate intrinsic properties, they should be able to justify how tactile evidence supports the inference and how that inference informs subsequent control decisions. While structured tactile representations [11] and interpretable policy learning [12] improve transparency at the feature or policy level, they rarely connect calibrated property estimation to operator-facing explanations grounded in tactile measurements. In line with Adebayo et al.’s xAI framework for robotics [13] and prior work emphasizing user-oriented explanations [14], practical systems should provide transparent reasoning, uncertainty reporting, and sensor-grounded rationales. Without these capabilities, trust and adoption in human-facing settings remain limited. These gaps motivate a framework that treats hardness estimation as a first-class perceptual objective within a structured, multimodal, and explainable interaction loop.

To address this need, we propose **TactEx**, an explainable multimodal robotic interaction framework for touch-based hardness estimation (Fig. 1). TactEx integrates language grounding, visual perception, controlled tactile exploration, and statistically grounded hardness inference in a modular pipeline. Given a natural-language query, the robot localizes the referenced objects, performs top-down probing with a high-resolution tactile sensor, and estimates hardness from the resulting contact sequence. A large language model (LLM) then produces concise, sensor-grounded explanations and comparisons, such as ripeness or hardness rankings

*Equal Contribution. All authors are with the Department of Bioengineering, Imperial-X Initiative, Imperial College London, London, United Kingdom. Corresponding: d.zhang17@imperial.ac.uk.

across the queried fruits.

TactEx is designed around three principles: (i) multimodal interaction, where touch is treated as a primary perceptual signal coupled with exploratory contact; (ii) human-like hardness reasoning, enabling fine-grained comparative estimation with confidence reporting; and (iii) explainability, providing intermediate outputs and uncertainty-aware rationales grounded in tactile evidence. Together, these components support accurate and interpretable hardness estimation in contact-rich, human-facing settings.

This paper makes the following contributions:

- 1) **Data-efficient visuo-tactile hardness regression.** We introduce a pretrained ResNet50 + LSTM hardness estimation pipeline that achieves RMSE 4.3 and $\rho = 0.88$ using only $N=280$ fine-tuning samples, with statistically significant ripeness ranking across five fruit types.
- 2) **Language-conditioned servoing for touch perception.** We demonstrate that text-prompted Grounded-SAM outperforms YOLO in touch placement accuracy, leading to improved downstream hardness estimation reliability.
- 3) **End-to-end explainable multimodal interaction.** We present a language-grounded framework that converts user queries into visuo-tactile actions and sensor-grounded explanations, achieving high object- and scenario-level success across four interaction complexity tiers.

II. RELATED WORK

A. Hardness and Ripeness Estimation

While several works treat hardness as a classification problem using GelSight or traditional tactile sensors [10], [15], discrete labels lack the resolution required for objects that cluster tightly at the high end of the Shore 00 scale. Because many fruits and vegetables fall within this narrow band, classification systems often fail to distinguish subtle mechanical variations, such as the difference between an unripe and a ripe banana. This limitation necessitates a regression formulation capable of resolving these fine-grained compliance differences.

Physics and feature-engineered methods aim to estimate hardness by tracking physical changes in tactile images or force signals. Yuan et al. proposed a numerical model based on the changing brightness in the tactile image and the force variations [16], and Liao et al. leveraged force dynamics for ripeness tracking [5]. Although effective in controlled settings, these approaches rely on hand-crafted features, impose shape or pose constraints, show limited generalizability across objects and contact regimes, and rarely integrate other modalities.

Deep learning approaches, such as the widely adopted VGG16-LSTM GelSight baseline by Yuan et al. [3], have demonstrated strong performance in foundational robotic applications. However, these unimodal systems do not report statistically validated rankings for fine-grained properties like

ripeness, restricting their real-world applicability. Similarly, recent work by Nam et al. [9] explored continuous tactile regression but encountered performance degradation when assessing harder objects (upper Shore 00 range) and did not incorporate an explainable, language-conditioned interface.

B. Language Grounded Multimodal Models

Recent work increasingly integrates language with vision and tactile sensing. In these systems, language is used to interpret user requests and to generate responses [15], [17], [18]. For example, Ueda et al. demonstrate that a robot equipped with tactile sensors can leverage the zero-shot capabilities of vision-language models to recognize objects. Tactile-VLA shows that augmenting vision-language models with touch improves the translation of user intent into precise physical actions [19]. Force-VLA reports performance gains in contact-rich manipulation [20], and VTLa further improves challenging insertion tasks [21]. Despite these advances, the focus remains primarily on action execution, with less attention to improving perception and language-conditioned reasoning about object properties. In this study, we plan to integrate language and vision to support perception by helping to locate and filter objects in the scene.

III. METHODS

We present **TactEx** (“The Tactile Explainer”), a modular and explainable multimodal robotic framework that *perceives* and *verbally explains* object hardness and fruit ripeness (Fig. 1). The system comprises three stages: (A) *Input and Interface*, (B) *Back-end Processing*, and (C) *Outputs*.

A. System Interface

The user interface (Fig. 1(A)) is implemented in Streamlit, providing a lightweight and interactive front end for human-robot interaction. Upon launch, the interface displays a live RGB stream from the robot-mounted camera, showing fruits and vegetables arranged on a workspace. A chat window below the video feed accepts free-form natural-language prompts (e.g., “Which fruit is the softest?” or “How ripe are the banana and the lemon?”).

Submitting a query initiates the back-end processing pipeline (Fig. 1(B)). The system first interprets the user’s language instruction to determine the relevant objects. It then executes perception, manipulation, tactile inference, and language generation in sequence, as described below.

a) Vision and Object Grounding: We employ the Grounded Segment-Anything Model (GSAM) for zero-shot, text-promptable object identification in unstructured household environments [22]. Conditioned on the user’s language query, GSAM returns instance segmentation masks and corresponding object locations (Fig. 1(B1)). We use these outputs to initialise and guide visual servoing, allowing accurate end-effector alignment and positioning for subsequent tactile probing. We compare GSAM against a conventional baseline that combines YOLO-based object detection with language-conditioned selection [23]. In this pipeline, YOLO provides

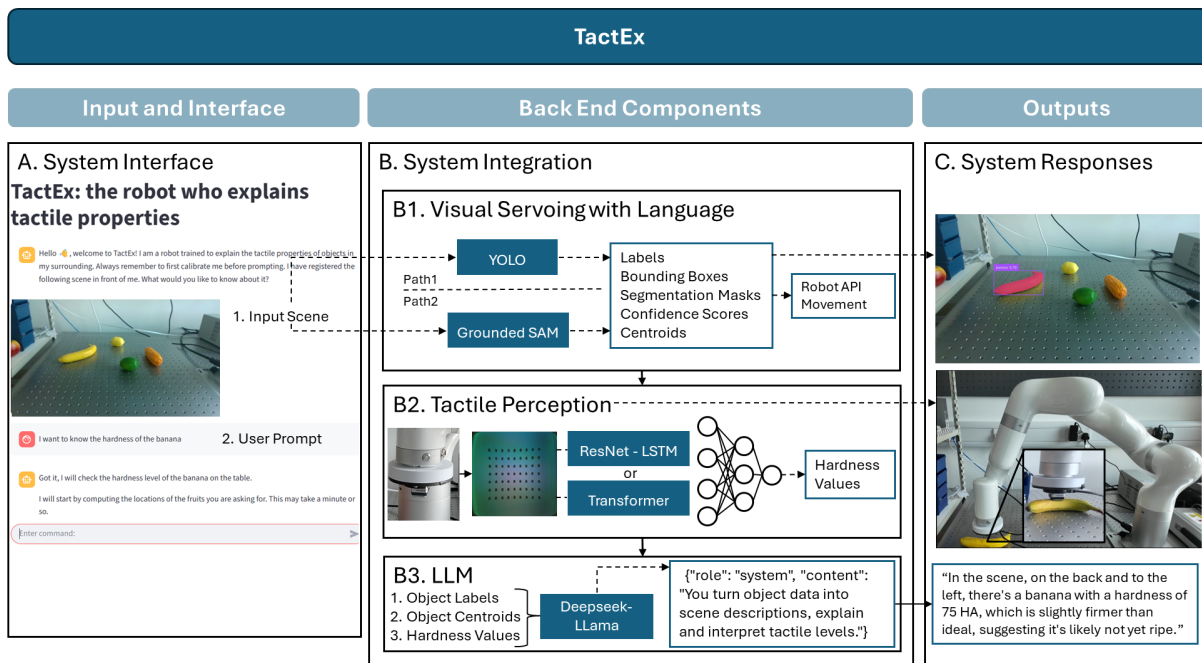


Fig. 1. Overview of TactEx (“The Tactile Explainer”), a multimodal framework for fruit ripeness explanation. Users interact via a chat interface (A), objects are localized with YOLO or GSAM (B1), hardness is estimated with a GelSight sensor (B2), and an LLM composes the final response from the fruit names, locations and hardness values (B3–C). The three components are detailed in III-B.1, section III-B.2 and section III-B.3, respectively.

candidate bounding-box detections, while the language instruction is used to select the subset of detections matching the query.

b) Tactile Acquisition and Hardness Estimation: Once the target objects are localised, the robot performs a controlled top-down pressing motion using a GelSight-Mini tactile sensor to capture a sequence of contact images. The resulting spatiotemporal tactile data are processed by either a ResNet-LSTM architecture or a Transformer-based regressor to estimate object hardness (Fig. 1(B2)).

c) Language Generation and Output: Finally, the estimated hardness values are combined with object labels and spatial context to produce a concise natural-language explanation via a large language model (Fig. 1(B3)). The interface then presents the results to the user (Fig. 1(C)), including object positions, predicted hardness values, and inferred ripeness expressed in user-friendly terms.

B. System Components

1) Visual Servoing Methods: We employ an eye-to-hand configuration for visual servoing. An Intel RealSense depth camera is mounted above the workspace, with intrinsic parameters $f_x = 608.5$, $f_y = 606.9$, $c_x = 309.4$, and $c_y = 213.8$. Workspace coordinates are expressed in metric units (mm) by projecting pixel detections into 3D space using camera intrinsics and extrinsics obtained via calibration with a 7×5 ChArUco checkerboard (square size: 2.5 cm; marker size: 1.8 cm).

The user’s request is first parsed using a lightweight NLP module to extract the target fruit classes and intended comparison (e.g., hardness ranking). These extracted object labels condition the subsequent perception stage. We compare two segmentation approaches: (i) a YOLOv8-based

instance segmentation model and (ii) the Grounded Segment-Anything Model (GSAM) [22].

The YOLOv8 model was trained on a custom dataset of 60 images containing up to 12 fruits and vegetables per scene for 100 epochs. Ground-truth annotations were manually generated using Roboflow. Training was performed with default hyperparameters (pretrained=True, learning rate=0.01, optimizer=SGD). At inference time, a confidence threshold of 0.40 was selected to balance false positives and missed detections. Since YOLO predicts all detectable object classes, NLP-based filtering is applied to retain only those objects referenced in the user’s query. In contrast, GSAM requires no task-specific training. A higher confidence threshold (0.60) is applied during inference to reduce spurious segmentations. Also, NLP filtering is used to remove objects not mentioned in the request.




Fig. 2. Example of the Grounded SAM procedure: (a) original scene, (b) object detection with bounding box, (c) results from SAM with inner mask for computing the centroid.

To mitigate the effect of noisy depth measurements on 3D localisation, both pipelines employ a multi-stage robustness procedure: (i) For GSAM, Grounding DINO first predicts a bounding box from the language prompt (Fig. 2(b)). (ii) Instance masks are then obtained from YOLO or GSAM. (iii) To suppress boundary artefacts, mask edges are detected using a Canny operator and dilated with a 3×3 kernel (two

iterations), producing an inner mask that excludes noisy contours (Fig. 2(c)). (iv) Depth at each pixel within the refined mask is stabilised by computing the median across 10 consecutive frames. (v) The 3D centroid is finally computed as the median of the filtered mask coordinates projected into metric space. This procedure improves localisation stability and reduces the influence of depth outliers on visual servoing performance.

TABLE I. Model hyperparameters for base models. CNN: CNN output; LSTM/TF: hidden dim of LSTM or Transformer; nH,FF: attention heads / feedforward dimension; FC/Out: fully connected layers / output.

Model	CNN	LSTM/TF	nH	FF	FC/Out
VGG16-LSTM	4096	512	-	-	256→128→32→1
ResNet50-LSTM	2048	512	-	-	256→128→32→1
ConvNext-TF	256	256	4	512	128→1

2) *Tactile Perception Methods*: We compare three tactile perception models for hardness estimation. Each model combines (i) a CNN backbone for per-frame feature extraction from GelSight RGB images and (ii) a temporal encoder to aggregate information across the contact sequence.

We adopt a VGG16-LSTM architecture as a baseline, following early work on vision-based tactile perception [3]. The CNN encodes each tactile frame into a feature vector, which is then integrated over time by an LSTM to predict hardness. Motivated by evidence that residual networks improve representation quality in tactile tasks (e.g., slip detection and grasp outcome prediction), we replace VGG16 with a ResNet50 backbone while keeping the LSTM temporal module fixed [24], [25]. This isolates the effect of the CNN backbone on hardness estimation performance. Finally, we evaluate a non-recurrent temporal encoder by pairing a ConvNeXt backbone with a Transformer. In contrast to LSTMs, Transformers model long-range dependencies via self-attention over the entire sequence, potentially improving robustness to subtle temporal cues during pressing.

Each model takes as input a sequence of $T \in \{2, 4\}$ RGB tactile images resized to 224×224 . Table I summarises the principal hyperparameters. For LSTM-based models, we use three stacked LSTM layers.

To avoid model collapsing, which means that all predictions are mapped to the mean of the training set, three extra measures in model training were taken. First, dropout layers (factor 0.1 to 0.2) were added to the LSTM and fully connected layers. Additionally, the AdamW optimiser used a higher learning rate in later layers ($1e-3$) in comparison with early layers ($5e-5$). Finally, a custom loss function (Equ 1) was constructed based on the mean squared error (MSE) between the predictions (p) and the labels (l) and a penalty for low variability (Var) within the predictions:

$$\mathcal{L} = \text{MSE}(p, l) + 4 \min \left(\frac{1}{\text{Var}(p) + 10^{-6}}, 1000 \right) \quad (1)$$

3) *LLM*: For generating responses to user requests, we employed a DeepSeek-R1-Distill-Llama-70B model, accessed via Groq. The model was configured with a temperature of 0.1 to reduce hallucination while keeping natural flow.

The system was assigned the following role: “You turn object data into scene descriptions, explain and interpret tactile levels.” The prompt included 10 rules, covering different aspects for a more precise response. These included, but were not limited to: (i) object location descriptions translated from workspace coordinates via a rule-based mapping (e.g., left/center/right, front/center/back); (ii) ideal ripeness thresholds for bananas, limes, and lemons (defined empirically by comparing ripe and unripe fruits to reference hardness objects); and (iii) writing style guidance to ensure concise, fluent, and operator-friendly language.

IV. EXPERIMENTS

A. Dataset Description

1) *Visual Servoing Dataset*: We evaluate in tabletop scenes containing 1 to 5 fruits. The dataset includes 40 annotated instances (4 per each of 10 fruit types) with bounding boxes, segmentation masks, and centroids, enabling detection, segmentation, and centroid accuracy assessment.

2) *Tactile Perception Dataset*: We use three splits: (i) a pretraining set drawn from an online GelSight dataset of approximately 5,000 objects spanning the Shore 00 scale, where contact frames are detected at $\text{SSIM} < 0.90$ relative to a no-contact reference and 8-frame clips are extracted [3], [26]; (ii) a fine-tuning set collected with a marker-based GelSight-Mini mounted on a uFactory 850 robotic arm; and (iii) a validation set collected with the same setup. All collected sequences contain 8 images with 0.25 mm inter-frame steps. For our collected data, contact is defined by $\text{SSIM} \leq 0.96$ and mean marker displacement > 2 pixels [9], [27]. The difference in contact criteria across datasets reflects minor sensor and illumination differences.




Fig. 3. Data collection: images were compared to a reference image. If the contact criteria were met, 8 images were captured and transformed into a 2 or 4 image sequence.

The fine-tuning objects are five rubber cubes (66 to 80 HA), an elastic band (88 HA), and a glasses pouch (62 HA). This selection covers most of the expected fruit hardness range (60 to 90 HA). For each object, we record 40 poses by varying x - y position by ± 5 mm and yaw by 0 to 45° , yielding 280 samples in total. The validation set includes three fruit pairs (mango, lime, tomato) and two fruit trios (banana, avocado) at distinct ripeness stages. For each individual fruit, we collect 20 samples.

We convert each 8-frame contact sequence into shorter image sequences. For 2-frame sequences we use the 2nd and 8th frames; for 4-frame sequences we additionally include the 4th and 6th frames (Fig. 3). These selections capture early and peak deformation while maintaining coverage of

the contact trajectory. Finally, the difference between the selected images and the first contact image was used as input.

3) **LLM Dataset:** We construct an evaluation set of 100 prompts based on randomly generated scenes featuring 1 to 6 fruits. Each prompt references objects present in the validation tactile set so that language outputs can be grounded in measured hardness. Prompts span common intents, including locate-and-measure requests, pairwise or list-wise comparisons, ranking by hardness or ripeness, and filtering by attributes such as color or position. For each prompt, the expected output includes referenced object labels, per-object hardness estimates or intervals and an ordered list when applicable.

B. Experimental Setup

1) **Visual Servoing Setup:** Each scene is queried through the interface with the prompt “*I want to know the hardness of the [fruit]*”, which triggers detection and centroid computation as described in Sec. III-B.1.

2) **Tactile Perception Setup:** Backbones are initialised from ImageNet, then pretrained (80 epochs) on the online data and fine-tuned (15 epochs) on collected data.

All images were augmented by horizontal flips and colour jitter to increase model robustness to orientation variations and lighting conditions. The latter employed a brightness, contrast and saturation variation of 10% and hue by 1%. Furthermore, a cosine LR scheduler (patience 2, factor 0.2), weight decay of $1e-4$ and batch size of 8 were used.

3) **LLM Setup:** An LLM-as-a-judge is used to evaluate the prompts. Within this setup, a dedicated LLM model (Llama-4-Maverick-17b-128e-Instruct) is instructed to score the prompt from 1 to 5 based on three metrics [28].

TABLE II. Scenario complexity breakdown. [object] = chosen items, [property] = hardness, ripeness or softness. Complexity is determined by the number of objects, their distinctness, and whether fruits are explicit.

ID	Prompt	# Obj.	# Distinct	Explicit?	Complexity
1	Identify [property] of [object].	1	1	Yes	Low
2	Identify the most [property] [object] in scene.	2	1	Yes	Medium
3	Summarize the [property] of [object], [object] and [object].	3	3	Yes	Med-High
4	Summarize the [property] of all fruits in scene.	5	3-5	No	High

4) **Complete Framework Setup:** The complete framework is evaluated based on four scenarios varying in complexity [24], [29]. The complexity level is determined by three factors: 1) number of objects prompted, 2) number of distinct objects requested and 3) language reasoning. The latter one is determined by whether the fruits are explicitly mentioned in the request. Overall complexity is then categorised as low, medium and high, as shown in Table II.

C. Evaluation Metrics

1) **Visual Servoing Metrics:** Both models are evaluated using three metrics: confidence score, segmentation score (intersection over union or IoU) and the distance between the ideal midline of the fruit and the computed centroid. The ground-truth masks were manually annotated using

Roboflow. These three metrics were compared using an independent t-test with unequal variances ($\alpha = 0.01$), as Levene’s test revealed p-values below 0.01. The null hypothesis of these t-tests is that there is no difference in the means between YOLO and GSAM. Although normality (Shapiro-Wilk) was not met for all groups, we assume the central limit theorem holds true given the large sample size (40).

Additionally, the errors were compared using a one-sample t-test ($\alpha = 0.01$) to a threshold of 5mm to test if tactile perception would be valid. This threshold was indeed the deviation from the centre we allowed during data collection of the tactile models (see Sec. IV-A.2). Finally, also the success rate (SR) was noted. In this case, success is defined as the model hovering toward the correct fruit.

2) **Tactile Perception Metrics:** The test set for pretraining comprised 20% of the online data (N=962). For fine-tuning, the model run 7 times with a leave-one-out procedure across the object. In order to select the most optimal model after pretraining and fine-tuning, the root mean squared error (RMSE), coefficient of determination (R^2) and spearman correlation (ρ) were noted. The latter one reflects the model’s ability to keep ranks between the objects.

As it is difficult to determine the ground-truth Shore 00 value for fruits [3], a test was set up to determine if predictions followed the ranks. Since the Shapiro-Wilk test revealed non-normality in some data groups, a non-parametric Wilcoxon rank-sum test ($\alpha = 0.01$) was conducted to analyse whether, within one fruit sort, the median on the harder fruit was significantly higher than the softer fruit. As medians are compared, the interquartile range (IQR) will be given. In case of multiple comparisons (bananas and avocados), a Bonferroni-Holms correction was applied.

3) **LLM Metrics:** The LLM-as-a-judge is instructed to score the prompt from 1 to 5 based on three metrics: (i) Accuracy: are the objects, hardnesses and (relative) positions correctly described? (ii) Completeness: are all objects from the request mentioned and is info given if an object was not found? Is the ripeness interpreted for the correct cases? (iii) Clarity and Coherence: is the description understandable, concise and fluent?

4) **Complete Framework Metrics:** Each scenario is executed 10 times, with all fruit equally present across the runs. In these scenarios, tactility values will not be tested again. Contrarily, an object-level (OL-SR) and scenario-level success rate (SL-SR) will be used. The OL-SR is defined as the average percentage of fruits the model was able to accurately identify, measure and communicate through the LLM. The SL-SR is more restrictive: it is defined as the percentage of times the total scenario was correctly executed and communicated. For instance, if 1 of 5 objects is mislocated, the OL-SR would be 4/5 while SL-SR 0. Finally, the latencies are reported as the average among the succeeded trials per scenario. A breakdown in latencies will reveal which steps take the longest.

TABLE III. RMSE, R^2 and ρ for different models under various conditions. Lower RMSE, higher R^2 and Spearman ρ are better. Results are shown after pretraining (80 epochs on online data) and fine-tuning (15 epochs on collected data). ResNet50-LSTM3: model employing a ResNet50 as CNN backbone and 3 LSTM layers. The baseline models use 2 contact images.

Model	Pretraining			Fine-tuned		
	RMSE	R^2	ρ	RMSE	R^2	ρ
<i>Main Results</i>						
ResNet50-LSTM3	7.18	0.93	0.95	4.30	0.73	0.88
Transformer	6.87	0.93	0.94	6.23	0.44	0.77
VGG-LSTM3	27.98	-0.01	-0.03	20.99	-5.32	0.02
<i>Effect When Training on Half Scale</i>						
ResNet50-LSTM3	9.11	0.63	0.77	5.01	0.64	0.86
Transformer	9.03	0.65	0.81	4.33	0.73	0.89
<i>Effect of Different ResNet Backbone</i>						
ResNet34-LSTM3	6.83	0.94	0.96	6.50	0.41	0.89
ResNet101-LSTM3	7.46	0.92	0.95	7.79	0.16	0.86
<i>Effect of More Contact Images (4)</i>						
ResNet50-LSTM3	7.13	0.93	0.96	8.80	-0.11	0.73
Transformer	6.86	0.93	0.95	6.27	0.44	0.76
<i>Effect of LSTM Depth (1 layer)</i>						
ResNet50-LSTM1	7.43	0.93	0.95	10.63	-0.65	0.78
<i>Effect of Direct Training</i>						
ResNet50-LSTM3	-	-	-	9.40	0.05	0.21

D. Main Results and Analysis

1) **Visual Servoing:** In case the objects were identified, the independent t-test revealed that the confidence score of the YOLO model (0.921, 95% CI:[0.886, 0.956]) was significantly higher ($t=12.84$, $p < 0.01$, $n_1=39$, $n_2=36$) than the score for the GSAM model (0.645, [0.606, 0.685]). This is reflected in a SR of 0.9 for YOLO and 0.85 for GSAM. In the other cases, either the object was not found or another object had a higher confidence score than the target fruit.

However, in the correct cases, the segmentation score (IoU) of GSAM (0.942, [0.927, 0.956]) was significantly higher ($t=9.01$, $p < 0.01$, $n_1=39$, $n_2=36$) than the score from YOLO (0.786, [0.752, 0.820]). Although the errors for YOLO (7.194mm, [5.534, 8.855]) were statistically not different ($t=1.39$, $p=0.17$, $n_1=39$, $n_2=36$) than GSAM (5.645mm, [4.314, 6.981]), only GSAM succeeded in having an error not statistically different than 5mm ($t=1.35$, $p=0.19$, $n = 34$). In contrast, YOLO noted a statistically higher error than 5mm ($t=2.80$, $p < 0.01$, $n=36$). This alignment with the tactile training setup makes the GSAM model more reliable for integration with the tactile perception model in TactEx.

2) **Tactile Perception:** The main results in Table III reveal that the VGG16-LSTM was very prone to model collapsing, despite the extra measures taken. More specifically, the ResNet50-LSTM3 and Transformer baseline models in Table III avoided model collapsing, thereby indicating that the applied measures were effective. The ResNet50-LSTM3 model achieved the best performance (lowest RMSE, highest R^2 and ρ) of the three models, both after pretraining and fine-tuning. The ResNet50-LSTM3 results after pretraining and fine-tuning are visualized in Fig. 6. Due to its superior performance among the primary results, the ResNet50-LSTM3 was selected for validation in the fruit ranking scenario.

The validation results are reported in Table IV. The Wilcoxon rank-sum tests showed that all comparisons were statistically significant ($p < 0.01$). This demonstrates that the model can correctly interpret which fruit is harder, thereby mimicking human touch. This statistical significance

TABLE IV. Median and interquartile range of predictions (20 samples per individual fruit) from the full-range trained ResNet50-LSTM3 on different fruit pairs and trios. Hard, Medium, and Soft correspond to the empirical ripeness stages. Wilcoxon rank-sum test indicates whether the harder fruit is significantly ($p < 0.01$) harder than the softer fruits.

Condition	Median	25th	75th	Wilcoxon
<i>Mango</i>				
Hard (1)	79.47	72.89	84.59	$u = 343$, $p < 0.01$
Soft (0)	67.75	65.60	72.78	
<i>Lime</i>				
Hard (1)	64.13	63.78	64.64	$u = 285$, $p = 0.011$
Soft (0)	63.84	63.73	63.93	
<i>Tomato</i>				
Hard (1)	71.02	65.69	79.91	$u = 308$, $p < 0.01$
Soft (0)	64.14	63.13	65.98	
<i>Banana</i>				
Hard (2)	72.63	67.50	82.53	2 vs 1: $u = 288$, $p < 0.01$
Medium (1)	66.87	66.31	67.62	
Soft (0)	63.05	62.85	63.89	2 vs 0: $u = 368$, $p < 0.01$
<i>Avocado</i>				
Hard (2)	65.25	64.02	65.92	2 vs 1: $u = 299$, $p < 0.01$
Medium (1)	63.54	63.39	64.12	
Soft (0)	61.73	60.97	62.15	2 vs 0: $u = 373$, $p < 0.01$

is something not earlier discovered in literature [5], [16]. We dedicate this result to the change in backbone, training strategy and study of the optimal architecture, as explained in the ablation study (section IV-E).

3) **LLM:** The LLM-as-a-judge scores on 5 reveal a solid performance of the LLM answer: the accuracy was 4.19 ± 0.59 , completeness 4.94 ± 0.24 and conciseness and clarity 4.92 ± 0.44 .

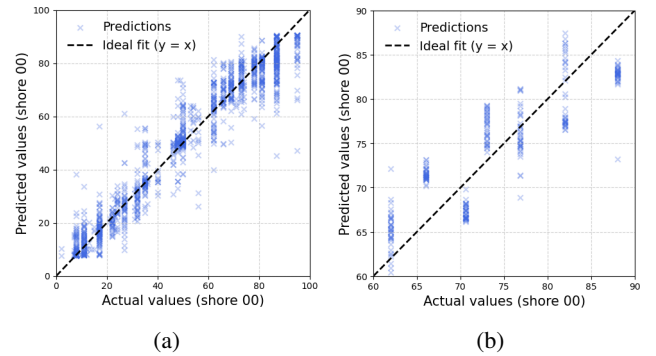


Fig. 6. Results of tactile predictions from the main ResNet50-LSTM3 model after pretraining (a) and fine-tuning (b). This is the model that will eventually be implemented within TactEx.

4) **Integrated Framework:** The success rate are visualized in Fig. 4. The mismatch between SL-SR and OL-SR in Fig. 4 reveals that both models have problems with certain fruits, more specifically a kiwi. This strongly impacts the SL-SR. The effect was further amplified in more complex scenarios where kiwis appeared more frequently. Indeed, discarding the kiwi cases improved the performance dramatically. For YOLO, this likely reflects the low amount of kiwi images (15) and too little training (100 epochs). We encourage future researchers to include at least 25 images per class to avoid these class imbalance issues and train for 150 epochs.

The decrease across higher complexity is consistent with literature [29]. The SL-SR in Fig. 4 demonstrates that when the fruits are explicitly stated, strong performance is achieved. This shows the model's ability to select the hardest fruit on the table. In the case of a banana, lime and lemon, the




Fig. 4. Success rates of the TactEx framework across four interaction scenarios of increasing complexity (Sc1-Sc4), as defined by Table II. SL-SR: Scenario-Level Success Rate, OL-SR: Object-Level Success Rate, w/o: without, Sc: Scenario.




Fig. 5. Latency breakdown across the different scenarios (Sc1-Sc4), as defined by Table II. Tactile collection takes the longest in both models, while GSAM also suffers from heavy computational delays during centroid computation. Sc: Scenario, Coll: Collection.

LLM interprets these hardships as ripe or unripe with high accuracy, making it practical in household applications. In the fourth scenario however, the SL-SR drops significantly, even without kiwis. This is explained by the fact that the prompt did not mention the fruits explicitly. In contrary, a list of 20 fruits and vegetables was prompted in GSAM. However, without defined fruits of interest, we cannot use the NLP logic to filter out mistakes. This ultimately leads to more misdetections, with GSAM more affected than YOLO, as shown by the OL-SR. The latter is explained by the lower confidence score reported earlier, stemming from the fact that it is not specifically trained on these targets, unlike YOLO.

In general, however, GSAM outperforms YOLO by 10 percentage point. This is explained by the lower error and better segmentation reported earlier, which makes tactile prediction more reliable as it adheres more to the tactile training setup. This further illustrates the need for longer training and more data in YOLO. In contrast, GSAM is more versatile and deployable into new applications compared to YOLO as it does not require any training. The combination of GSAM and tactile perception makes TactEx an optimal solution for household, agricultural, and industrial applications.

E. Ablation Study

1) *Ablation Results:* We conducted an ablation study to assess key design choices: number of contact images (2 vs. 4), LSTM depth (1 vs. 3), and CNN backbone (ResNet34, ResNet50, ResNet101). We also compared direct training on the collected samples against the pretraining strategy, and examined pretraining restricted to the upper half of the Shore 00 scale (> 39) to test its effect after fine-tuning. The results of these variants are summarized in Table III.

2) *Analysis:* Table III suggests a new effective training strategy for future hardness estimation models. Four findings support this statement. Firstly, the directly trained ResNet50-LSTM3 performed poorly (RMSE 9.40) compared to the

pretraining strategy (RMSE 4.30). Second, while the pre-trained main models achieve slightly higher RMSE than related works (6.87 vs 5.18), the R^2 values are consistent with state-of-the-art models [3], [9]. Third, the fine-tuned models demonstrate that even with a low amount of collected data the model seems to generalize well to a new robot setup. Indeed, the spearman correlations remain high, with only minor inconsistencies visible in Fig. 6. This suggests that the model is great at ranking hardness values, with minimal deviations from the true values (RMSE 4.30). Fourth and finally, none of the visualized models in Fig. 6 reveal systematic degradation for harder objects, suggesting that the Gelsight sensor is better suited than the TacTip used by Nam et al. (2024) [9].

With regard to the optimal architecture, the ablation results were influenced by the glass pouch object. Nevertheless, we consider the conclusions from the ablation study as valid, as they are mostly based on the spearman correlation. Three design choices emerged as most effective: (1) four contact images do not improve the model after pretraining and fine-tuning suggesting that earlier research by Nam et al. (2024) was right in that the first and last contact image are most crucial [9], (2) three LSTM layers better capture the subtle variation between contact images and (3) while we believe the other ResNet backbones also work for our use case ($\rho \geq 0.80$), the ResNet50 seems to balance complexity and robustness more optimally (lower RMSE and higher R^2).

Considering the pretraining range, results are mixed: the Transformer model benefited from pretraining on half the range ($\rho = 0.89$ vs. 0.77), whereas the ResNet50-LSTM3 model did not ($\rho = 0.86$ vs. 0.88). This is likely due to the Transformer’s attention mechanism, which may emphasize localized patterns when the training range is narrower. Therefore, our recommendation is as follows: if the application of interest has a predefined target range, it is worthy to investigate whether pretraining on that range may help.

V. CONCLUSION AND FUTURE WORK

We presented **TactEx**, an explainable multimodal robotic framework that integrates vision, language grounding, tactile exploration, and robotic control to estimate and communicate object hardness in natural language. Across the evaluated scenarios, the ResNet50-LSTM3 architecture demonstrated reliable tactile hardness regression, preserving expected hardness rankings with statistically significant differences under non-parametric testing. These results support the feasibility of interactive, language-driven fruit ripeness assessment using vision-guided tactile sensing. The modular integration of GSAM-based object grounding, tactile inference, and LLM-based explanation enables component-level replacement and adaptation to new sensors, object categories, and tasks without redesigning the overall system. Future work should explore tighter perception-action coupling, including pre-processing visual scenes prior to user prompts, incremental scene updating, and closed-loop tactile exploration policies that adapt probing online.

REFERENCES

- [1] Z. Su, J. A. Fishel, T. Yamamoto, and G. E. Loeb, "Use of tactile feedback to control exploratory movements to characterize object compliance," *Frontiers in neurobotics*, vol. 6, p. 7, 2012.
- [2] D. Zhang, W. Fan, J. Lin, H. Li, Q. Cong, W. Liu, N. F. Lepora, and S. Luo, "Design and benchmarking of a multimodality sensor for robotic manipulation with gan-based cross-modality interpretation," *IEEE Transactions on Robotics*, vol. 41, pp. 1278–1295, 2025.
- [3] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson, "Shape-independent hardness estimation using deep learning and a gelsight tactile sensor," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2017, p. 951–958. [Online]. Available: <http://dx.doi.org/10.1109/ICRA.2017.7989116>
- [4] J. Li, S. Dong, and E. Adelson, "Slip detection with combined tactile and visual information," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7772–7777.
- [5] Z. Liao, Y. Du, J. Duan, H. Liang, and M. Y. Wang, "Quantitative hardness assessment with vision-based tactile sensing for fruit classification and grasping," 2025. [Online]. Available: <https://arxiv.org/abs/2505.05725>
- [6] W. Fan, H. Li, and D. Zhang, "Crystaltac: Vision-based tactile sensor family fabricated via rapid monolithic manufacturing," *Cyborg and Bionic Systems*, vol. 6, p. 0231, 2025.
- [7] W. Fan, M. Yang, Y. Xing, N. F. Lepora, and D. Zhang, "Tac-vgnn: A voronoi graph neural network for pose-based tactile servoing," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 10 373–10 379.
- [8] L. Fu, G. Datta, H. Huang, W. C.-H. Panitch, J. Drake, J. Ortiz, M. Mukadam, M. Lambeta, R. Calandra, and K. Goldberg, "A touch, vision, and language dataset for multimodal alignment," in *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [9] S. Nam, T. Jack, L. Y. Lee, and N. F. Lepora, "Softness prediction with a soft biomimetic optical tactile sensor," in *2024 IEEE 7th International Conference on Soft Robotics, RoboSoft 2024*. Institute of Electrical and Electronics Engineers Inc., 2024, pp. 121–126.
- [10] J. Chen, A. Kshirsagar, F. Heller, M. G. Andreu, B. Belousov, T. Schneider, L. P. Y. Lin, K. Doerschner, K. Drawing, and J. Peters, "Investigating active sampling for hardness classification with vision-based tactile sensors," 2025. [Online]. Available: <https://arxiv.org/abs/2505.13231>
- [11] W. Fan, H. Bo, Y. Lin, Y. Xing, W. Liu, N. Lepora, and D. Zhang, "Graph neural networks for interpretable tactile sensing," in *2022 27th International Conference on Automation and Computing (ICAC)*. IEEE, 2022, pp. 1–6.
- [12] D. Zhang, Q. Li, Y. Zheng, L. Wei, D. Zhang, and Z. Zhang, "Explainable hierarchical imitation learning for robotic drink pouring," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 3871–3887, 2021.
- [13] A. Adebayo, O. Ajayi, N. Chukwurah, and A. Pub, "Explainable ai in robotics: A critical review and implementation strategies for transparent decision-making," *Journal Of Multidisciplinary Research*, vol. 05, pp. 26–32, 03 2024.
- [14] R. Setchi, M. B. Dehkordi, and J. S. Khan, "Explainable robotics in human-robot interactions," *Procedia Computer Science*, vol. 176, pp. 3057–3066, 2020, knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 24th International Conference KES2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1877050920321001>
- [15] Z. Guo, H. Chen, X. Mai, Q. Qiu, G. Ma, Z. Kappassov, Q. Li, and N. Chen, "Robotic perception with a large tactile-vision-language model for physical property inference," 2025. [Online]. Available: <https://arxiv.org/abs/2506.19303>
- [16] W. Yuan, M. A. Srinivasan, and E. H. Adelson, "Estimating object hardness with a gelsight touch sensor," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 208–215.
- [17] S. Ueda, A. Hashimoto, M. Hamaya, K. Tanaka, and H. Saito, "Visuo-tactile zero-shot object recognition with vision-language model," in *International Conference on Intelligent Robots and Systems*. IEEE/RSJ, 9 2024. [Online]. Available: <http://arxiv.org/abs/2409.09276>
- [18] X. Zhao, M. Li, C. Weber, M. B. Hafez, and S. Wermter, "Chat with the environment: Interactive multimodal perception using large language models," in *IEEE International Conference on Intelligent Robots and Systems*. Institute of Electrical and Electronics Engineers Inc., 2023, pp. 3590–3596.
- [19] J. Huang, S. Wang, F. Lin, Y. Hu, C. Wen, and Y. Gao, "Tactile-vla: Unlocking vision-language-action model's physical knowledge for tactile generalization," 2025. [Online]. Available: <https://arxiv.org/abs/2507.09160>
- [20] J. Yu, H. Liu, Q. Yu, J. Ren, C. Hao, H. Ding, G. Huang, G. Huang, Y. Song, P. Cai, C. Lu, and W. Zhang, "Forcevla: Enhancing vla models with a force-aware moe for contact-rich manipulation," 2025. [Online]. Available: <https://arxiv.org/abs/2505.22159>
- [21] C. Zhang, P. Hao, X. Cao, X. Hao, S. Cui, and S. Wang, "Vtla: Vision-tactile-language-action model with preference learning for insertion manipulation," 2025. [Online]. Available: <https://arxiv.org/abs/2505.09577>
- [22] T. Ren, S. Liu, A. Zeng, J. Ling, H. Cao, K. Li, J. Chen, X. Huang, F. Yan, and Y. Chen, "Grounded sam: Assembling open-world models for diverse visual tasks international digital economy academy (idea) & community," International Digital Economy Academy (IDEA) & Community, Tech. Rep., 1 2024. [Online]. Available: <https://github.com/IDEA-Research/Grounded-Segment-Anything>
- [23] Q. Bai, S. Li, J. Yang, Q. Song, Z. Li, and X. Zhang, "Object detection recognition and robot grasping based on machine learning: A survey," *IEEE Access*, vol. 8, pp. 181 855–181 879, 2020.
- [24] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine, "The feeling of success: Does touch sensing help predict grasp outcomes?" 2025. [Online]. Available: <https://arxiv.org/abs/1710.05512>
- [25] J. Gao, Z. Huang, Z. Tang, H. Song, and W. Liang, "Visuo-tactile-based slip detection using a multi-scale temporal convolution network," 2023. [Online]. Available: <https://arxiv.org/abs/2302.13564>
- [26] X. Zhang, T. Yang, D. Zhang, and N. F. Lepora, "Tactpalm: A soft gripper with a biomimetic optical tactile palm for stable precise grasping," *IEEE Sensors Journal*, 2024.
- [27] J.-i. Lee, S. Lee, H.-M. Oh, B. R. Cho, K.-H. Seo, and M. Y. Kim, "3d contact position estimation of image-based areal soft tactile sensor with printed array markers and image sensors," *Sensors*, vol. 20, no. 13, 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/13/3796>
- [28] J. Gu, X. Jiang, Z. Shi, H. Tan, X. Zhai, C. Xu, W. Li, Y. Shen, S. Ma, H. Liu, S. Wang, K. Zhang, Y. Wang, W. Gao, L. Ni, and J. Guo, "A survey on llm-as-a-judge," 2025. [Online]. Available: <https://arxiv.org/abs/2411.15594>
- [29] L. Verbaan, Y. B. Eisma, and R. van Leeuwen, "Perception and control with large language models in robotic manipulation developing and assessing an integrated large language model system on environmental and task complexity," Ph.D. dissertation, Delft University of Technology, 8 2024.