Datasheet - EgoPressure

Yiming Zhao[1] Taein Kwon[1] Paul Streli[1] Marc Pollefeys[1,2] Christian Holz[1]
[1]ETH Zürich [2]Microsoft MR&AI Lab, Zürich

This document is based on *Datasheets for Datasets* by Gebru *et al.* [8].

MOTIVATION

For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

The ability to perceive touch, contact, and pressure during hand-object interactions is vital for numerous applications in Augmented Reality (AR), Virtual Reality (VR), and robotic manipulation. Specifically, estimating these physical properties from an egocentric perspective is essential for facilitating real-world tasks in physical environments.
Researchers have aimed to estimate users' forces to better understand human actions, forming a basis for modeling and controlling intelligent agents [5], [22]. Instead of focusing solely on contact, prior work has estimated the pressure forces applied during hand interactions, which is crucial for robotic grasping tasks [17] and can provide an additional control dimension for input [20]. To estimate pressure from monocular images, visual cues such as fingernail alterations [3], [18] or object deformations [3] during press events have been used. Changes in object trajectory and interaction forces [6], [14], [19] also offer insights but are ineffective with static objects like tables and walls. Accurate pressure labels for training usually require instrumenting the user's hands with gloves [2], [15], [25] or the surface with force sensors [10], [19], [23], ideally flexible or conforming to various shapes [1], [13], [16]. However, this can alter the visual appearance and tactile features of the hands and surface, affecting interaction and limiting generalization to bare hands and uninstrumented surfaces. Grady et al. [9], [10] collected two datasets with ground-truth pressure maps using a Sensel Morph [12] pressure sensor to train a neural network for estimating contact regions on surfaces from single RGB images. However, their method relies on an external static camera and good visibility of the corresponding fingertips.
With EgoPressure, we aim to bridge this gap by offering a dataset that includes egocentric views, utilizing head-mounted cameras to better understand human interactions from this perspective. EgoPressure is a novel dataset of touch contact and pressure interaction from an egocentric perspective, complemented with hand pose meshes and fine-grained pressure intensities for each contact.

The results of this project are intended to improve the interaction between human hands and digital interfaces, particularly in AR/VR environments. As such, we foresee these results bringing significant benefits in terms of enhancing user experience and facilitating more natural and intuitive interactions in digital spaces.

Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

This dataset was created by SIPLAB (Sensing, Interaction & Perception Lab) and CVG (Computer Vision and Geometry Group) at the Department of Computer Science of ETH Zürich.

What support was needed to make this dataset? (e.g., who funded the creation of the dataset? If there is an associated grant, provide the name of the grantor and the grant name and number, or if it was supported by a company or government agency, give those details.)

Resources from the corresponding research groups at ETH Zürich fund this project.

Any other comments?

No.

COMPOSITION

What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)? Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.

How many instances are there in total (of each type, if appropriate)?

There are 21 participants in total, with recordings from 9 devices (7 static cameras, 1 head-mounted egocentric camera, and 1 Sensel Morph pressure-sensing touchpad, hereafter referred to as 'touchpad').

Does the dataset contain all possible instances or

is it a sample (not necessarily random) of instances from a larger set? If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances because instances were withheld or unavailable).

The full dataset will be publicly released upon acceptance of our associated submission to the NeurIPS 2024 Datasets and Benchmarks Track. Reviewers will have partial access during the review process.

What data does each instance consist of? "Raw" data (e.g., unprocessed text or images) or features? In either case, please provide a description.

The data for a single participant consists of 64 sequences (31 interaction gestures and one calibration gesture for each hand) with synchronized, undistorted images from all cameras and pressure data in a binary file (.bin) captured by the touchpad. Additionally, it includes the tracked camera poses of the head-mounted egocentric camera as a .json file for each frame in each sequence. The camera parameters are provided in a .json file for each sequence.

Is there a label or target associated with each instance? If so, please provide a description.

We provide high-fidelity hand pose and mesh data during interactions. Our automatic annotation method focuses on frames of interest, specifically those immediately before, during, and after interaction with the touchpad, capturing hand pose and mesh data. This selective annotation targets the crucial moments when hand poses are most likely to correlate with pressure changes. We use the MANO [11], [21] model for annotation, capturing details about finger and palm positioning, orientation, and shape.

Is any information missing from individual instances? If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

No information is missing from individual participant captures.

Are relationships between individual instances made explicit (e.g., users' movie ratings, social network links)? If so, please describe how these relationships are made explicit.

We label data captured by a given participant, but there are no special relationships between individuals' data.

Are there recommended data splits (e.g., training, development/validation, testing)? If so, please provide a description of these splits, explaining the rationale behind them.

When splitting the data, we recommend using a leave-one-subject-out (LOSO) cross-validation split. Data from each participant should not be split across different sets to ensure the model remains unbiased. This is critical in studies where individual variability could influence the model's learning and performance. Keeping each participant's data confined to one split (training, validation, or testing) ensures that the model's evaluation is a true test of its ability to generalize to new, unseen participants.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description.

Overall, we did not notice any significant sources of noise or errors in our data collection process. We observed some glitching in a few camera frames, but the number of affected frames is very low; in the entire recording from one participant, only 3 or 4 frames were impacted. This minimal disruption is unlikely to significantly affect the overall quality or outcomes of our analysis.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

We provide annotated hand pose and meshes using MANO, which is widely used [26]–[28] and available under a single-user, non-exclusive, non-transferable, free of charge right (`https://mano.is.tue.mpg.de/license.html`).

Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals' non-public communications)? If so, please provide a description.

No.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.

No.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

Yes.

Does the dataset identify any subpopulations (e.g., by age, gender)? If so, please describe how these subpopulations are identified and provide a description of their respective distributions within the dataset.

All demographic information (such as sex, age, weight, and height) along with the sensor and video data are pseudonymized by assigning a numeric code to each participant. Personal data (sex, age, weight, and height) is stored separately from the sensor and video data and is accessible only to the primary researchers involved in the study. Only global summary statistics about this information are published.

Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset? If so, please describe how.

No, participants' faces and voices were not recorded.

Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)? If so, please provide a description.

Yes, the videos of the participants' hands could potentially be used to infer racial or ethnic origins.

Any other comments?

The EgoPressure project has received approval from the ETH Zürich Ethics Commission as proposal EK 2023-N-228. This approval includes both the data collection and the public release of the dataset. All participants provided explicit written consent for recording their sessions, creating the dataset, and releasing it.

COLLECTION

How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

All data was directly observable and measured by sensors.

Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created. Finally, list when the dataset was first published.

Each recording session for a single participant lasted around one hour, encompassing instructions, tutorials, breaks, and the actual data capture of the touchpad interactions.

What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)? How were these mechanisms or procedures validated?

To capture accurate ground-truth labels for hand pose and pressure from egocentric views, we constructed a data capture rig that integrates eight Microsoft Azure Kinect cameras and a pressure-sensitive touchpad (Sensel Morph). Seven cameras were stationary, positioned on a capture rig with distances ranging from 0.337 m to 0.631 m from the touchpad, recording color frames at 2560×1440 resolution and depth frames in WFOV 2×2 binned mode (512×512 resolution) at 30 fps. The eighth camera, an egocentric camera, was helmet-mounted, capturing video at 1920×1080 resolution. The touchpad, measuring 240×169.5 mm, was mounted on a tripod head, and four different texture overlays (white, green, dark wood, light wood) were used across participants. All cameras and the touchpad were connected to two workstations (Intel Core i7-9700K, Nvidia GeForce RTX 3070), with timestamps synchronized via a Raspberry Pi CM4 using PTP, which also triggered all cameras simultaneously at 30 fps.

Tracking of the head-mounted camera's pose was facilitated by nine active infrared markers around the touchpad, extracted from the infrared images using thresholding, and computed using Perspective-N-Points (PnP) algorithms.

Validation of the setup was conducted through specially developed software that streamed RGB frames and tracked poses from the head-mounted camera. Participants were instructed to mimic gestures shown in a video demo for consistent data collection. To improve data segmentation during post-processing, participants wore black sport sleeves to enhance the contrast between the hand and arm. Each session focused on the hand interacting with the touchpad, with instructions to keep the other hand out of view, ensuring that the data collected was both accurate and relevant to the study's objectives.

What was the resource cost of collecting the data? (e.g. what were the required computational resources, and the associated financial costs, and energy consumption - estimate the carbon footprint. See Strubell *et al.* [24] for approaches in this area.)

The recording equipment is reusable and was acquired for use in previous studies. The annotation process requires substantial computational resources, utilizing high-performance graphics cards shared by the Euler cluster at ETH Zürich. The entire annotation process for each participant takes around 64 hours, resulting in an energy consumption of approximately 19.2 kWh.

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?

—

Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?

The authors of this dataset were the experimenters who accompanied the participants throughout the study. Participants were recruited on a voluntary basis and received a small gift as compensation for their participation.

Were any ethical review processes conducted (e.g., by an institutional review board)? If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

Yes, the study was approved by the ethics board of ETH Zürich (EK 2023-N-228). This included a review concerning the physical safety of the participants, data security, and general ethical considerations.

Does the dataset relate to people? If not, you may skip the remainder of the questions in this section.

Yes.

Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?

We recorded the data in question directly.

Were the individuals in question notified about the data collection? If so, please describe (or show with screenshots or other information) how notice was provided, and provide a link or other access point to, or otherwise reproduce, the exact language of the notification itself.

Yes, participants explicitly chose to join the data recording and signed an information and consent form explaining the study setup, all recorded data, and how the data would be used and published.

Did the individuals in question consent to the collection and use of their data? If so, please describe (or show with screenshots or other information) how consent was requested and provided, and provide a link or other access point to, or otherwise reproduce, the exact language to which the individuals consented.

Yes, all participants signed a consent form that informed them about the data collection. They agreed to the recording and publication of the data.

If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses? If so, please provide a description, as well as a link or other access point to the mechanism (if appropriate)

Generally, participants can revoke their consent at any time by submitting an explicit written request via email to the authors.

Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis)been conducted? If so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.

No formal data protection impact analysis was conducted. However, the dataset is released in an anonymized form and includes only footage of hand-surface interactions captured in a controlled indoor setting, featuring neutral activities and gestures.

Any other comments?

No.

PREPROCESSING / CLEANING / LABELING

Was any preprocessing/cleaning/labeling of the data done (e.g.,discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)? If so, please provide a description. If not, you may skip the remainder of the questions in this section.

Yes, the data was synchronized across all devices. We developed an automatic method to annotate the hand pose based on the MANO [11], [21] model, implemented using DIB-R [4] and the NVIDIA Kaolin library [7].

Was the "raw" data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the "raw" data.

Yes, all raw data was saved.

Is the software used to preprocess/clean/label the instances available? If so, please provide a link or other access point.

The full annotation software may be shared upon request, given the acceptance of the associated submission.

Any other comments?

No.

USES

Has the dataset been used for any tasks already? If so, please provide a description.

Currently, there is no published work related to the dataset. We describe and demonstrate the usage of the dataset for egocentric hand pressure estimation in our associated submission.

Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.

Not applicable. We plan to list papers that make use of the dataset on the EgoPressure website.

What (other) tasks could the dataset be used for?

The dataset supports research on hand pose and mesh estimation from egocentric, static, single, and multi-view RGB-D cameras. Specifically, it captures hands interacting with rigid surfaces, improving the robustness of hand pose and pressure estimation in contact scenarios, which is vital for augmented reality, robotics, and human-computer interaction. Additionally, researchers can estimate and study hand poses using the associated pressure maps from the touchpad.

Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses? For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

Future users should be aware of potential biases in the dataset (e.g., based on age, given that the participants' age range spans 23-32 years). It is important to avoid uses that could result in unfair treatment or undesirable harms, such as stereotyping or quality of service issues. To mitigate these risks, users should validate the dataset's applicability across different demographic groups and contexts.

Are there tasks for which the dataset should not be used? If so, please provide a description.

The dataset must not be used for any analysis related to the race of any participant.

Any other comments?

No.

DISTRIBUTION

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.

Yes, the dataset will be publicly available for non-commercial, academic use according to the CC BY-NC-SA license.

How will the dataset be distributed (e.g., tarball on website, API, GitHub)? Does the dataset have a digital object identifier (DOI)?

Instructions and a link for downloading the tarball files from ETH servers will be posted on a website hosted by ETH Zürich for long-term support.

When will the dataset be distributed?

We will release the complete dataset to the wider public upon acceptance of our associated publication at a peer-reviewed venue.

Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

The dataset is released under the Creative Commons CC BY-NC-SA license.

Have any third parties imposed IP-based or other

restrictions on the data associated with the instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.

No.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation.

No.

Any other comments?

No.

MAINTENANCE

Who is supporting/hosting/maintaining the dataset?

The dataset is hosted on ETH Zürich servers with long-running maintenance intended for long-term availability.

How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

For any question, please contact the authors of dataset,

- Yiming Zhao `yimzhao@ethz.ch`
- Taein Kwon `taein.kwon@inf.ethz.ch`
- Paul Streli `paul.streli@inf.ethz.ch`

and corresponding author: Prof. Dr. Christian Holz `christian.holz@inf.ethz.ch`.

Is there an erratum? If so, please provide a link or other access point.

No.

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

If the authors discover or are made aware of issues with the dataset, it will be updated, and users will be informed through the project website, which will include a changelog.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.

No.

Will older versions of the dataset continue to be supported/hosted/maintained? If so, please describe how. If not, please describe how its obsolescence will be communicated to users.

The project website will host older versions of the dataset and code or provide the necessary information to roll back any changes detailed in the changelog.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

The dataset was captured using a custom data capture setup to ensure high annotation fidelity. Currently, there are no processes or plans to support extensions from outside contributors. Researchers interested in collaborating may add to the dataset in the future, but any additions must be made in a way that does not compromise compatibility with the original dataset.

Any other comments?

No.

REFERENCES

[1] Raunaq Bhirangi, Tess Hellebrekers, Carmel Majidi, and Abhinav Gupta. Reskin: versatile, replaceable, lasting tactile skins. In *5th Annual Conference on Robot Learning*, 2021.

[2] Gereon H Büscher, Risto Kõiva, Carsten Schürmann, Robert Haschke, and Helge J Ritter. Flexible and stretchable fabric-based tactile sensor. *Robotics and Autonomous Systems*, 63:244–252, 2015.

[3] Nutan Chen, Göran Westling, Benoni B Edin, and Patrick van der Smagt. Estimating fingertip forces, torques, and local curvatures from fingernail images. *Robotica*, 38(7):1242–1262, 2020.

[4] Wenzheng Chen, Jun Gao, Huan Ling, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *Advances In Neural Information Processing Systems*, 2019.

[5] Sammy Christen, Muhammed Kocabas, Emre Aksan, Jemin Hwangbo, Jie Song, and Otmar Hilliges. D-grasp: Physically plausible dynamic grasp synthesis for hand-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20577–20586, 2022.

[6] Kiana Ehsani, Shubham Tulsiani, Saurabh Gupta, Ali Farhadi, and Abhinav Gupta. Use the force, luke! learning to predict physical forces by simulating effects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 224–233, 2020.

[7] Clement Fuji Tsang, Maria Shugrina, Jean Francois Lafleche, Towaki Takikawa, Jiehan Wang, Charles Loop, Wenzheng Chen, Krishna Murthy Jatavallabhula, Edward Smith, Artem Rozantsev, Or Perel, Tianchang Shen, Jun Gao, Sanja Fidler, Gavriel State, Jason Gorski, Tommy Xiang, Jianing Li, Michael Li, and Rev Lebaredian.

Kaolin: A pytorch library for accelerating 3d deep learning research. `https://github.com/NVIDIAGameWorks/kaolin`, 2022.

[8] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumeé III, and Kate Crawford. Datasheets for Datasets. *arXiv:1803.09010 [cs]*, January 2020.

[9] Patrick Grady, Jeremy A Collins, Chengcheng Tang, Christopher D Twigg, Kunal Aneja, James Hays, and Charles C Kemp. Pressurevision++: Estimating fingertip pressure from diverse rgb images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 8698–8708, 2024.

[10] Patrick Grady, Chengcheng Tang, Samarth Brahmbhatt, Christopher D Twigg, Chengde Wan, James Hays, and Charles C Kemp. Pressurevision: Estimating hand pressure from a single rgb image. In *European Conference on Computer Vision*, pages 328–345. Springer, 2022.

[11] Yana Hasson, Gül Varol, Dimitris Tzionas, Igor Kalevatykh, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *CVPR*, 2019.

[12] Sensel Inc. Sensel morph., 2024.

[13] Hong-Ki Kim, Seunggun Lee, and Kwang-Seok Yun. Capacitive tactile sensor array for touch screen application. *Sensors and Actuators A: Physical*, 165(1):2–7, 2011.

[14] Zongmian Li, Jiri Sedlar, Justin Carpentier, Ivan Laptev, Nicolas Mansard, and Josef Sivic. Estimating 3d motion and forces of person-object interactions from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8640–8649, 2019.

[15] PPS UK Limited. Tactileglove - hand pressure and force measurement., 2023.

[16] Yiyue Luo, Yunzhu Li, Pratyusha Sharma, Wan Shou, Kui Wu, Michael Foshey, Beichen Li, Tomás Palacios, Antonio Torralba, and Wojciech Matusik. Learning human–environment interactions using conformal tactile textiles. *Nature Electronics*, 4(3):193–201, 2021.

[17] Priyanka Mandikal and Kristen Grauman. Learning dexterous grasping with object-centric visual affordances. In *2021 IEEE international conference on robotics and automation (ICRA)*, pages 6169–6176. IEEE, 2021.

[18] Stephen A Mascaro and H Harry Asada. Measurement of finger posture and three-axis fingertip touch force using fingernail sensors. *IEEE Transactions on Robotics and Automation*, 20(1):26–35, 2004.

[19] Tu-Hoa Pham, Nikolaos Kyriazis, Antonis A Argyros, and Abderrahmane Kheddar. Hand-object contact force estimation from markerless visual tracking. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2883–2896, 2017.

[20] Philip Quinn, Wenxin Feng, and Shumin Zhai. Deep touch: Sensing press gestures from touch image sequences. *Artificial Intelligence for Human Computer Interaction: A Modern Approach*, pages 169–192, 2021.

[21] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, November 2017.

[22] Pierre Schumacher, Daniel Häufle, Dieter Büchler, Syn Schmitt, and Georg Martius. Dep-rl: Embodied exploration for reinforcement learning in overactuated and musculoskeletal systems. *arXiv preprint arXiv:2206.00484*, 2022.

[23] Pressure Mapping Sensors. Tekscan., 2024.

[24] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and Policy Considerations for Deep Learning in NLP. *arXiv:1906.02243 [cs]*, June 2019.

[25] Subramanian Sundaram, Petr Kellnhofer, Yunzhu Li, Jun-Yan Zhu, Antonio Torralba, and Wojciech Matusik. Learning the signatures of the human grasp using a scalable tactile glove. *Nature*, 569(7758):698–702, 2019.

[26] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 581–600. Springer, 2020.

[27] Lixin Yang, Kailin Li, Xinyu Zhan, Fei Wu, Anran Xu, Liu Liu, and Cewu Lu. Oakink: A large-scale knowledge repository for understanding hand-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20953–20962, 2022.

[28] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. Freihand: A dataset for markerless capture of hand pose and shape from single rgb images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 813–822, 2019.