PressureVision: Estimating Hand Pressure from a Single RGB Image

Patrick Grady¹⁽⁶⁾, Chengcheng Tang²⁽⁶⁾, Samarth Brahmbhatt³⁽⁶⁾, Christopher D. Twigg²⁽⁶⁾, Chengde Wan²⁽⁶⁾, James Hays¹⁽⁶⁾, and Charles C. Kemp¹⁽⁶⁾

¹ Georgia Institute of Technology ² Meta Reality Labs ³ Intel Labs

Abstract. People often interact with their surroundings by applying pressure with their hands. While hand pressure can be measured by placing pressure sensors between the hand and the environment, doing so can alter contact mechanics, interfere with human tactile perception, require costly sensors, and scale poorly to large environments. We explore the possibility of using a conventional RGB camera to infer hand pressure, enabling machine perception of hand pressure from uninstrumented hands and surfaces. The central insight is that the application of pressure by a hand results in informative appearance changes. Hands share biomechanical properties that result in similar observable phenomena, such as soft-tissue deformation, blood distribution, hand pose, and cast shadows. We collected videos of 36 participants with diverse skin tone applying pressure to an instrumented planar surface. We then trained a deep model (PressureVisionNet) to infer a pressure image from a single RGB image. Our model infers pressure for participants outside of the training data and outperforms baselines. We also show that the output of our model depends on the appearance of the hand and cast shadows near contact regions. Overall, our results suggest the appearance of a previously unobserved human hand can be used to accurately infer applied pressure. Data, code, and models are available online¹.

1 Introduction

Humans often interact with their surroundings by applying pressure with their hands. Given the importance of hand pressure, methods that enable the machine perception of this quantity could have broad applications. Traditionally, measuring the pressure a hand exerts has been accomplished with physical sensors that sit between the hand and contact surface. This includes sensors worn on the hand, such as pressure-sensitive gloves, and sensors mounted to the environment, such as arrays of pressure sensors.

While physical sensors are accurate and robust, they have drawbacks. Sensors between the hand and the environment alter surface properties, changing

 $^{^{1}}$ https://github.com/facebookresearch/pressurevision



Fig. 1. PressureVisionNet takes an RGB image as input and outputs a pressure image with estimates of pressure applied by the hand to a planar surface. Each pixel of the pressure image is an estimate of hand pressure for the corresponding pixel in the RGB image. Black, purple, and yellow represent zero, low, and high pressure.

appearance and interfering with contact mechanics relevant to human manipulation and tactile sensing. Sensors attached to hands can also be uncomfortable. Sensors mounted to environments require large numbers of sensing elements to cover modest areas, such tabletops, at high resolution and can be difficult to apply to varied surfaces.

Cameras have the potential to economically cover surfaces with virtual pressure sensors at high spatial resolution and do so without requiring the application of cumbersome instrumentation. Much like markerless pose estimation has enabled new applications by inferring human kinematics from RGB images, *visual hand pressure estimation* has the potential to be widely applied. For example, modern augmented and virtual reality headsets have cameras that could be used to perceive hand pressure, turning a wall into a giant touchscreen, or a tabletop into piano keyboard.

As a first step towards general visual hand pressure estimation, we investigate the feasibility of estimating hand pressure from RGB images under controlled conditions. To the best of our knowledge, our work is the first to demonstrate the feasibility of this approach. In particular, we address two critical questions: 1) Can appearance-based inference estimate hand pressure? 2) Can visual hand pressure inference generalize to previously unseen hands from a diverse population? We use a planar pressure sensing array to measure ground truth pressure during natural interactions from 36 participants with diverse skin tones.

We developed PressureVisionNet to infer hand pressure from a single RGB image. PressureVisionNet performed well with unseen participants, providing evidence that the visual signals used for pressure estimation are shared across people. Our sensitivity analysis also indicates that PressureVisionNet relies on the appearance of the hand close to regions of contact, where we expect applied pressure to most influence hand appearance. Similarly, PressureVisionNet performs poorly at estimating pressure for visually occluded parts of the hand. We also provide a demonstration of PressureVisionNet inferring plausible pressure given conditions outside of the training data. Our demonstrations involve a



Fig. 2. Fingers display visible cues indicative of applied pressure. In the no-contact case, shadows are diffuse and the finger can be less sharp due to motion blur. As the finger makes contact, the skin at the fingertip loses color, the distal joint hyperextends, the texture of the skin changes, and the finger pad expands.

smartphone camera observing an uninstrumented tabletop with ambient office lighting being contacted by people outside of the training set.

Our central insight is that hands display subtle visual cues that indicate the presence of pressure. As the hand applies pressure to a surface, the surface applies equal and opposite pressure to the hand. As such, the hand itself can serve as a physical probe that changes in predictable ways when pressure is applied to it, and the camera can observe this probe from a distance to infer the associated pressure. The shape of the hand can also result in informative shadows cast on the surface.

Hands are complex, non-rigid appendages with multiple types of soft tissues and fluids surrounding an interior skeleton that is often modeled as having 27 degrees of freedom. Pressure applied to the hand deforms tissues, moves fluids, and changes the configuration of the joints. The active application of pressure also involves contraction of hand muscles and tensile force applied by muscles in the forearm via tendons. These biomechanical processes result in changes to the appearance of the hand in visible light, including changes to surface geometry, color, and texture (Figure 2).

In summary, our paper makes the following contributions:

- We propose a novel task, visual hand pressure estimation.
- We present PressureVisionNet and show that it can infer hand pressure from a single RGB image and generalize to new people.
- We release PressureVisionDB, a dataset of 36 participants with paired pressure and image data.
- We release our trained models and code.

2 Related Work

Physical Sensors for Force Measurement Sensors to directly measure the force that a hand exerts generally fall into two categories: sensors on the hand and sensors on the surface being touched. Gloves instrumented with flexible

force sensors can directly measure the force that the hand exerts [7,54]. While sensors on the hand work for a wide range of objects, they can inhibit natural manipulation, reduce the person's tactile sensation, and often do not cover the entire hand. Commercially available systems are also expensive [46,56].

Alternatively, objects can be instrumented to capture force information [2,6,39,45]. While these methods can accurately measure pressure, they change the properties of the surfaces they cover, can be challenging to mount on curved surfaces, and scale poorly to larger surfaces.

A variety of fingertip pressure sensors have been developed for robotic grippers. These internal sensors sometimes observe the deformation of a soft exterior [58,65] or monitor an internal fluid [59].

Inferring Contact and Force from Vision One class of methods uses physics to infer the contact forces given the contact points and object trajectory. Methods proposed to determine the contact points include neural networks [17,34] or by combining markerless tracking of the hands with mesh-object intersection [44,45,47]. These methods complement our own because they can infer contact that is occluded or out-of-view, but methods that rely on the accelerations of hand-held objects cannot generalize to interactions with relatively immobile surfaces like tables.

Additional work studies force application in general human-environment interactions such as sitting and standing [52,68,69]. Clever *et al.* use depth data to regress the amount of pressure between a human at rest and a mattress [12]. They use a neural network to estimate the pose of a body on the bed in addition to pressure.

Predicting Pressure using Hand and Surface Appearance As the hand applies force, blood in the surrounding tissue is displaced. This effect is visible at the fingertip and underneath the fingernail, where a whitening of the tissue can be observed. Various techniques have been proposed to estimate fingertip force using optical sensors focused on this effect [9,36,37].

The soft tissue in the palm and pads of the fingers deforms under applied force [43]. This deformation is accompanied by expansion in other areas, often visible as a widening of the pads of the fingers. Hwang *et al.* [30] use surface deformation to infer contact force, but support only unoccluded point contacts. Johnson *et al.* [31] use the deformation of the hand to perform a high-resolution reconstruction of the object surface.

Cast shadows provide important cues in human perception of depth in a 3D scene [26,27,29]. Researchers have used visual observations of shadows for closed-loop control of a robot [18].

It has additionally been demonstrated that videos of faces can be used to identify if participants are squeezing an object tightly [1].

Amplifying Imperceptible Visual Cues Various techniques have been presented to extract subtle cues from images or video which are typically not noticed by the human eye. Wu *et al.* present Eulerian Video Magnification [60]. They magnify periodic signals to make them clearly visible, including minute changes in skin color due to bloodflow. Davis *et al.* [14] are able to amplify tiny subpixel vibrations in videos to reconstruct audio from visual data.

Contact for Grasping and Pose Estimation Contact between hands and objects is important for grasping and manipulation. Estimating where contact should occur is often a first step for planning a robotic grasp [4,10,50,51] or creating anthropomorphic animation [32,64,67]. For pose estimation during hand-object interaction [19,21], using contact to enforce consistency between hands and objects improves pose plausibility and accuracy [20,22,23,48,57]. In a broader context, contact is used when reasoning about people interacting with the surrounding environment [11,25,40,53,68].

Contact for Human-Computer Interaction Detecting contact with ordinary surfaces can be useful for building human-computer interfaces. MR-Touch [61] uses depth data to identify touches on flat surfaces for the Microsoft HoloLens. The method first fits a 3D plane to the wall or tabletop surface and thresholds the distance between the plane and the fingertip. TapID [38] uses subtle accelerations from a pair of IMU sensors on the wrist to detect contact between fingers and surfaces. TouchAnywhere [41] detects touch via heuristics by estimating the intersection between finger and its shadow, but only detects binary contact and does not include pressure ground truth.

Table 1. Datasets for contact between hands and surfaces: GRAB [55] infers dynamic contact via geometric models and marker-based motion capture. TactileGloves [54] measures dynamic pressure with a sensorized glove. ContactDB [3] and ContactPose [5] measure static contact via thermal imaging. Ours, PressureVisionDB, measures pressure with a sensorized plane and provides registered and synchronized RGB images.

	Diverse Surfaces	Bare Hands	Images	Pressure
GRAB [55]	 ✓ 	×	×	×
TactileGloves [54]	 ✓ 	×	×	 ✓
ContactDB [3]	 ✓ 	~	×	×
ContactPose [5]	 ✓ 	~	 	×
PressureVisionDB (Ours)	×	 ✓ 	~	 ✓

Hand-Object Contact Datasets GRAB [55] uses optical motion capture to capture hand and object pose and indirectly infers contact using mesh proximity. ContactDB [3] measures high-resolution contact on a variety of objects. After participants grasp an object, the thermal imprint from the hand is measured from the object surface using a thermal camera. ContactPose [5] adds paired images,

object poses, and hand poses to ContactDB. ContactDB and ContactPose do not provide pressure measurements and only provide contact for static grasps.

As shown in Table 1, our dataset is unique compared with these datasets, providing RGB images paired with high-resolution, dynamic pressure images.

3 The PressureVisionDB Dataset

We collected PressureVisionDB, a novel dataset of bare hands interacting with a pressure-sensitive object. The dataset includes 16 hours of data collected from 36 participants with pressure data and RGB images from four synchronized cameras, totalling 64 hours of RGB video data.

3.1 The Capture Setup

We assemble a rigid frame to which we attached cameras, lights, and a pressure sensing surface (Figure 3a). Participants reached through one side of the cube-shaped frame to place their hands on the planar sensing surface.



Fig. 3. a) The data capture rig used to collect PressureVisionDB. Participants reach into the front opening to make contact with the elevated pressure sensor. b) The sensorized plane records a pressure image, which is projected into image space for training and evaluation. Pressure is visualized as bright blue regions. c) PressureVisionDB includes participants with diverse skin tones, ages, and genders under varied lighting conditions.

Pressure Sensor To measure hand pressure, we used a Sensel Morph [39] (Figure 3b). The sensor is a planar surface featuring a grid of force-sensitive resistor (FSR) pixels. The sensor produces a 185x105 "pressure image", with each pixel having a pitch of 1.25mm. The sensor produces diffuse readings below

0.5 kPa, which we use as the minimum effective pressure, and the 99th percentile pressure recorded in the dataset is 82 kPa.

We covered the sensor surface with a white vinyl overlay for most participants, and covered it with a wood-textured overlay for seven participants. We mounted the sensor above the table surface to allow participants to reach underneath the sensor for pinching and grasping actions (Figure 3a). PressureVisionDB only measures pressure across this single planar surface, and we leave the capture of hand pressure on more diverse shapes to future work. Additionally, the sensor measures pressure normal to the plane, so we only consider this component of pressure.

Cameras The capture setup uses four synchronized and calibrated OptiTrack Prime Color cameras to capture 1080p RGB frames. We mounted the cameras to provide overhead views of the scene along with three OptiTrack eStrobe light sources. Each light can be turned on or off to achieve 8 different lighting conditions. The lighting condition was changed after each participant.

Participants PressureVisionDB was collected from a diverse set of adults (Figure 4) with a range of skin tones, ages, and genders (Figure 3c). Participants self-reported their age and gender, and their skin tone were measured objectively. A Pantone RM200 colorimeter was used to classify the participant's skin tone in the Pantone SkinTone Guide [42]. These skin tones were then divided into three categories: light, medium, and dark.

Protocol We developed a protocol to capture a variety of hand-surface interactions. Participants were asked to perform 36 different actions with one hand and then the other. Prior to each action, participants were shown a text description and an image of the action. The list of actions includes pressing with a single finger (index, thumb, or other), pressing with the whole hand, applying tangential force, grasping the edge of the sensor, and drawing with a finger. For many actions, participants were prompted to use one of three force levels: high force, low force, and no contact. For the no contact condition, participants moved their hand very close to the sensor, as if they intended to touch the sensor, but did not make contact.

Ethics Participants gave informed consent and were compensated for their time. The data was captured by a third-party corporation that specializes in data collection with human participants. The capture rig was designed so that no images of the participants' torsos or faces were collected. The third-party did not provide personally identifiable information to the research team.

At the inception of our project, we were concerned about the potential for darker skin tones to reduce performance [13,33,8]. Melanin and hemoglobin both influence skin color in visible light and melanin concentration varies widely across people [70]. As such, we prioritized recruiting participants with a wide range of skin tones (Figure 3c).



Fig. 4. PressureVisionDB includes data from participants with diverse skin tones.

4 Estimating Pressure and Contact from RGB

We designed a deep network to infer hand contact and pressure. From an input RGB image, the network estimates a pressure image of the same size. The pressure image provides a pressure estimate for each pixel of the RGB image.

4.1 Network Architecture

We designed PressureVisionNet to estimate the location and magnitude of pressure between a hand and surface. The network is an encoder-decoder architecture which inputs a single RGB image, I, and outputs a pressure map, $\hat{P} = f(I)$. During training and evaluation, the ground truth pressure data from the sensor is projected into image space (Figure 3b). Consequently, the network estimates hand pressure for each pixel of the input image, which results in an output pressure image that is the same size as the input image.

Pressure estimation is treated as a classification problem. The pressure range is divided into nine bins placed in logarithmic space [15]. The network infers the pressure bin for each pixel of the output pressure image. We experimented with direct regression of pressure scalars, but found that this was outperformed by the classification approach.

We cropped images from each camera to include a margin around the pressure sensor and resized the images to 480x384 pixels. PressureNet uses a SE-ResNeXt50 [24,28,62] encoder, with weights from pretraining on ImageNet [16]. A feature pyramid network (FPN) [35,63] decoder produces the output image. A cross-entropy loss was used during training. During inference, the network runs at 53 FPS using an RTX 3090 GPU.

4.2 Evaluation Metrics

We considered four types of contact and pressure metrics. Contact is a binary quantity that describes if the hand and sensor are touching, while pressure is a scalar describing how hard the hand and sensor are pressing against each other. Contact maps C are generated by thresholding the pressure map P at a low value, $P_{th} = 1.0$ kPa. Values greater than this are marked as *in contact*.



Fig. 5. Results from PressureVisionNet on held-out participants. Some images have been magnified for clarity. PressureVisionNet can accurately infer the location and magnitude of pressure from an RGB image. PressureVisionNet correctly identifies the lack of contact in "no-contact" actions (right, top two). The bottom row illustrates a common failure mode: no pressure is estimated for parts of the hand that are occluded.

- **Temporal Accuracy** To evaluate the *temporal* accuracy with which the onset and termination of contact are estimated, if *any* contact is present in the estimated and ground truth contact maps, \hat{C} and C, the frame is marked as in contact. A frame is marked correct if the presence of contact is consistent in estimated and ground truth frames.
- Contact IoU To evaluate the *spatial* and *temporal* accuracy of estimated contact, we computed the intersection over union (IoU) between the the binary contact images. This metric does not consider the magnitude of the estimated pressure, and is an upper bound on Volumetric IoU.
- Volumetric IoU We propose the Volumetric IoU (Figure 6), a novel metric that extends Contact IoU to evaluate the *magnitudes* of pressure estimates in addition to their *spatial* and *temporal* accuracy. Each 2D pressure image is converted into a 3D "pressure volume", where the height of the volume is equal to the amount of pressure at that pixel. The Volumetric IoU can be calculated as:

$$IoU_{vol} = \frac{\sum_{i,j}^{i,j} \min(P_{i,j}, \hat{P}_{i,j})}{\sum_{i,j}^{i,j} \max(P_{i,j}, \hat{P}_{i,j})}$$
(1)

 Mean Absolute Error To evaluate the accuracy of estimated pressure in physical units, we calculate mean absolute error (MAE) over each pixel. As most of the dataset pressure images consist of zeros, these numbers are close to zero.



Fig. 6. Volumetric IoU quantifies the agreement between two pressure images by treating them as pressure volumes.

4.3 Dataset Splits

The dataset consists of 36 participants. Prior to using the dataset to develop PressureVisionNet, we selected 6 participants spanning skin tone and demographics for a held-out test set. We used data from the remaining 30 participants for training and validation.

5 Results

Our primary goal was to investigate the potential to infer hand pressure from a single RGB image. In addition to characterizing overall performance, we focused on three questions. First, we considered how well performance can generalize to people outside of the training set, since feasibility depends on the existence of shared properties across hands and human behavior. Second, we considered whether the appearance of the hand and its cast shadows were used for inference. Third, we considered whether a trained model can produce reasonable estimates for images captured in an unseen environment.



Fig. 7. a) A 3D pose estimator is used to estimate hand pose. b) As monocular pose estimators have difficulty in estimating true hand scale, hand scale is swept through a range of values (selected scales are visualized, optimal scale is colored). c) Ground truth contact is compared to contact calculated by mesh intersection.

11

5.1 Baseline Models

Since inferring hand pressure from an RGB image is a new task, there were no existing methods available for direct comparison. To provide context for our numeric results, we created a baseline model for comparison.

3D models of the human body have been used to infer contact and pressure in other contexts [55,20,12]. Our use of a single RGB image makes this approach more challenging. Prior work used multiple cameras, depth cameras, and motion capture systems to obtain high fidelity 3D models. We used FrankMocap [49] to produce a 3D pose estimate with a free parameter for the scale of the hand. As the true hand scale is not known, for each sequence we swept the hand scale with a discretization that corresponds to sub-millimeter adjustment in depth, and find the scale that maximizes the Contact IoU with respect to ground truth (Figure 7). Contact was estimated by finding intersection between the hand and sensor meshes.

5.2 Can inference succeed with new people?

 Table 2. PressureVisionNet outperforms the 3D Pose Baseline on a test set of unseen participants.

Method	Temporal Acc	Contact IoU	Vol. IoU	MAE
Zero Guesser	53.7%	0.0%	0.0%	51.9 Pa
3D Pose Baseline [49]	78.1%	13.0%	-	-
${\it PressureVisionNet}$	96.2%	55.8%	41.3%	39.9 Pa

Table 2 shows performance on the \approx 490k images in our test set, which includes frames from all four cameras, five distinct lighting conditions, and 6 participants performing 36 actions. None of the participants in our test set were in our training set, so performance indicates how well our approach generalizes to new people.

Table 3. Participants were prompted to perform actions with various force levels. Estimated forces from PressureVisionNet correlate with ground truth (GT).

ronce nequested	mean GI	mean Est.
to Participant	Force	Force
High Force	8.16 N	$5.73 \ \mathrm{N}$
Low Force	3.24 N	3.63 N
No Contact	0.00 N	0.04 N

Force Requested Mean GT Mean Est.

PressureVisionNet outperformed our baseline model. Its discretized representation for pressure limits its best possible Volumetric IoU to 81% and MAE

to 10.9 Pa. The Zero Guesser always outputs a zero-pressure image. It achieved 53.7% temporal accuracy since no pressure was recorded for the majority of frames in the dataset. The 3D Pose Estimator had generally low performance. PressureVisionNet performed well with non-contact images and inferred higher pressures on average when participants were instructed to apply more force (Table 3 and Figure 5).

Skin Tone Skin tone is a significant source of variability across people. Due to the limited number of participants in each skin tone category, we used cross-validation. The training set is split into five folds, and the model is trained with one rotating fold left out for testing. As shown Table 4, the highest performance was with dark skin tones. We performed a two-sample Kolmogorov–Smirnov test between each pair of skin tone categories, and did not find a statistically significant difference.

Table 4. Performance across skin tones is found via cross validation. We did not find a statistically significant difference between categories.

Skin Tone	Vol. IoU
Light	38.5%
Medium	37.2%
Dark	39.0%

5.3 Is hand-related appearance used for inference?

We expected the appearance of the hand and cast shadows to provide visual cues for pressure inference. Our results suggest that this is the case. In addition to providing evidence for the underlying information used for inference, these results reduce the likelihood that PressureVisionNet is exploiting information specific to our data that would not be generalizable, such as unintended sensor artifacts.

Sensitivity Analysis We conducted a sensitivity analysis to identify spatial regions that strongly influence PressureVisionNet's output. Our method is similar to the Occlusion Sensitivity method from Zieler and Fergus [66]. We divide the input RGB image into a 48x48 grid of cells. For each cell i, j, we create a new image $\mathbf{B}_{i,j}(I)$ replacing the cell contents with its average color. We provide this image as an input to the network and calculate how much the output pressure changes. We create a sensitivity image by normalizing the following image, S, to be between 0 and 1.

$$S_{i,j} = ||f(\mathbf{B}_{i,j}(I)) - f(I)||_2$$
(2)



PressureVision: Estimating Hand Pressure from a Single RGB Image 13

Fig. 8. Left: PressureVisionNet is sensitive to the appearance of the hand and shadows near regions of contact. Right: Results of projecting sensitivity images onto hand meshes and averaging over many frames for selected actions.

Figure 8 shows examples of resulting sensitivity images that indicate the model uses the appearance of the hand and cast shadows near regions of contact. To show this objectively, we also projected the sensitivity images from multiple frames onto the hand meshes generated for the 3D Pose Baseline. Figure 8 shows the results for three different actions that confirm that PressureVisionNet is highly sensitive to blurring of the hand near regions of contact.

We also found a common failure mode that corroborates this result. As shown in the bottom row of Figure 5, PressureVisionNet typically guesses zero pressure for parts of the hand that are occluded from view.

Dependence on Actions We found that PressureVisionNet's performance depended on the participant's action. For example, performance was highest with a single index finger action and lower with actions that apply pressure with the palm (Table 5). Two relevant factors are likely the visibility of the part of the hand near the region of contact and the tendency for larger areas of ground truth pressure to result in larger errors.

Table 5. The performance of PressureVisionNet on selected actions. Actions with lower visibility and larger contact areas tend to result in lower performance.

Action name	Vol. IoU
Press index, pull towards	59.4%
Press all fingers	42.1%
Press all fingers and palm	35.4%
Press palm	33.2%

14 P. Grady et al.



Fig. 9. PressureVisionNet may generalize to new environments. Results shown in b) were captured with a smartphone camera in unaltered office lighting on a tabletop. Pressure estimates can be accumulated over time to allow writing with a finger.

5.4 Is inference reasonable with new conditions?

We considered whether PressureVisionNet can produce reasonable estimates when given new RGB images acquired with conditions that differ from Pressure-VisionDB. For this question, we provide preliminary evidence based on images of hands not included in the dataset. We captured additional data with a smartphone camera in unaltered office lighting. The smartphone camera had different focal length, resolution, and color characteristics than the dataset cameras. The smartphone was mounted with a tripod to observe a normal white tabletop.

Figure 9 demonstrates that PressureVisionNet has some ability to generalize to new environments. The system can generalize to images with multiple hands, and can run over multiple frames in a video while accumulating inferred pressure estimates. In the bottom image, an author used the system to write with their index finger first with low pressure, then with high pressure. This illustrates a potential application and provides evidence that PressureVisionNet is not overly sensitive to the camera, the illumination, or the contact surface.

6 Conclusion

We collected a novel dataset, PressureVisionDB, and developed a deep model, PressureVisionNet, to infer hand pressure from a single RGB image. Using this model, we provided evidence that hand pressure can be accurately inferred from a single RGB image of a hand from a previously unobserved person. Our results suggest that the appearance of regions of the hand near regions of contact are especially informative when inferring pressure.

Acknowledgements We thank Kevin Harris, Steve Miller, and Steve Olsen for their help in data collection, and Robert Wang, Minh Vo, Tomas Hodan, Amy Zhao, Kenrick Kin, Mark Richardson, and Cem Keskin for their advice.

References

- Asadi, H., Zhou, G., Lee, J.J., Aggarwal, V., Yu, D.: A computer vision approach for classifying isometric grip force exertion levels. Ergonomics 63(8), 1010–1026 (2020)
- 2. Bhirangi, R., Hellebrekers, T., Majidi, C., Gupta, A.: Reskin: versatile, replaceable, lasting tactile skins. Conference on Robot Learning (CoRL) (2021)
- Brahmbhatt, S., Ham, C., Kemp, C.C., Hays, J.: ContactDB: Analyzing and predicting grasp contact via thermal imaging. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- Brahmbhatt, S., Handa, A., Hays, J., Fox, D.: ContactGrasp: Functional multifinger grasp synthesis from contact. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2386–2393. IEEE (2019)
- 5. Brahmbhatt, S., Tang, C., Twigg, C.D., Kemp, C.C., Hays, J.: ContactPose: A dataset of grasps with object contact and hand pose. In: IEEE European Conference on Computer Vision (ECCV) (August 2020)
- 6. Brahmbhatt, S.M.: Grasp contact between hand and object: Capture, analysis, and applications. Ph.D. thesis, Georgia Institute of Technology (2020)
- Büscher, G.H., Kõiva, R., Schürmann, C., Haschke, R., Ritter, H.J.: Flexible and stretchable fabric-based tactile sensor. Robotics and Autonomous Systems 63, 244– 252 (2015)
- Cavazos, J.G., Phillips, P.J., Castillo, C.D., O'Toole, A.J.: Accuracy comparison across face recognition algorithms: Where are we on measuring race bias? IEEE Transactions on Biometrics, Behavior, and Identity Science 3(1), 101–111 (2020)
- Chen, N., Westling, G., Edin, B.B., van der Smagt, P.: Estimating fingertip forces, torques, and local curvatures from fingernail images. Robotica 38(7), 1242–1262 (2020)
- Chu, F.J., Xu, R., Vela, P.A.: Real-world multiobject, multigrasp detection. IEEE Robotics and Automation Letters 3(4), 3355–3362 (2018)
- Clever, H.M., Erickson, Z., Kapusta, A., Turk, G., Liu, C.K., Kemp, C.C.: Bodies at Rest: 3d human pose and shape estimation from a pressure image using synthetic data. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2020)
- 12. Clever, H.M., Grady, P., Turk, G., Kemp, C.C.: BodyPressure inferring body pose and contact pressure from a depth image. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
- Cook, C.M., Howard, J.J., Sirotin, Y.B., Tipton, J.L., Vemury, A.R.: Demographic effects in facial recognition and their dependence on image acquisition: An evaluation of eleven commercial systems. IEEE Transactions on Biometrics, Behavior, and Identity Science 1(1), 32–41 (2019)
- Davis, A., Rubinstein, M., Wadhwa, N., Mysore, G.J., Durand, F., Freeman, W.T.: The visual microphone: passive recovery of sound from video. ACM Transactions on Graphics 33(4), 79:1–79:10 (2014)
- 15. Dehaene, S.: The neural basis of the weber-fechner law: a logarithmic mental number line. Trends in Cognitive Sciences **7**(4), 145–147 (2003)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 248–255. IEEE (2009)
- Ehsani, K., Tulsiani, S., Gupta, S., Farhadi, A., Gupta, A.: Use the force, Luke! Learning to predict physical forces by simulating effects. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 224–233 (2020)

- 16 P. Grady et al.
- Fitzpatrick, P.M., Torres-Jara, E.R.: The power of the dark side: using cast shadows for visually-guided touching. In: 4th IEEE/RAS International Conference on Humanoid Robots, 2004. vol. 1, pp. 437–449. IEEE (2004)
- Garcia-Hernando, G., Yuan, S., Baek, S., Kim, T.K.: First-person hand action benchmark with RGB-D videos and 3d hand pose annotations. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 409–419 (2018)
- Grady, P., Tang, C., Twigg, C.D., Vo, M., Brahmbhatt, S., Kemp, C.C.: Contact-Opt: Optimizing contact to improve grasps. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1471–1481 (2021)
- Hampali, S., Rad, M., Oberweger, M., Lepetit, V.: Honnotate: A method for 3d annotation of hand and object poses. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3196–3206 (2020)
- Hasson, Y., Tekin, B., Bogo, F., Laptev, I., Pollefeys, M., Schmid, C.: Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 571–580 (2020)
- Hasson, Y., Varol, G., Tzionas, D., Kalevatykh, I., Black, M.J., Laptev, I., Schmid, C.: Learning joint reconstruction of hands and manipulated objects. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11807– 11816 (2019)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
- Holden, D., Kanoun, O., Perepichka, M., Popa, T.: Learned motion matching. ACM Transactions on Graphics 39(4), 53–1 (2020)
- Hu, H.H., Gooch, A.A., Creem-Regehr, S.H., Thompson, W.B.: Visual cues for perceiving distances from objects to surfaces. Presence: Teleoperators & Virtual Environments 11(6), 652–664 (2002)
- Hu, H.H., Gooch, A.A., Thompson, W.B., Smits, B.E., Rieser, J.J., Shirley, P.: Visual cues for imminent object contact in realistic virtual environments. In: Proceedings Visualization 2000. pp. 179–185. IEEE (2000)
- Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7132–7141 (2018)
- Hubona, G.S., Wheeler, P.N., Shirah, G.W., Brandt, M.: The relative contributions of stereo, lighting, and background scenes in promoting 3d depth visualization. ACM Transactions on Computer Human Interaction 6(3), 214–242 (1999)
- 30. Hwang, W., Lim, S.: Inferring interaction force from visual information without using physical force sensors. Sensors **17**(11), 2455 (2017)
- Johnson, M.K., Adelson, E.H.: Retrographic sensing for the measurement of surface texture and shape. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1070–1077. IEEE (2009)
- 32. Karunratanakul, K., Yang, J., Zhang, Y., Black, M.J., Muandet, K., Tang, S.: Grasping field: Learning implicit representations for human grasps. In: 8th International Conference on 3D Vision. pp. 333–344. IEEE (2020)
- 33. Krishnapriya, K., Albiero, V., Vangara, K., King, M.C., Bowyer, K.W.: Issues related to face recognition accuracy varying based on race and skin tone. IEEE Transactions on Technology and Society 1(1), 8–20 (2020)
- Li, Z., Sedlár, J., Carpentier, J., Laptev, I., Mansard, N., Sivic, J.: Estimating 3d motion and forces of person-object interactions from monocular video. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8640–8649 (2019)

- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2117–2125 (2017)
- Mascaro, S.A., Asada, H.H.: Photoplethysmograph fingernail sensors for measuring finger forces without haptic obstruction. IEEE Transactions on Robotics and Automation 17(5), 698–708 (2001)
- Mascaro, S.A., Asada, H.H.: Measurement of finger posture and three-axis fingertip touch force using fingernail sensors. IEEE Transactions on Robotics and Automation 20(1), 26–35 (2004)
- Meier, M., Streli, P., Fender, A., Holz, C.: Tapid: Rapid touch interaction in virtual reality using wearable sensing. In: 2021 IEEE Virtual Reality and 3D User Interfaces (VR). pp. 519–528. IEEE (2021)
- 39. Morph: Sensel Morph haptic sensing tablet, www.sensel.com/pages/ the-sensel-morph, Last accessed on 2020-02-25
- Narasimhaswamy, S., Nguyen, T., Nguyen, M.H.: Detecting hands and recognizing physical contact in the wild. Advances in Neural Information Processing Systems 33 (2020)
- Niikura, T., Watanabe, Y., Ishikawa, M.: Anywhere surface touch: utilizing any surface as an input area. In: Proceedings of the 5th Augmented Human International Conference. pp. 1–8 (2014)
- 42. Pantone LLC: Pantone SkinTone Guide (2012)
- Pérez-González, A., Vergara, M., Sancho-Bru, J.L.: Stiffness map of the grasping contact areas of the human hand. Journal of Biomechanics 46(15), 2644–2650 (2013)
- 44. Pham, T.H., Kheddar, A., Qammaz, A., Argyros, A.A.: Towards force sensing from vision: Observing hand-object interactions to infer manipulation forces. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2810–2819 (2015)
- Pham, T.H., Kyriazis, N., Argyros, A.A., Kheddar, A.: Hand-object contact force estimation from markerless visual tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence 40(12), 2883–2896 (2017)
- 46. Pressure Profile Systems: PPS TactileGlove, www.pressureprofile.com/ body-pressure-mapping/tactile-glove
- 47. Rogez, G., Supancic, J.S., Ramanan, D.: Understanding everyday hands in action from RGB-D images. In: IEEE International Conference on Computer Vision (ICCV). pp. 3889–3897 (2015)
- Romero, J., Kjellström, H., Kragic, D.: Hands in action: Real-time 3D reconstruction of hands in interaction with objects. In: 2010 IEEE International Conference on Robotics and Automation. pp. 458–463. IEEE (2010)
- 49. Rong, Y., Shiratori, T., Joo, H.: Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In: IEEE International Conference on Computer Vision Workshops (2021)
- Rosales, C., Porta, J.M., Ros, L.: Global optimization of robotic grasps. Proceedings of Robotics: Science and Systems VII (2011)
- Saxena, A., Driemeyer, J., Ng, A.Y.: Robotic grasping of novel objects using vision. The International Journal of Robotics Research 27(2), 157–173 (2008)
- Scott, J., Ravichandran, B., Funk, C., Collins, R.T., Liu, Y.: From image to stability: learning dynamics from human pose. In: IEEE European Conference on Computer Vision (ECCV). pp. 536–554. Springer (2020)

- 18 P. Grady et al.
- Starke, S., Zhao, Y., Komura, T., Zaman, K.: Local motion phases for learning multi-contact character movements. ACM Transactions on Graphics 39(4), 54–1 (2020)
- Sundaram, S., Kellnhofer, P., Li, Y., Zhu, J.Y., Torralba, A., Matusik, W.: Learning the signatures of the human grasp using a scalable tactile glove. Nature 569(7758), 698–702 (2019)
- Taheri, O., Ghorbani, N., Black, M.J., Tzionas, D.: GRAB: A dataset of wholebody human grasping of objects. In: IEEE European Conference on Computer Vision (ECCV) (2020)
- 56. TekScan: TekScan Grip System, www.tekscan.com
- 57. Tzionas, D., Ballan, L., Srikantha, A., Aponte, P., Pollefeys, M., Gall, J.: Capturing hands in action using discriminative salient points and physics simulation. International Journal of Computer Vision 118(2), 172–193 (2016)
- Ward-Cherrier, B., Pestell, N., Cramphorn, L., Winstone, B., Giannaccini, M.E., Rossiter, J., Lepora, N.F.: The tactip family: Soft optical tactile sensors with 3dprinted biomimetic morphologies. Soft Robotics 5(2), 216–227 (2018)
- Wettels, N., Santos, V.J., Johansson, R.S., Loeb, G.E.: Biomimetic tactile sensor array. Advanced Robotics 22(8), 829–849 (2008)
- Wu, H., Rubinstein, M., Shih, E., Guttag, J.V., Durand, F., Freeman, W.T.: Eulerian video magnification for revealing subtle changes in the world. ACM Transactions on Graphics 31(4), 65:1–65:8 (2012)
- Xiao, R., Schwarz, J., Throm, N., Wilson, A.D., Benko, H.: MRTouch: adding touch input to head-mounted mixed reality. IEEE Transactions on Visualization and Computer Graphics 24(4), 1653–1660 (2018)
- 62. Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5987–5995 (2017)
- 63. Yakubovskiy, P.: Segmentation models pytorch (2020)
- 64. Ye, Y., Liu, C.K.: Synthesis of detailed hand manipulations using contact sampling. ACM Transactions on Graphics (TOG) **31**(4), 41 (2012)
- Yuan, W., Dong, S., Adelson, E.H.: Gelsight: High-resolution robot tactile sensors for estimating geometry and force. Sensors 17(12), 2762 (2017)
- Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: IEEE European Conference on Computer Vision (ECCV). vol. 8689, pp. 818–833. Springer (2014)
- Zhang, H., Ye, Y., Shiratori, T., Komura, T.: Manipnet: neural manipulation synthesis with a hand-object spatial representation. ACM Transactions on Graphics (TOG) 40(4), 1–14 (2021)
- Zhang, Y., Hassan, M., Neumann, H., Black, M.J., Tang, S.: Generating 3d people in scenes without people. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6193–6203 (2020)
- Zhu, Y., Jiang, C., Zhao, Y., Terzopoulos, D., Zhu, S.C.: Inferring forces and learning human utilities from videos. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3823–3833 (2016)
- Zonios, G., Bykowski, J., Kollias, N.: Skin melanin, hemoglobin, and light scattering properties can be quantitatively assessed in vivo using diffuse reflectance spectroscopy. Journal of Investigative Dermatology 117(6), 1452–1457 (2001)