ContactDB: Analyzing and Predicting Grasp Contact via Thermal Imaging Supplementary Material

Samarth Brahmbhatt[1], Cusuh Ham[1], Charles C. Kemp[1] and James Hays[1,2]
[1]Institute for Robotics and Intelligent Machines, Georgia Tech [2]Argo AI
{samarth.robo, cusuh, charlie.kemp, hays}@gatech.edu

Abstract

This document provides supplementary material for our submission. We compare ContactDB heatmaps qualitatively against the crowdsourced tactile saliency maps from [1]. We discuss the extent of heat dissipation while scanning the object, and potential sources of error in observing contact through the thermal camera and the texture mapping process. Lastly, we list the 50 objects used in ContactDB and the instructions given to participants for grasping the subset of 27 objects with the 'use' post-grasp intent. ContactDB can be explored interactively at https://contactdb.cc.gatech.edu.

1. Comparison to Tactile Mesh Saliency [1]

Qualitatively, the closest work to ContactDB that we've found is [1], which collects contact saliency information through crowd-sourcing by pairwise comparison of surface points. Figure 1(b) compares common objects from both datasets. Notably, data from [1] lacks clear finger-marks and resembles averaged contact maps. That data may be less accurate because it relies on self-reporting. For example, our data shows that people rarely contact the bottom half of the wine glass stem, whereas [1] shows high saliency for the entire stem.

2. Heat Dissipation During Data Collection

Scanning takes 18 s for a 360° rotation. Owing to the consistent use of hand-warmers and PLA material for 3D printed objects, thermal prints take more than 35 s to diffuse significantly (See Fig. 1(a)). Heat conduction across the surface of the plate does not seem to be a significant source of variation between 0 s and 18 s, since the prints are comparable in size and lack strongly blurred edges. This shows that the dissipation of finger heat on the object surface produces minimal artifacts in the contact maps presented in the paper. We operate the turntable motor at the maximum possible speed that avoids high centrifugal force and wear-and-tear.

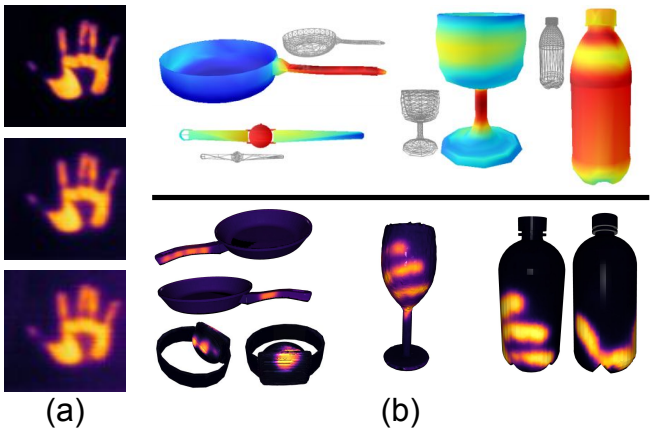

Figure 1: (a) Heat dissipation in the thermal images. Top-bottom: 0s, 18s, 35s. (b) Contact information collected by online crowd-sourcing ([1], top row) and ContactDB (ours, bottom row).

3. Accuracy of Texture Mapping

As discussed in Section 3.3 of the paper, thermal images from 9 views and corresponding object pose estimates are used in a texture mapping algorithm to produce a final mesh textured with a contact map. The whole process has multiple potential sources of error: calibration of the intrinsics and extrinsics of the Kinect v2 and thermal camera, inaccuracy in 3D printing the object, errors in object pose estimates due to noise/distortion in the Kinect depth maps, artifacts introduced by the texture mapping algorithm, etc. As such, the accuracy of this process can be different for different objects and sessions. In Figure 2, we attempt to quantify this error for one instance where we precisely heated a spot on the front button of the PS controller using a heated pencil-top eraser. In this case, we observed a final geometric error of 4.4 mm.

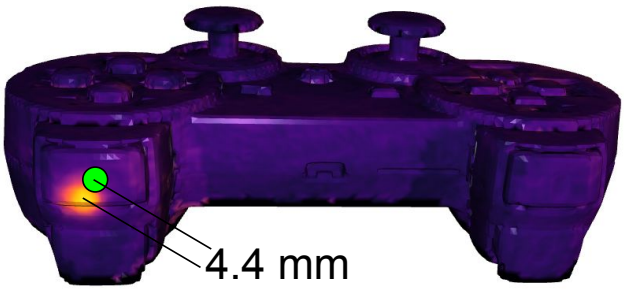

Figure 2: Geometric error of the texture mapping process. The spot on the front button shown in green was precision-heated with a warm pencil-top eraser.

4. List of Objects

Table 1 shows a list of all 50 objects in ContactDB, along with information about the which of these objects are included in the two functional grasping categories, and the specific 'use' instructions.

Object	handoff	use	use instruction
airplane	✓		
alarm clock	✓		
apple	✓	✓	eat
banana	✓	✓	peel
binoculars	✓	✓	see through
bowl	✓	✓	drink from
camera	✓	✓	take picture
cell phone	✓	✓	talk on
cube (small)	✓		
cube (medium)	✓		
cube (large)	✓		
cup	✓	✓	drink from
cylinder (small)	✓		
cylinder (medium)	✓		
cylinder (large)	✓		
door knob		✓	twist to open door
elephant	✓		
eyeglasses	✓	✓	wear
flashlight	✓	✓	turn on
flute	✓	✓	play
hammer	✓	✓	hit a nail
hand		✓	shake
headphones	✓	✓	wear
knife	✓	✓	cut
light bulb	✓	✓	screw in a socket
mouse	✓	✓	use to point and click
mug	✓	✓	drink from
pan	✓	✓	cook in
piggy bank	✓		
PS controller	✓	✓	play a game with
pyramid (small)	✓		
pyramid (medium)	✓		
pyramid (large)	✓		
rubber duck	✓		
scissors	✓	✓	cut with
sphere (small)	✓		
sphere (medium)	✓		
sphere (large)	✓		
Stanford bunny	✓		
stapler	✓	✓	staple
toothbrush	✓	✓	brush teeth
toothpaste	✓	✓	squeeze out toothpaste
torus (small)	✓		
torus (medium)	✓		
torus (large)	✓		
train	✓		
Utah teapot	✓	✓	pour tea from
water bottle	✓	✓	open
wine glass	✓	✓	drink wine from
wristwatch	✓		
Total	**48**	**27**	

Table 1: List of objects in ContactDB and specific ‘use’ instructions

References

[1] Manfred Lau, Kapil Dev, Weiqi Shi, Julie Dorsey, and Holly Rushmeier. Tactile mesh saliency. *ACM Transactions on Graphics (TOG)*, 35(4):52, 2016. 1