Evaluating Grasping Performance in "Unstructured" Tasks

Aaron M. Dollar
Yale University

Performance Evaluation of Hands

- Oliver: "Need to get the human out of the loop in evaluation"
- One step further: Need to decouple as many subsystems as possible in evaluation
  - E.g. sensing/perception and control schemes (and human) don't affect results
  - Easier to do in analysis/simulation, harder to do in experimental evaluation

Performance Evaluation of Hands

- Some high-level (experimental) approaches:
  - "open loop" testing scenarios
    - Can do this with adaptive hands!
  - Use robot arm where possible
    - Make scenario more like the real thing rather than benchtop testing
  - Wide variation of object and hand/object interaction properties

Performance Evaluation of Hands

- Object properties
  - What range of size, shape, weight, etc. can be accommodated?

Performance Evaluation of Hands

- Hand-object interaction properties
  - Position/orientation of the object relative to the hand
  - Force applied to the object
    - During reaching/hand closing
    - After grasp

Overview

- Ways that we've examined performance
  - Simulated Error
  - Contact/Disturbance Force
  - Pullout Force
  - Standard Objects and Tasks
  - Comparison to Human Performance

Overview

- Ways that we've examined performance
  - Simulated Error
  - Contact/Disturbance Force
  - Pullout Force
  - Standard Objects and Tasks
  - Comparison to Human Performance

  - Many of these have been used before/frequently by others.
    - Apologies in advance for the lack of "thoroughness", especially in terms of citations

Evaluation Approach 1: Simulated Error

Simulated Error

- Basic idea:
  - Arrange object in a known position/orientation that would generally not be attempted under no error
    - E.g. place off-center

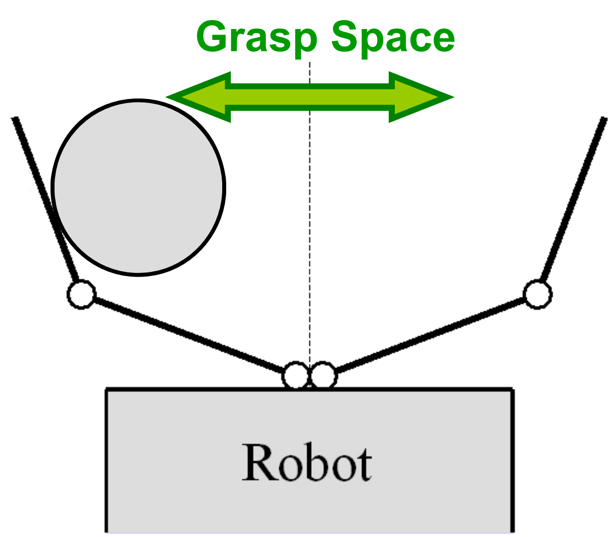

Simulated Error

- Basic idea:
  - Arrange object in a known position/orientation that would generally not be attempted under no error
    - E.g. place off-center
  - Examine range of positions/orientations over which the object can be successfully grasped
  - Do with a wide range of objects

Simulation Results

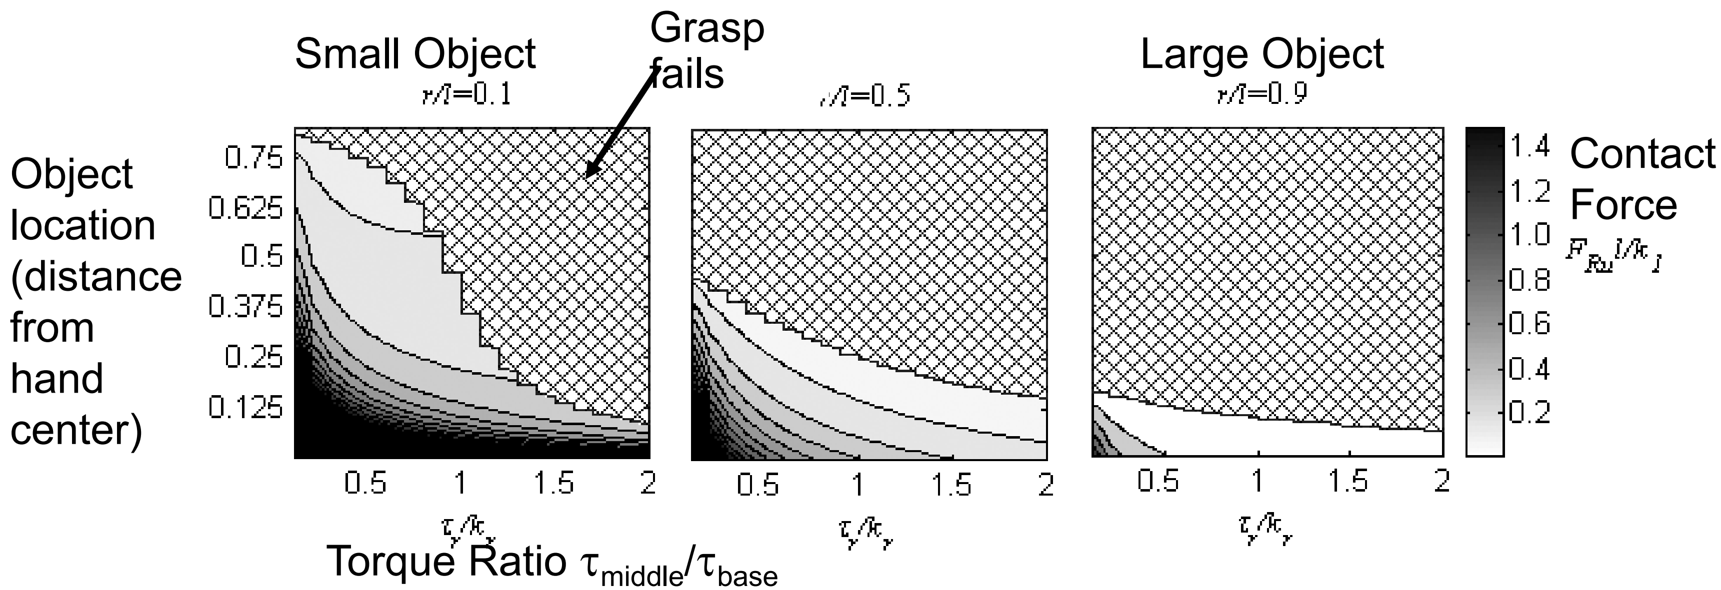

Dollar and Howe, IJRR 2010

Simulated Error - Experimental

- For experiments: decouple subsystems:
  – Get the human out of the loop
    - No teleop, placing objects in the hand, etc.
  – Do “open loop” testing
    - Place the object in a known location/orientation
    - Command the arm to go to a “wrong” position/orientation
    - Actuate/control hand open-loop
  – Examine range of success (objects, positions, etc.)

SDM Hand Performance Results

- Experimental testing of SDM hand:
  - Range of objects that could be grasped (O.C.)
  - How far from "center" could they be grasped (E.C.)
  - How much disturbance force applied during reach and grasp (O.C.+E.C.)

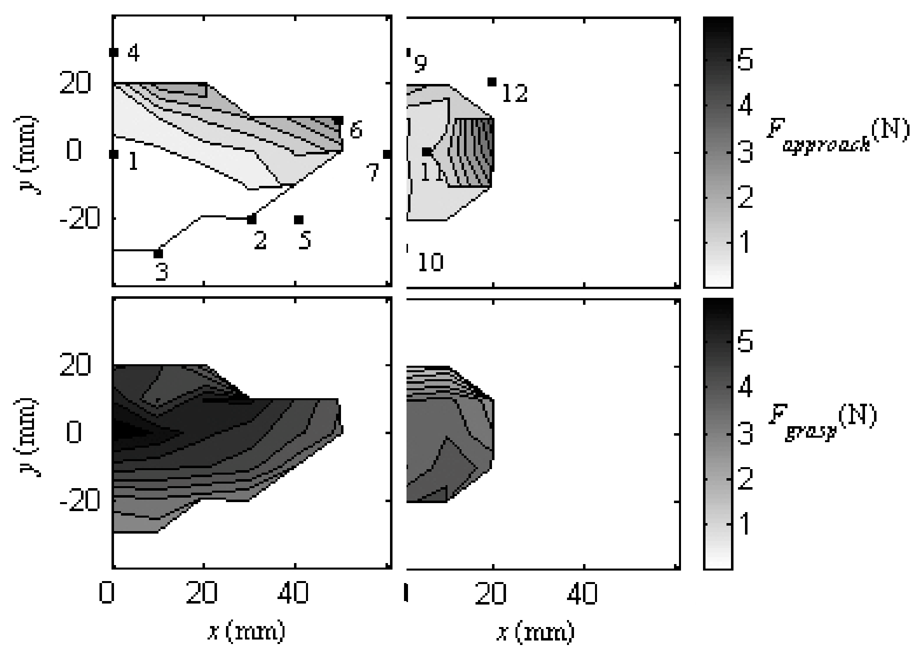

Analysis of Results

- Assume a normal distribution *z* of object position from expected position
  - Low σ for good sensing
  - High σ for poor sensing

$$z(x, x_t, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-x_t)^2}{2\sigma^2}}$$

$$p(x, x_t, \sigma) = \int_{-\infty}^{x} z(x')dx'$$

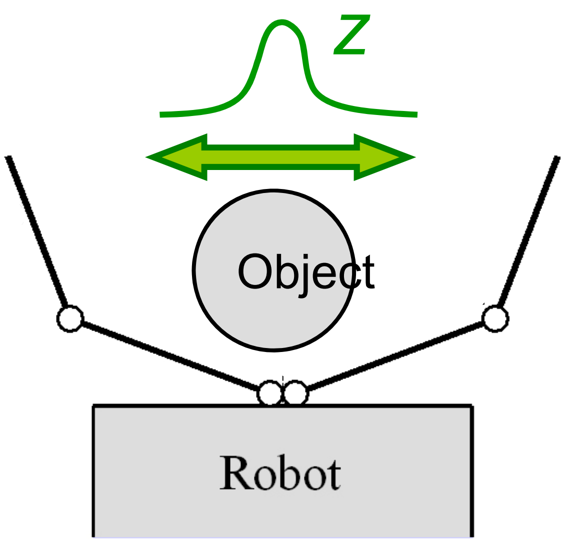

Evaluation Approach 2: Contact/Disturbance Force

Experiments

- Objects mounted to force/torque sensor
- Measure net contact force on object
  - Would tend to disturb objects

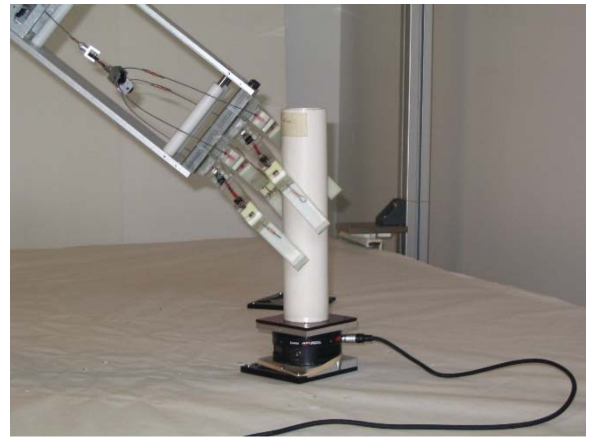

SDM Hand Performance Results

- Experimental testing of SDM hand:
  - Range of objects that could be grasped (O.C.)
  - How far from “center” could they be grasped (E.C.)
  - How much disturbance force applied during reach and grasp (O.C.+E.C.)

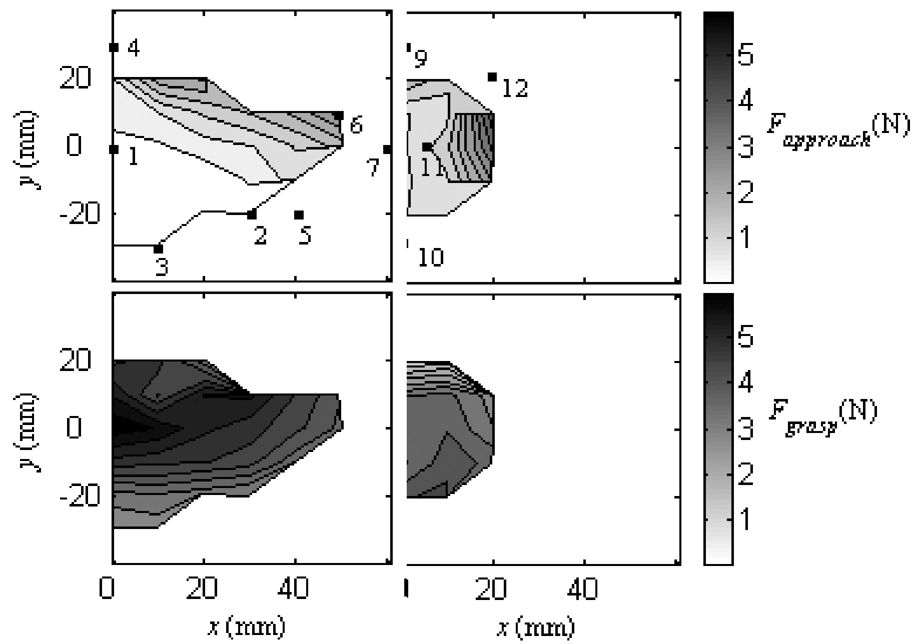

"Objects of Daily Living"

TABLE III
OBJECTS OF DAILY LIVING, ASSOCIATED ADLS, AND PHYSICAL PROPERTIES

Object	Categories	Source(s)	Mass (g)	Dims. (cm)	Veneer	Granite	Linoleum
Food Preparation							
bag of coffee beans, paper	D1, P1	[36]	n/a	n/a	n/a	n/a	n/a
baking pan (non-stick metal)	D1, P1, D2	[34]	351.9	21x11x8	0.105±0.006	0.139±0.013	0.069±0.007
bottle cap, metal	D1, P1, D2	[24, 31]	n/a	n/a	n/a	n/a	n/a
bowl, glass	D1, P1, D2	[28, 31]	545.1	18x8	0.223±0.009	0.124±0.006	0.163±0.003
box of crackers, cardboard	D1, P1	[37]	194.6	6x13x20	0.536±0.015	0.702±0.015	0.514±0.005
eating utensil, stainless steel	D1, P1, D2	*most sources	47.6	18x4x1	0.206±0.023	0.124±0.006	0.134±0.007
can of preserved food, steel	D1, P1		473.9	7x11	0.363±0.005	0.219±0.012	0.207±0.010
bowl, ceramic	D1, P1, D2	[28, 31]	479.3	13x8	0.236±0.006	0.111±0.009	0.266±0.011
juice carton (empty), paper	D1, P1, D2	[34]	74.5	10x10x24	0.257±0.011	0.303±0.040	0.252±0.013
coffee can (full), tin	D1, P1	[24]	397.4	10x18	0.329±0.016	0.163±0.008	0.219±0.016
dinner plate, ceramic	D1, P1, D2	[28]	798	27x3	0.350±0.011	0.222±0.004	0.349±0.011
drinking straw, plastic	D1, P1	[28]	n/a	n/a	n/a	n/a	n/a
beverage bottle, glass (empty)	D1, P1	[31, 32, 36]	213.7	6x24	0.325±0.030	0.171±0.020	0.168±0.018
beverage bottle, glass (full)	D1, P1	[31, 32, 36]	597.1	6x24	0.307±0.008	0.182±0.010	0.150±0.009
jar, glass	D1, P1, D2	[25, 34]	289	7x16	0.173±0.010	0.113±0.008	0.184±0.012
jar lid, steel	D1, P1, D2	[25, 30, 34]	n/a	n/a	n/a	n/a	n/a

Properties of ODLs

- For frictional properties, need to specify both object and surface properties
- Tested objects on common household surfaces
  - wood veneer, granite, polished metal, linoleum, unfinished wood, glass

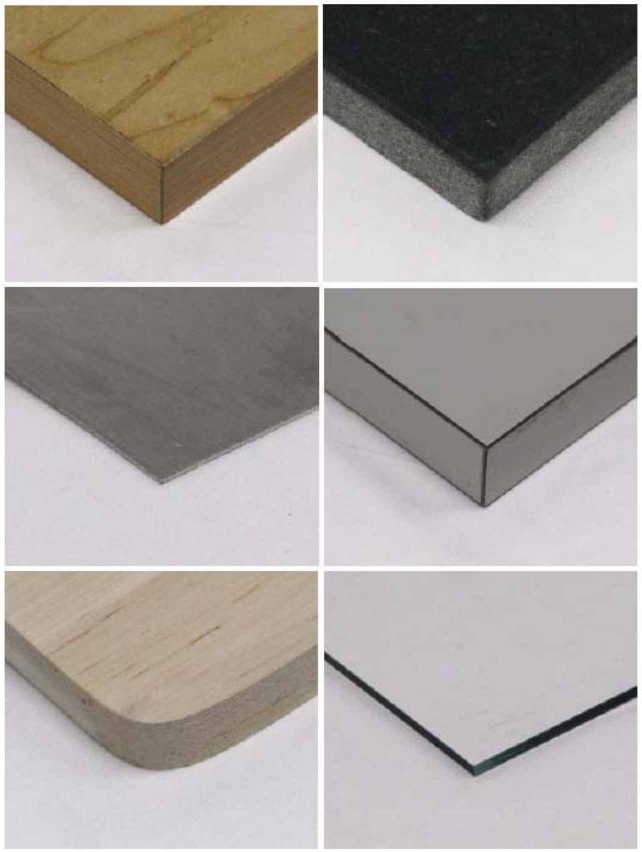

Frictional properties test setup

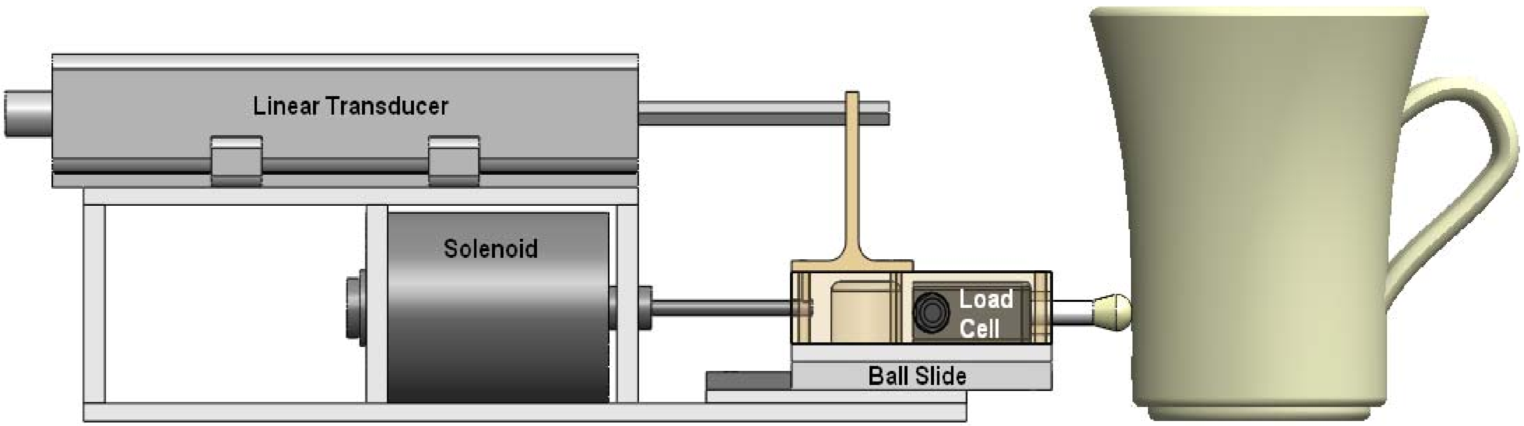

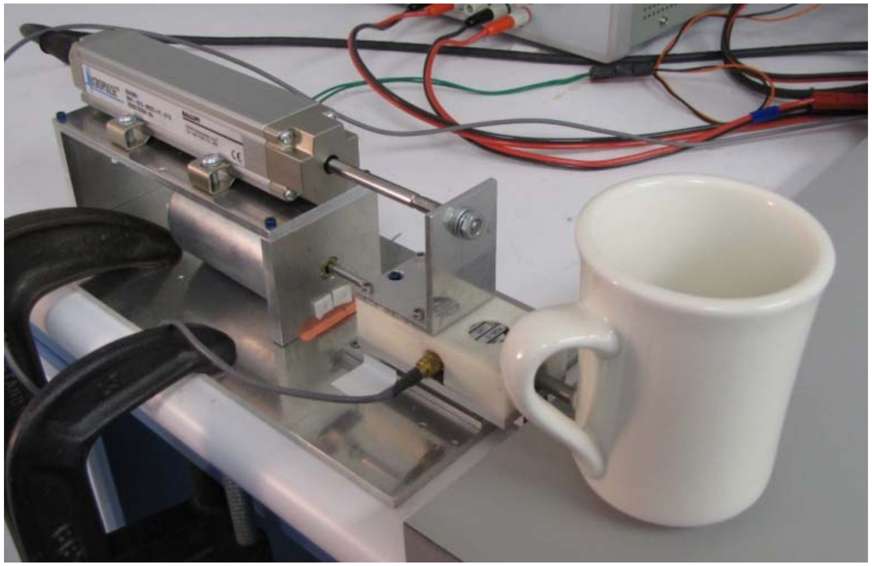

Frictional properties results

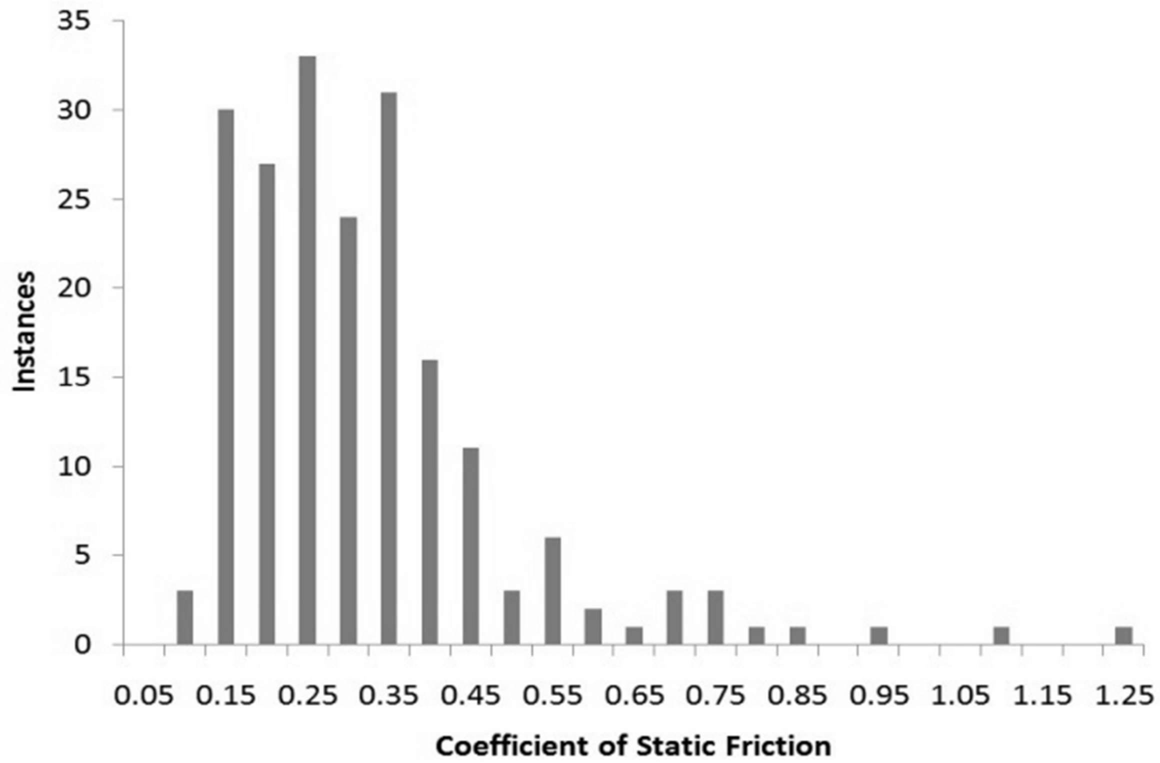

SDM Hand Performance Results

- Experimental testing of SDM hand:
  - Range of objects that could be grasped (O.C.)
  - How far from "center" could they be grasped (E.C.)
  - How much disturbance force applied during reach and grasp (O.C.+E.C.)

How much disturbance force is too much?

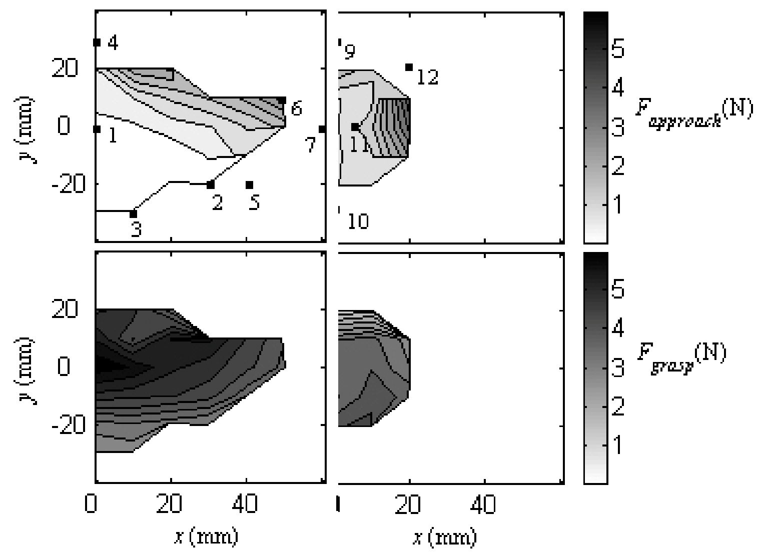

Frictional properties results

- 50% of objects → 0.463 N or less
- 75% of objects → 0.116 N or less
- 90% of objects → 0.051 N or less

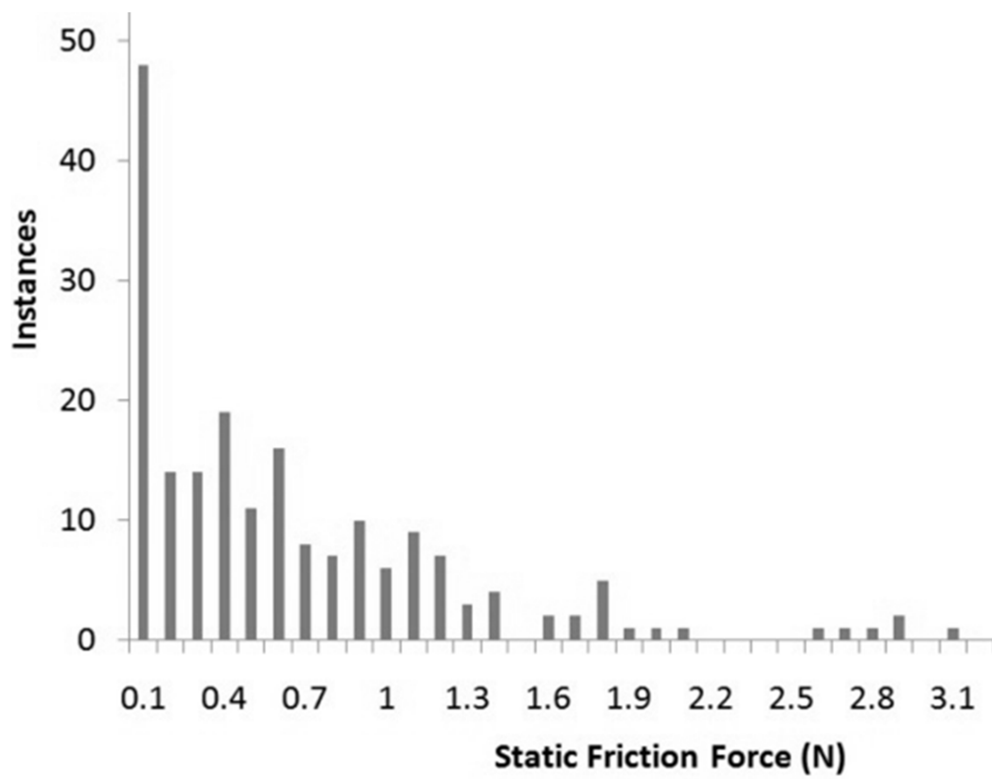

Frictional properties results

- 50% of objects → 0.463 N or less
- 75% of objects → 0.116 N or less
- 90% of objects → 0.051 N or less

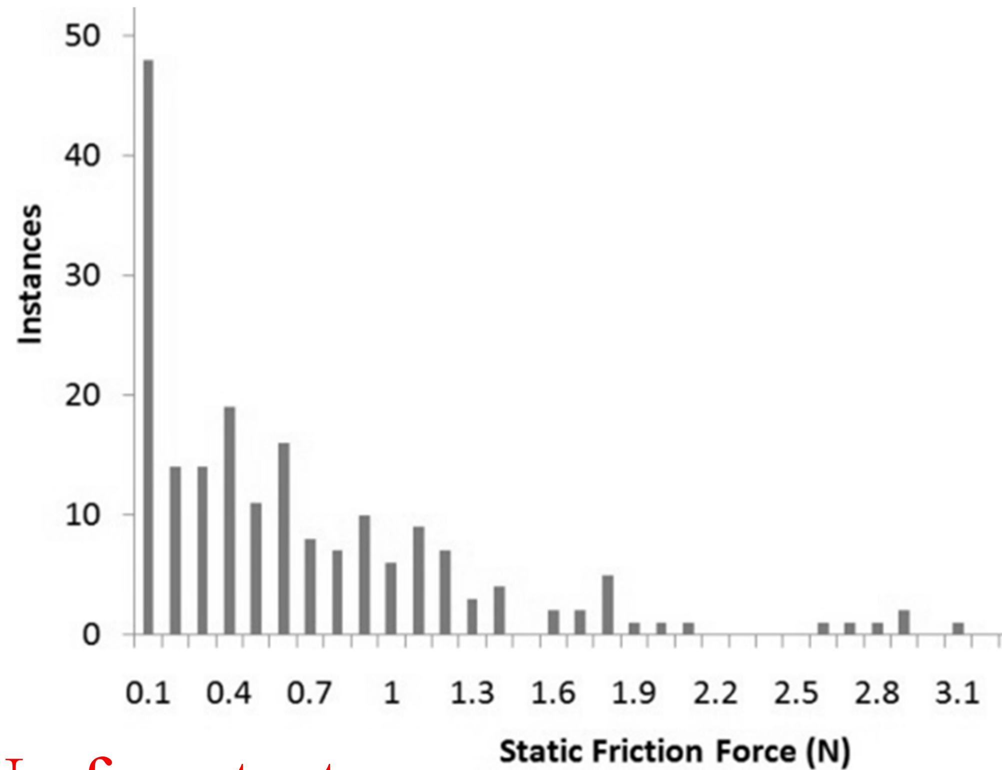

During reach/grasp, ~0.05N of contact force will move 90% of common objects, potentially causing grasp to fail.

Evaluation Approach 3: Pullout Force

(see Joe Falco's talk)

Evaluation Approach 4: Standard Objects and Tasks

YCB
Object and Model Set

Benchmarking in Manipulation Research


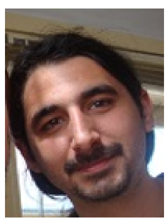

Prof. Aaron Dollar
Dr. Berk Calli
Yale University

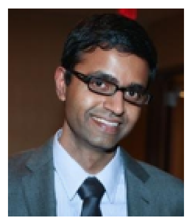
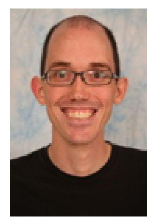

Prof. Sidd Srinivasa
Aaron Walsman
Carnegie Mellon Univ.


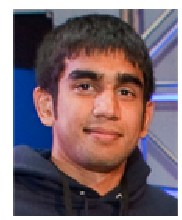

Prof. Pieter Abbeel
Arjun Singh
Univ. of California, **B**erkeley

YCB Object and Model Set

Benchmarking in Manipulation Research

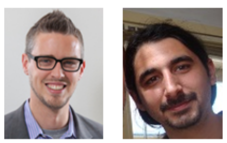

Prof. Aaron Dollar
Dr. Berk Calli
Yale University

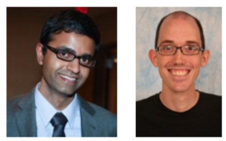

Prof. Sidd Srinivasa
Aaron Walsman
Carnegie Mellon Univ.

Prof. Pieter Abbeel
Arjun Singh
Univ. of California, Berkeley

Google

Kurt Konolige, James Bruce, Naresh Rajkumar

.

Project Overview

- Set of Real Physical Objects distributed to large number of research groups (~100 so far)
- Hi-Res RGBD scans and geometric models
- Task and Benchmark Protocols for standardized procedures

Background and Motivation

- How can we quantitatively compare approaches in manipulation research?
  - Simulation (good start, but finite value due to lack of realism)
  - Experiment
- Experimental validation is generally done in an *ad hoc* way
  - Objects and tasks are not standard
  - Incomplete detail (for replicability) is generally given
  - Procedures are not generalizable across platforms

→ Very few accepted benchmarking procedures

Our Approach

- Physical objects widely distributed/available
  - *has been a hurdle limiting previous attempts
  - Available for purchase at cost
    - Free distribution of initial 50 kits
- Sufficient data provided for using the objects
  - Physical properties, RGBD scans, and geometric models
- Detailed task protocols and benchmarking procedures
  - Allow replicable procedures and quantification of performance
  - Give framework and samples, have community contribute/update

*much of this will be an evolving effort with community feedback and interaction

YCB Object Selection Priorities

- Variety:
  - Physical properties: shape, size, deformability, and texture.
  - Grasping and manipulation difficulty: simple geometric shapes (e.g. boxes and spheres) vs. higher shape complexity (e.g. spring clamps, spatula, banana, toys)
- Use: Range of manipulation tasks
  - Simple grasping
  - Common tasks (e.g. pitcher/cup, hammer/nails)
  - Complex assemblies (e.g. toy airplane)
  - Standardized tests (e.g. box and blocks, 9-hole peg)

YCB Object Selection Priorities

- Durability:
  - Avoid objects that are fragile or perishable.
  - To increase longevity, choose objects that are likely to remain in circulation and change little over time.
- Cost:
  - Keep as low as possible to broaden accessibility
    - Commercial items only (i.e. no custom fabrication)
- Portability:
  - Fit in single box within airline size/weight limits

Food Items


CHEEZ-IT
ORIGINAL
Pringles
Master Chef
Domino
SUGAR
40% MORE
French's
Classic YELLOW
JELL-O
JELL-O
SPAM
Campbell's
Tomato SOUP
StarKist

Kitchen Items

Tool Items

Shape Items

Manipulation Tests

Assembly objects

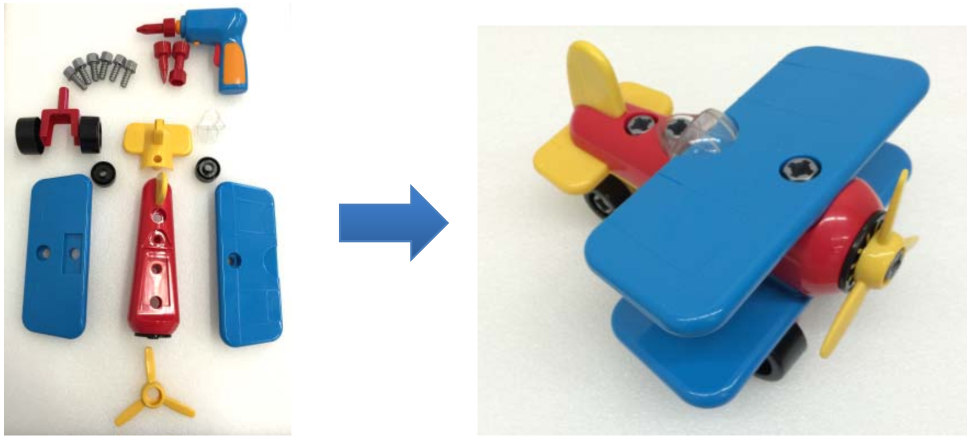

Other YCB Features

- (very) Hi-Res RGB-D scans of the objects
  - Meshes and mesh code available
- Object size/mass properties
- Templates/Instructions for developing detailed task and reporting specifications
  - *Really* need more community involvement in this aspect

How to get it

- www.ycbbenchmarks.org
- Purchase orders via website
  - Object cost + shipping for research labs (~$500)
- Data files (very large) through Amazon

Evaluation Approach 5: Comparison to Human Capabilities

Comparison to Human Performance

- Basic Idea: *Quantify* human hand usage and function
- Compare robot performance to it

The Yale Human Grasping Dataset

Ian M. Bullock, Thomas Feix, and
Aaron M. Dollar

Project Overview

- Subjects wore head-mounted camera
- Recorded and analyzed ~8 hours of dense grasping using per subject
- 4 subjects: 2 Housekeepers, 2 Machinists

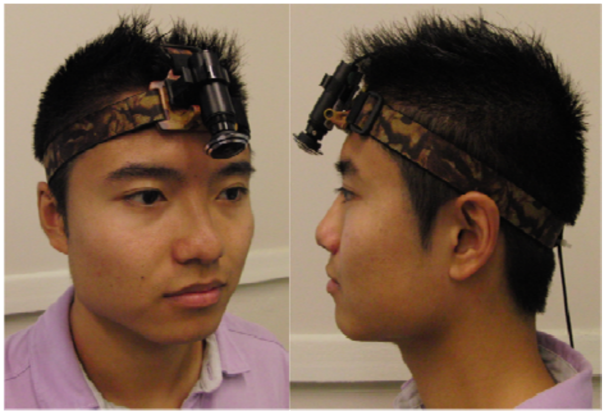

Project Overview

- ~32,000 (?) unique grasps total analyzed
  - Coded manually by trained raters
  - Took a few years ;)
- Recorded:
  - Grasp type
  - Object properties
  - Task type
  - Time duration

Project Overview

- Dataset includes:
  – Video files
  – Video stills for each grasp
  – Lots of coded data for each:

Table 2. Overview of all fields in the tagged dataset

Parameter	Description	Data	
Video	Number of the video file	Video number from 1-179	
Time Stamp	Time stamp of the grasp in the video file	Video timestamp in hh:mm:ss format	
Duration	Length of the grasp instance	Duration in seconds	
Subject	Participant profession and number	Machinist 1/2, Housekeeper 1/2	
BlackRatio	Proportion of instance blacked out for privacy	Ratio between 0 (all visible) and 1 (all black)	
Grasp	The grasp type according to (Feix et al. 2009)	no grasp, one of 33 grasp types	Grasp
OppType	Opposition type of the grasp (Mackenzie & Iberall 1994)	Pad, Palm, Side, NG	
PIP	Power, intermediate or precision grasp	Power, Intermediate, Precision, NG	
Object	High level object name	no object, object name	Object
A	Longest object dimension	Length in cm	
B	Intermediate object dimension	Length in cm	
C	Shortest object dimension	Length in cm	
Grasped Dimension	Dimension along which object is grasped	a/b, a/b/c, b, c, b/c, floppy, CCObj, NG	
Rigidity	Rigidity of the object	rigid, fragile, squeezable, floppy, CCObj, NG	
Roundness	Dimensions along which object is round	a, abc, c, non-round, floppy, CCObj, NG	
Mass	Mass of the object	Value in g	
CCObj	True (1) if Cannot Classify Object (see Section 3)	0, 1, NG	
Shape	Basic shape class, according to (Zingg 1935)	equant, oblate, prolate, bladed, CCObj, NG	
Type	Object type, as defined in (Feix et al. 2013b)	11 object types, CCObj, NG	
Task	High level task name	no task, brief task description	Task
Force	Type of forces required for task	weight, interaction, CCTask, NG	
Constraint	Constraints of the task	11 constraint types, CCTask, NG	
Class	Function class of the task	hold, feel, use, CCTask, NG	
CCTask	True (1) if Cannot Classify Task (see Section 3)	0, 1, NG	

NG = No Grasp, CCTask/CCObj = Cannot Classify

URL

https://www.eng.yale.edu/grablab/humangrasping/

Human Precision Manipulation

- Examine human subjects' ability to manipulate an object grasped in the fingertips
  - Various sizes of objects
  - 2-5 fingers in contact with the object
- Record position and orientation workspace of the object
  - Trakstar sensors
  - Screen-based feedback

Human Precision Manipulation

- Measuring workspace of fingertip-based manipulation

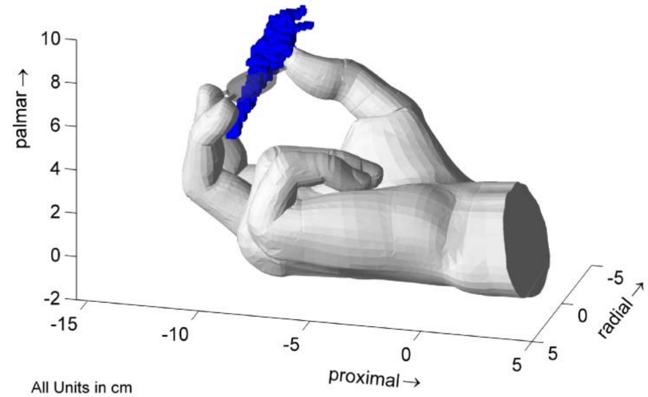

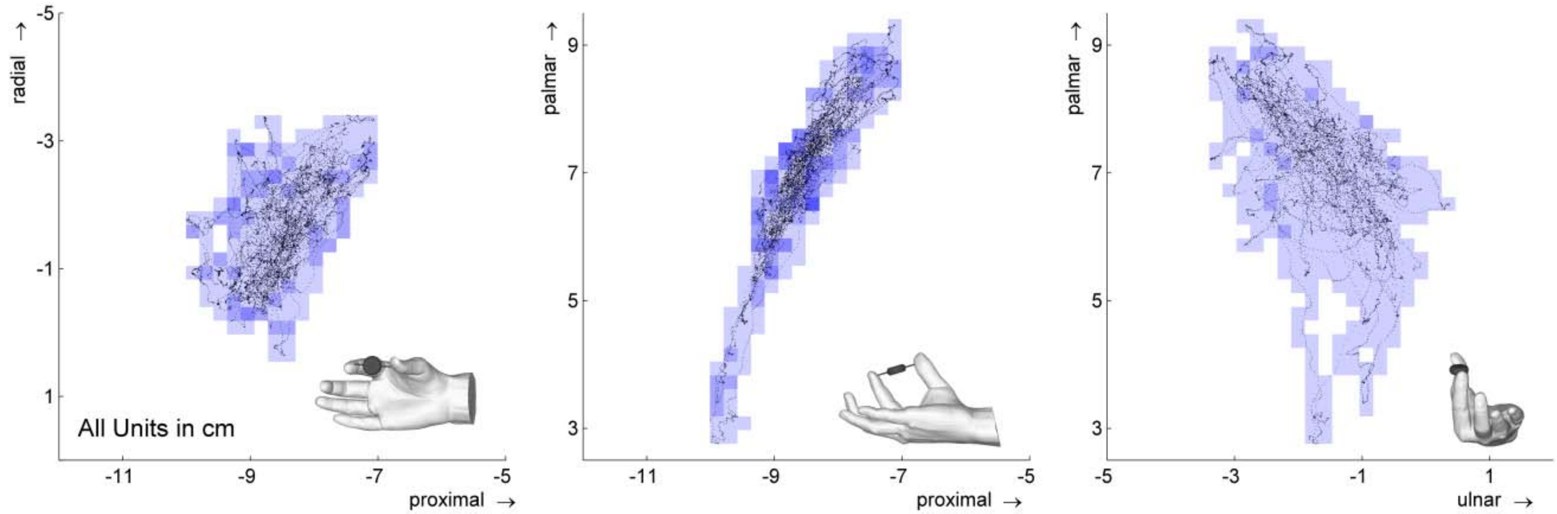

Human Precision Manipulation

- Workspace volume is one possible quantification for comparing performance
  - 2-finger case:
    - 68% of points within 1.05 cm of the centroid and 95% within 2.31 cm
  - 3-finger case:
    - 68% of points within 0.94 cm of the centroid and 95% of points within 2.19 cm

Human Precision Manipulation

- PCA to find major axis
- Major axis direction is significant, others not
  - ~axisymmetric

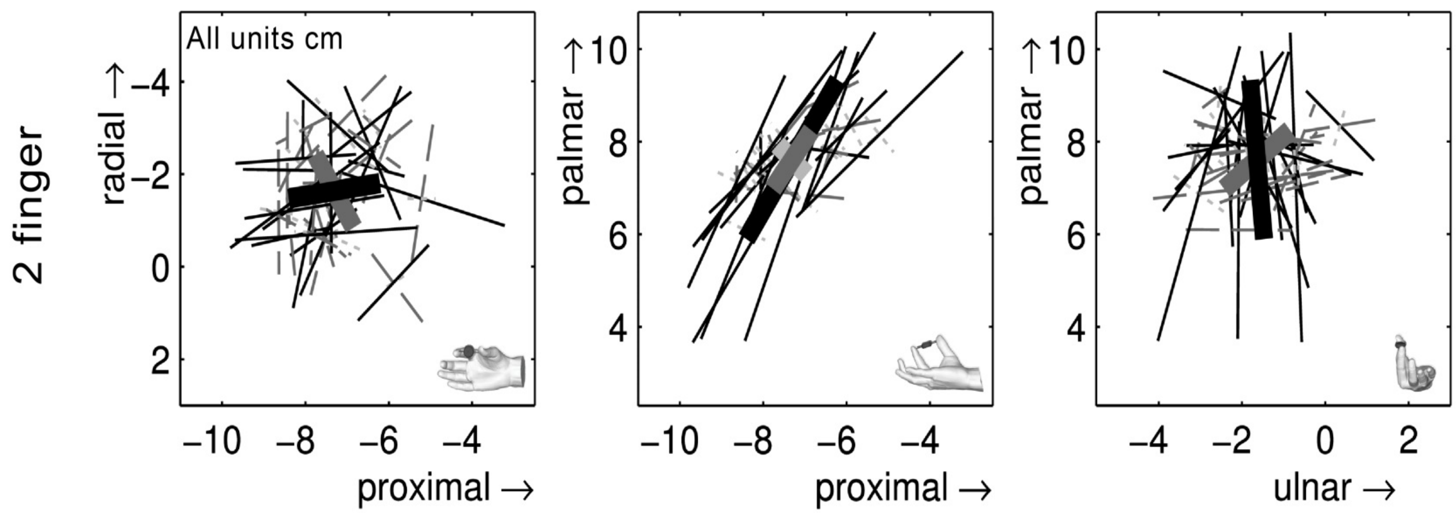

Bullock et. Al, TOH 2015

Questions/Comments?