# A. Unfiltered ORBIT dataset

# A.1. Dataset collection protocol

The unfiltered ORBIT dataset was collected in two phases, the first with blind/low-vision collectors based only in the UK, and the second with blind/low-vision collectors globally [44]. Collectors were recruited via blind charities and networks, and were screened for blindness/low-vision before beginning the collection task. Following this, they were sent instructions for collecting the videos using an accessible iOS app (see Appendix A.2).

## A.1.1 Phase 1 protocol

Collectors were asked to record videos for at least 10 objects, including at least 2 large or immovable objects. For each of these objects, they were asked to record two types of videos, train and test which corresponded to clean and clutter, respectively. Collectors were asked to adopt specific video recording techniques to help capture their objects in-frame.

For clean videos, the collector was asked to place the object on a clear surface with no other objects present and record a video of the object, repeating this on at least 6 different surfaces. Collectors were given step-by-step instructions to record these using a zoom-out technique:

- 1. Keep one hand on the surface next to the object as an anchor point to help aim the camera at the object
- 2. Hold the phone in your other hand and bring it as close as possible to the object
- 3. Start recording, then slowly draw the phone away from the object until it reaches your body at shoulder height
- 4. Rotate the object so that a different side is facing you, while returning the phone close to your anchor hand
- 5. Repeat this at least 3 times for different sides of the object, then stop recording

We helped collectors time each rotation by playing a tick sound every 5 seconds, and asking them to record a new side of the object between ticks. To prevent inadvertently long videos, we automatically ended recordings after 2 minutes.

For large or immovable objects, collectors were asked to position themselves relative to the object so that they were recording different aspects or angles of the object. Again, they were asked to place an anchor hand on three different parts or positions of the object, and then draw their phone slowly towards their body. We asked them to record the object from 3 different angles, repeated this twice, to provide a total of 6 clean videos.

For clutter videos, collectors were asked to construct a cluttered scene based on a real-life situation in which the recognizer might be used — for example, a surface on which someone else may have placed an object. The scene needed to include at least 5 other distractor objects that did not include any of the collector's selected objects. For immov-

able objects like a front door, we asked them to position some likely objects around or in front of it, such as packages, umbrellas, or shoes. Collectors were asked to record their clutter videos using 2 different techniques: a zoom-out and a panning technique, recording at least one video with each technique for each of their objects. For the zoom-out technique, similar to the clean videos, they were asked to place an anchor hand in the scene, and then draw their phone slowly towards their body, before ending the recording. For the panning technique, collectors were asked to place an anchor hand in the scene, making sure the selected object was not directly in front of them. They were then asked to remove their anchor hand, and pan over the scene at shoulder height from right to left by turning their upper body.

### A.1.2 Phase 2 protocol

In Phase 2, we reduced the number of objects to at least 5 per collector, with no specification for large or immovable objects. We also reduced the number of videos to 5 clean (taken on 5 different surfaces) and 2 clutter videos. For clean videos, collectors were asked to use the zoom-out technique, showing at least four sides of the object. For clutter videos, instead of creating a real-world scene as in Phase 1, collectors were asked to record the object, including its surroundings, in a scene where it would normally be found. For simplicity, the panning technique was dropped, and collectors were asked to record both clutter videos using only the zoom-out technique. For the second clutter video, they were simply asked to record it from a different position. Finally, we reduced the cut-off video time to 30 and 20 seconds for clean and clutter videos, respectively.

# A.2. ORBIT Camera iOS app

The dataset was collected using a custom-built iOS app (see Figure A.1) that followed Apple's accessibility guidelines and was tested with blind/low-vision users. On first use, the collector was asked to read an accessible participant information and consent form, and then provide their explicit consent. Once provided, collectors could access the app's home screen and start to add objects to their library. This was done by entering an object's name into a text box at the top of the home screen. This name then became the label for all that object's videos. The collector could then navigate to a second screen to i) record new videos, and ii) review videos already recorded for that object. Here, the collector could also delete and re-record videos already stored. Status indicators were available for each recorded video, including whether it was uploaded to the data server, validated, and visible in the public dataset.

All contributed videos were uploaded to a data server which provided administrative functionality. Administrators were able to set-up data collection periods, review videos uploaded by collectors, and mark them as reviewed and safe for export. All videos were manually checked to ensure i) the labeled object was present in the video, ii) the video did not contain any personally identifiable information (PII), and iii) the object as well as object name were appropriate. If a video did not meet all of these criteria, it was removed from the server and the collector received a notification via the app to re-record the video. All PII obtained via the consent screen was encrypted, and decryption keys were held by administrators outside the server.

Code for the app and back-end data server is open-sourced at https://github.com/orbit-ally.



Figure A.1: ORBIT Camera iOS app: Informed consent screen (left). 'Things' screen showing a collector's personal objects (center-left). 'Thing' screen for a specific object (center-right). 'Thing' screen after recording a video (right).

## A.3. Unfiltered dataset summary

The unfiltered ORBIT dataset contains 4733 videos (3,161,718 frames, 97GB) of 588 objects, collected by 97 people who are blind/low-vision [31]. We summarize it in Table A.1 and describe it in more detail below.

**Number of collectors** 97 blind/low-vision collectors contributed to the unfiltered ORBIT dataset.

Numbers of videos and objects Collectors contributed a total of 588 objects and 4,733 videos. 3,356 of these videos showed the object in isolation, referred to as *clean* videos, while 1,377 showed the object in realistic, multiobject scenes, referred to as *clutter* videos. Of the clutter videos, 873 were recorded with the zoom-out technique, and 504 with the panning technique (see video techniques in Appendix A.1). Each collector contributed on average 6.1 ( $\pm$ 3.8) objects, with 34.6 ( $\pm$ 41.5) clean videos, 9.0 ( $\pm$ 10.5) clutter-zoom-out videos, and 5.2 ( $\pm$ 10.5) clutter-pan videos per object. Figure A.2 shows the number of objects (A.2a) and number of videos per collector (A.2b).

**Types of objects** For summarization purposes, we clustered objects based on object similarity and show their long-tailed distribution in Figure A.7a. The clustering algorithm and full contents of each cluster are included in Appendix D.

## **B. ORBIT benchmark dataset preparation**

The ORBIT benchmark dataset is a subset of the unfiltered ORBIT dataset and was selected to meet the requirements of the ORBIT benchmark (Section 4). The following sections detail the steps to prepare the benchmark dataset.

#### **B.1.** File structure

The following summarizes the file hierarchy of the OR-BIT benchmark dataset:

> mode (train/validation/test)

> video  $oldsymbol{v}_i \in \mathcal{V}_n^{\kappa}$ 

# **B.2. Video filtering**

From the 97 collectors who submitted 4,733 videos (3,356 clean, 873 clutter, 504 clutter-pan) in the unfiltered dataset, the videos of 77 collectors were selected for the benchmark dataset — a total of 3,822 videos (2,996 clean, 826 clutter, 0 clutter-pan) of 486 objects (see Table 2 and Figure 2). The following steps detail the video selection procedure:

- All clutter-pan videos were removed (504 videos, 4 objects, 0 collectors). This was done to maintain consistency because the panning technique was explored in the first, but not the second, phase of the dataset collection. Clutter-zoom-out videos are, therefore, referred to as simply clutter videos in the benchmark dataset.
- Videos shorter than 1 second were removed (20 videos, 0 objects, 0 collectors). Extremely short videos were assumed to be a mistaken recording.
- Objects with < 2 clean videos and < 1 clutter video were remove (387 videos, 101 objects, 20 collectors). These limits were the minimum numbers required for the benchmark evaluation.

### **B.3.** Train/validation/test sets

Of the 77 collectors remaining, 12 contributed only 1 object. To avoid discarding these videos, we merged these collectors together (or with others who had only contributed 2 objects) to enforce a minimum of 3 objects per collector<sup>2</sup>. This yielded an effective total of 67 benchmark users. These 67 users were then split into train/validation/test users  $(44/6/17, respectively^3)$  according to the following criteria:

<sup>&</sup>lt;sup>2</sup>See merged collectors in data/orbit\_benchmark\_users \_to\_split.json in the code repository

 $<sup>^3</sup>See \ \mbox{data/orbit_benchmark_mode_splits.json}$  in the code repository

	Collectors	Objects	Videos	V	ideos per obje	et	]	Frames per vid	eo
				mean/std	25/75th perc.	min/max	mean/std	25/75th perc.	min/max
Total	97	588	4733	8.0/5.5	7.0/8.0	1.0/49.0	668.0/411.1	355.0/898.0	2.0/3600.0
Clean		576	3356	5.8/4.4	5.0/6.0	1.0/44.0	763.0/416.7	504.0/900.0	2.0/3600.0
Clutter-zoom-out		519	873	1.7/1.4	1.0/2.0	1.0/14.0	452.3/270.7	241.0/598.0	3.0/3596.0
Clutter-pan		330	504	1.5/1.6	1.0/1.0	1.0/12.0	408.9/310.4	186.0/521.0	8.0/2684.0
Per-collector	1	6.1/3.8	48.8/53.7	6.7/4.5	5.0/7.7	1.0/41.5	667.2/222.2	550.6/808.4	73.0/1467.2
Clean		6.1/3.8	35.1/41.2	5.0/3.6	4.0/6.0	1.0/36.6	774.9/287.1	648.9/898.7	134.5/1872.6
Clutter-zoom-out		6.0/3.8	10.3/10.6	1.8/1.4	1.0/2.0	1.0/12.0	465.1/191.8	329.2/597.1	73.0/1258.8
Clutter-pan		6.1/3.8	9.8/11.4	1.6/1.6	1.0/1.4	1.0/8.6	419.5/231.4	239.0/574.3	164.5/1208.4

Table A.1: Unfiltered ORBIT dataset.





(b) Number of videos (stacked by object) per collector.

Figure A.2: Number of objects and videos across 97 collectors in the unfiltered ORBIT dataset.

**Validation/test user criteria** We enforced that validation and test users should have at least 5 objects to ensure the test case was sufficiently challenging. In total, 53 out of 67 collectors met this criteria, from which we randomly sampled 17 test users and 6 validation users. Note, one test user (P204) contributed  $\sim 1$  order of magnitude more clutter videos per object than the average test user. Since the CLU- VE evaluation mode averages performance over all clutter videos pooled from all tasks across all test users (i.e.  $\mathcal{T}^{all}$ ), results are likely slightly skewed toward P204. This can be addressed by also reporting the per-user scores as in Figure 4. Note, P204's clean videos were inline with the average, thus CLE-VE performance remains unaffected.

**Train user criteria** The remaining 44 collectors were marked as train users (14 of which did not meet the above criteria). Note, one train user (P587) captured  $\sim$ 1 order of magnitude more clean (though not more clutter) videos than the average train user. Since models are trained on a fixed number of tasks per train user and each task is capped in the number of clean videos sampled per object (see Appendix G.2), the results are not skewed by P587.

## **B.4. Video pre-processing**

Videos were split into frames using ffmpeg at a rate of 30 FPS. Frames were also re-sized from  $1080 \times 1080$  pixels to  $84 \times 84$ . This was done for GPU memory purposes and future work will look toward scaling to larger images.

# C. Extended benchmark dataset summary

In Section 3, we summarized the benchmark dataset over all 67 users. Here, we report the same summaries but now broken down by train/validation/test users.

**Number of videos and objects.** Mirroring Table 2, we report the statistics for train/validation/test users in Table A.2. Extending Figures 2a and 2b, we report the number of objects and videos per user in Figures A.3 and A.4, respectively.

**Types of objects.** Following Appendix D, we include the object cluster histogram for all benchmark users in Figure A.7b, and for train (A.7c), validation (A.7d) and test (A.7e) users. Extending Figure 1, we include more examples of clean and clutter frames from the ORBIT benchmark dataset in Figure A.5.

**Bounding box annotations.** We provide bounding box annotations around the ground-truth object (as per the video label) in all clutter videos. We compute the proportion of each clutter video for which the target object is in-frame (identified by if the frame contained a bounding box), and report these for train/validation/test users in Figure A.6. We use the summarization of test users (Figure A.6c) as an upper bound on model performance in Table 4.

## **D.** Object clustering

The objects submitted by collectors varied widely in highlevel object categories. For summarization purposes, we grouped objects into clusters based on object similarity.

**Clustering algorithm** Each object label (as entered by the user in the app) was converted to lowercase, tokenized and parts-of-speech tagged (using NLTK). A weighted average of all nouns (NN, NNS) in the label was computed (each represented as a FastText [3] embedding)<sup>4</sup> where each weight was the (normalized) frequency of the noun across all nouns in the dataset. Intuitively, this encouraged more

common nouns to dominate the object label's embedding (and down-weighted typos and collectors' idiosyncrasies). We then clustered these embeddings using an agglomerative clustering algorithm based on pair-wise cosine distances (With a distance threshold of 0.4). This produced cluster proposals which we then manually tweaked to obtain the final clusterings shown in Figure A.7.

Clusters for Figure A.7a [97 users, 588 objects, 132 clusters] 1: keys, house keys, radar key, my keys, front door keys, keychain, door keys, key, set of keys, 2: remote, tv remote, apple tv remote, amazon remote control, television remote control, remote control, virgin remote control, tv remote control, samsung tv remote control, presentation remote, sky q remote, clickr, fire stick remote, control, television control, remote tv, remote comtrol, apple tv ii remote control, t. v. remote control, 3: black small wallet, my purse, my wallet, ladies purse, money pouch, coin purse, wallet for bus pass cards and money, i d wallet, ipod in wallet, walletv, wallet, purse, 4: symbol cane, black mobility cane, my cane, white mobility came, folded white cane, long cane, cane, folded cane, folded long guide cane, p939411 white cane, white came, visibility stick, mobility, mobility cane, white cane, 5: earphones, apple headphones, apple earpods, airpods, my airpod pros case, iphone air pods, airpods pro charging case, airpod case, airpods, my airpods, airpods case, earpods, apple airpods case, 6: mug, personal mug, cup again, coffee mug, toddler cup, styrofoam cup, my mug, my cup, 7: front door, back door, house door, front door to house, door, my front door, shed door, front door from outside, 8: blue headphones, my headphones, headphone, bose wireless headphones, headset, my sennheiser pxc 350-2, headphones, 9: phone, mobile phone, iphone 6, apple mobile phone, i phone 11 pro, iphone in case, iphone, cell phone, work phone, phone case, 10: watch, wrist watch, apple watch, apple wath, risk watch, my apple watch, 11: backpack, my work backpack, work bag, bag, shoulder bag, bum bag, back pack, work backpack, sports bag, sport bago, numb bag, 12: sunglasses, my wraparound sunglasses, dark glasses, 13: slippers, nike trainers, my shoes, boot, trainers, trainer shoe, slipper, my trainers, shoes, running shoes, 14: medication, blister pack of painkillers, tablets, money, aspirin vs tylenol, aspirin, pill bottle, my pill dosette, buckleys, box of tablets, 15: victor stream book reader, orbit braille reader and notetaker, victor reader stream, orbit reader 20 braille display, braillepen slim braille keyboard, braille orbit reader, solo audiobook player, braille note, my braille displat, rnib talking book envelope, 16: guide dog play cola, dogs lead, guide dog harness, retractable dog lead, leash, guide dog harness, dog lead, dogs, running tether, 17: blue tooth keyboard, bluetooth keyboard, my keyboard, apple wireless keyboard, mini bluetooth keyboard, portable keyboard, keyboard, 18: my water bottle, bottle, pop bottle, green water bottle, water bottle, bottle of water, clean canteen stainless

<sup>&</sup>lt;sup>4</sup>If the label had no nouns or only 1 word, the whole label was kept.

	Collectors	Objects	Videos	v	ideos per objec	et	I	Frames per vid	eo
				mean/std	25/75 <sup>th</sup> perc.	min/max	mean/std	25/75 <sup>th</sup> perc.	min/max
Total	44	278	2277	8.2/6.0	7.0/7.0	3.0/46.0	716.8/448.6	378.0/898.0	34.0/3600.0
Clean			1814	6.5/6.0	5.0/6.0	2.0/44.0	774.0/471.6	394.5/899.0	34.0/3600.0
Clutter			463	1.7/0.8	1.0/2.0	1.0/7.0	492.7/235.5	324.5/599.0	47.0/2412.0
Per-collector	1	6.3/2.6	51.8/55.2	7.6/4.8	6.5/7.4	3.4/38.4	735.5/232.0	629.7/805.7	213.1/1614.3
Clean			41.2/52.7	5.9/4.8	4.6/6.0	2.4/36.5	823.0/279.0	710.2/898.4	219.3/1872.6
Clutter			10.5/5.2	1.7/0.5	1.2/2.0	1.0/3.0	467.3/163.2	362.4/597.0	177.7/853.0
				(a) 7	Frain users.				
	Collectors	Objects	Videos	v	/ideos per obje	ct	H	Frames per vid	eo
	1			mean/std	25/75 <sup>th</sup> perc.	min/max	mean/std	25/75 <sup>th</sup> perc.	min/max
Total	6	50	347	6.9/1.0	7.0/7.0	4.0/10.0	638.4/227.6	561.0/775.5	33.0/2053.0
Clean			284	5.7/1.0	5.0/6.0	3.0/9.0	689.6/210.3	626.0/835.0	33.0/2053.0
Clutter			63	1.3/0.6	1.0/1.0	1.0/4.0	407.7/144.1	289.0/593.5	176.0/618.0
Per-collector	1	8.3/2.4	57.8/17.7	6.9/0.3	6.7/7.0	6.6/7.6	666.7/95.8	621.0/741.7	529.8/808.4
Clean			47.3/17.5	5.5/0.7	5.0/5.7	4.6/6.6	730.5/114.1	673.9/819.0	565.0/899.9
Clutter			10.5/0.8	1.4/0.4	1.0/1.8	1.0/2.0	409.8/133.2	313.3/528.4	250.5/597.9
				(b) Val	idation users.				
	Collectors	Objects	Videos	v	ideos per obje	et	I	Frames per vid	eo
	I	Ū		mean/std	25/75 <sup>th</sup> perc.	min/max	mean/std	25/75 <sup>th</sup> perc.	min/max
Total	17	158	1198	7.6/2.4	7.0/7.0	3.0/19.0	696.4/384.8	428.0/900.0	40.0/3596.0
Clean			898	5.7/1.0	5.0/6.0	2.0/10.0	791.6/352.7	621.0/926.0	51.0/3596.0
Clutter			300	1.9/2.3	1.0/2.0	1.0/13.0	411.5/332.5	175.0/598.0	40.0/3596.0
Per-collector	1	9.3/2.1	70.5/24.2	7.6/2.1	6.8/7.4	6.2/15.6	733.2/166.6	599.4/811.2	335.7/977.0
Clean			52.8/14.5	5.6/0.6	5.0/6.0	4.4/6.7	804.1/164.7	658.8/899.0	518.1/1095.6
Clutter			17.6/18.7	1.9/2.0	1.0/1.9	1.0/9.9	478.3/285.4	282.3/598.2	158.6/1258.8

(c) Test users.

Table A.2: The ORBIT benchmark dataset, grouped by train, validation, and test users.

steel water bottle, 19: reading glasses, glasses, eye glasses, spectacles, 20: adaptive washing machine, washing machine, tumble dryer, adaptive dryer, tumble dry setting washing machine, washing machine setting, 21: baseball cap, cap, orange skullcap, my tilly hat, hat, my tilly hat upside down, 22: microwave, toaster, kettle, one cup kettle, 23: digital dab radio, dab radio, speaker, radio, my bose bluetooth speaker, 24: wheely bin, wheelie bin, black bin, recycling bin, bin, brown wheelie bin, 25: hairbrush, comb, hair brush, 26: necklace, brown leather bracelet, ladies silver bracelet, favourite earings, silver male wedding band, clear quartz, 27: screwdriver, small space screwdriver, small screwdriver, spanner, wood phillips head, pex plumbers pliers, wood screw phillips head, 28: 12 measuring cup, 13 measuring cup, 14 measuring cup, 1 cup measuring cup, measuring spoon, measuring spoons, 29: lime green marker, yellow marker, pink marker, migenta marker, reptile green marker, pen, 30: digital recorder, dictaphone, pen friend, handheld police scanner, 31: face mask, covid mask, my purple mask, blue facemask, 32: airpod pro, single airpod, apple airpods, pro, 33: wireless earphones, bone conducting headset, bose earpods, my muse s headband, 34: phone stand, nobile phone stand, ipod stand, iphone stand, 35: memory stick, usb c dongle, usb stick, usb, 36: apple phone charger, phone charger, 37: scissors, secateurs, 38: car, my sisters car, 39: socks, sock, 40: fridge, fridge freezer indicator, 41: tv unit, flat screen television, tv, smartty, television, 42: prosecco, jd whisky bottle, bottle of alcoholic drink, vagabond ale bottle, 43: lighter, vape pen, cannabis vape battery, 44: pepper shaker, pink himalayan salt, mediterranean sea salt, salt grinder, 45: sunglasses case, glasses case, eyewear case, 46: lipstick, chap stick, lip balm, 47: tweezers, finger nail clipper, hook, 48: inhaler, insulin pen, spacer, 49: dining table setup, desk, garden table, 50: liquid level indicator, water level sensor, 51: bottle opener, corkscrew, 52: pint glass, wine glass, glass, 53: guide dog, my flee, 54: dog streetball, dog toy, adaptive tennis ball, 55: stylus, apple pencil, 56: my laptop, laptop, 57: ipad, ipad pro, 58: mouse, 59: local post box, post box, 60: black strappy vest, t-shirts, 61: gloves, winter gloves, 62: hand gel, lotion bottle, 63: perfume, deodorant, 64: hair clip, headband, 65: toothbrush, electric toothbrush, 66: alcohol wipe, skip prep, 67: my hearing aid, hearing aid, 68: magnifier, pocket magnifying glass, 69: blue artificial eye, green artificial eye, 70: hand saw, 71: tape measure, ruler, 72: electric sanding disc,



Figure A.3: Number of objects per collector in the ORBIT benchmark dataset, grouped by train, validation, and test users.

miter saw, 73: valve oil, slide grease, 74: cooker, 75: shelf unit with things, wardrobe, 76: knitting basket, washing basket, basket, 77: home phone, cordless phone, 78: skipping rope, spinner, 79: stairgate, back patio gate, 80: large sewing needle, knitting needle, 81: mayonnaise jar, mustard, 82: rice, chicken instant noodles, 83: tea, cranberry cream tea, 84: ottawa bus stop, bus stop sign, 85: journal, book, 86: security fob, 87: money clip, 88: headphone case, 89: ps4 controller, 90: compact disc, 91: garden wall, 92: garden shed, 93: litter and dog waste bin, 94: my clock, 95: sleep mask, 96: glasses cleaning wipe, 97: tissue box, 98: condom box, 99: make up, 100: clear nail varnish, 101: eye drops, 102: battery drill, 103: bed, 104: sofa, 105: cushion, 106: garden bench, 107: mirror, 108: table fan, 109: tred mill, 110: exercise bench, 111: watering can, watering, 112: vase with flowers, 113: veg peeler, 114: sharp knife, 115: wall plug, 116: embroidery thread cone, 117: hole punch, 118: pencil case, 119: paperclips, 120: uno cards, 121: baked bean tin, 122: banana, 123: cappuccino pods, 124: mountain dew can, 125: pinesol cleaner, 126: fish food, 127: cat, 128: dog poo, 129: my slate, 130: av tambourine, 131: grinder, 132: fluffy blanket

Clusters for Figure A.7b [67 users, 486 objects, 122 clusters] 1: keys, house keys, radar key, my keys, front door keys, keychain, door keys, key, set of keys, 2: remote, tv remote, apple tv remote, amazon remote control, television remote control, remote control, virgin remote control, tv remote control, samsung tv remote control, presentation remote, sky q remote, clickr, fire stick remote, control, television control, remote tv, 3: black small wallet, my purse, my wallet, ladies purse, money pouch, coin purse, wallet for bus pass cards and money, i d wallet, ipod in wallet, walletv, wallet, purse, 4: symbol cane, black mobility cane, my cane, white mobility came, folded white cane, long cane, cane, folded cane, folded long guide cane, p939411 white cane, white came, visibility stick, mobility, mobility cane, white cane, 5: earphones, apple headphones, apple earpods, airpods, my airpod pros case, iphone air pods, airpods pro charging case, airpod case, airpods, my airpods, airpods case, earpods, 6: blue headphones, my headphones, headphone, bose wireless headphones, headset, my sennheiser pxc 350-2, headphones, 7: phone, mobile phone, iphone 6, apple mobile phone, i phone 11 pro, iphone in case, iphone, cell phone, work phone, phone case, 8: mug, personal mug, cup again, coffee mug, toddler cup, styrofoam

```
obj 13 clutter = obj 13 clean
                              obj 12 clutter = obj 12 clean
                                                            obj 11 clutter = obj 11 clean
obj 10 clutter = obj 10 clean
                              obj 9 clutter
                                             obj 9 clean
                                                             obj 8 clutter
                                                                            obj 8 clean
               obj 7 clean
                                                             obj 5 clutter
obj 7 clutter
                              obj 6 clutter
                                              obj 6 clean
                                                                            obj 5 clean
obj 4 clutter
               obj 4 clean
                              obj 3 clutter
                                             obj 3 clean
                                                             obj 2 clutter
                                                                            obj 2 clean
obj 1 clutter
               obj 1 clean
```



Figure A.4: Number of clean and clutter videos (stacked by object) per collector in the ORBIT benchmark dataset, grouped by train, validation, and test users.



Figure A.5: Video clips from the ORBIT benchmark dataset. Each row comes from a different video (alternating clean and clutter). Rows 1,2: "dog toy". Rows 3,4: "dog lead". Rows 5,6: "white cane". Rows 7,8: "victor stream book reader". Rows 9,10: "toddler cup". Rows 11,12: "hearing aid".



Figure A.6: The proportion of the video the target object spends in- versus out-of-frame, averaged over all clutter videos per collector and grouped by train, validation, and test users.

cup, my mug, 9: front door, back door, house door, front door to house, door, my front door, shed door, 10: sunglasses, my wraparound sunglasses, dark glasses, 11: watch, wrist watch, apple watch, apple wath, risk watch, my apple watch, 12: victor stream book reader, orbit braille reader and notetaker, victor reader stream, orbit reader 20 braille display, braillepen slim braille keyboard, braille orbit reader, solo audiobook player, braille note, my braille displat, 13: blue tooth keyboard, bluetooth keyboard, my keyboard, apple wireless keyboard, mini bluetooth keyboard, portable keyboard, keyboard, 14: medication, blister pack of painkillers, tablets, money, aspirin vs tylenol, aspirin, pill bottle, my pill dosette, buckleys, 15: backpack, my work backpack, work bag, bag, shoulder bag, bum bag, back pack, work backpack, 16: slippers, nike trainers, my shoes, boot, trainers, trainer shoe, slipper, 17: reading glasses, glasses, eye glasses, 18: baseball cap, cap, orange skullcap, my tilly hat, hat, my tilly hat upside down, 19: guide dog play cola, dogs lead, guide dog harness, retractable dog lead, leash, guide dog harness, dog lead, dogs, 20: my water bottle, bottle, pop bottle, green water bottle, water bottle, 21: hairbrush, comb, hair brush, 22: adaptive washing machine, washing machine, tumble dryer, adaptive dryer, 23: digital recorder, dictaphone, pen friend, handheld police scanner, 24: wheely bin, wheelie bin, black bin, recycling bin, bin, 25: face mask, covid mask, my purple mask, blue facemask, 26: screwdriver, small space screwdriver, small screwdriver, spanner, wood phillips head, pex plumbers pliers, 27: 12 measuring cup, 13 measuring cup, 14 measuring cup, 1 cup measuring cup, measuring spoon, 28: lime green marker, yellow marker, pink marker, migenta marker, reptile green marker, 29: airpod pro, single airpod, apple airpods, pro, 30: wireless earphones, bone conducting headset, bose earpods, my muse s headband, 31: phone stand, nobile phone stand, ipod stand, iphone stand, 32: memory stick, usb c dongle, usb stick, usb, 33: scissors, secateurs, 34: socks, sock, 35: necklace, brown leather bracelet, ladies silver bracelet, favourite earings, 36: microwave, toaster, kettle, one cup kettle, 37: fridge, fridge



(a) All users (unfiltered dataset).



(d) Validation users (benchmark dataset).



Figure A.7: Histogram of object clusters (grouped by object similarity) in the ORBIT dataset (benchmark and unfiltered). Objects in each cluster are listed in Appendix D.

freezer indicator, 38: prosecco, jd whisky bottle, bottle of alcoholic drink, vagabond ale bottle, 39: digital dab radio, dab radio, speaker, 40: apple phone charger, phone charger, 41: sunglasses case, glasses case, eyewear case, 42: lipstick, chap stick, lip balm, 43: tv unit, flat screen television, tv, smarttv, 44: dining table setup, desk, garden table, 45: liquid level indicator, water level sensor, 46: bottle opener, corkscrew, 47: pint glass, wine glass, glass, 48: lighter, vape pen, cannabis vape battery, 49: pepper shaker, pink himalayan salt, mediterranean sea salt, 50: dog streetball, dog toy, adaptive tennis ball, 51: stylus, apple pencil, 52: mouse, 53: car, 54: local post box, post box, 55: black strappy vest, t-shirts, 56: gloves, winter gloves, 57: hand gel, lotion bottle, 58: perfume, deodorant, 59: tweezers, finger nail clipper, 60: hair clip, headband, 61: inhaler, insulin pen, 62: alcohol wipe, skip prep, 63: my hearing aid, hearing aid, 64: magnifier, pocket magnifying glass, 65: hand saw, 66: tape measure, ruler, 67: electric sanding disc, miter saw, 68: cooker, 69: shelf unit with things, wardrobe, 70: knitting basket, washing basket, basket, 71: stairgate, back patio gate, 72: large sewing needle, knitting needle, 73: mayonnaise jar, mustard, 74: rice, chicken instant noodles, 75: tea, cranberry cream tea, 76: ottawa bus stop, bus stop sign, 77: money clip, 78: headphone case, 79: ps4 controller, 80: my laptop, 81: ipad, 82: compact disc, 83: garden wall, 84: garden shed, 85: litter and dog waste bin, 86: my clock, 87: sleep mask, 88: glasses cleaning wipe, 89: tissue box, 90: condom box, 91: make up, 92: clear nail varnish, 93: toothbrush, 94: eye drops, 95: battery drill, 96: bed, 97: sofa, 98: cushion, 99: mirror, 100: table fan, 101: tred mill, 102: exercise bench, 103: skipping rope, 104: watering can, watering, 105: vase with flowers, 106: veg peeler, 107: sharp knife, 108: wall plug, 109: embroidery thread cone, 110: hole punch, 111: pencil case, 112: paperclips, 113: baked bean tin, 114: banana, 115: mountain dew can, 116: pinesol cleaner, 117: fish food, 118: dog poo, 119: my slate, 120: av tambourine, 121: grinder, 122: journal

Clusters for Figure A.7c [44 train users, 278 objects, 100 clusters] 1: keys, house keys, radar key, my keys, front door keys, keychain, door keys, key, set of keys, 2: remote, tv remote, apple tv remote, amazon remote control, television remote control, remote control, virgin remote control, tv remote control, samsung tv remote control, presentation remote, sky q remote, clickr, fire stick remote, control, television control, remote tv, 3: symbol cane, black mobility cane, my cane, white mobility came, folded white cane, long cane, cane, folded cane, folded long guide cane, p939411 white cane, white came, visibility stick, mobility, mobility cane, white cane, 4: earphones, apple headphones, apple earpods, airpods, my airpod pros case, iphone air pods, airpods pro charging case, airpod case, airpods, my airpods, airpods case, earpods, 5: black small wallet, my purse, my wallet, ladies purse, money pouch, coin purse, wallet for

bus pass cards and money, i d wallet, ipod in wallet, wallety, wallet, purse, 6: blue headphones, my headphones, headphone, bose wireless headphones, headset, my sennheiser pxc 350-2, headphones, 7: phone, mobile phone, iphone 6, apple mobile phone, i phone 11 pro, iphone in case, iphone, cell phone, work phone, phone case, 8: sunglasses, my wraparound sunglasses, dark glasses, 9: medication, blister pack of painkillers, tablets, money, aspirin vs tylenol, aspirin, pill bottle, my pill dosette, buckleys, 10: mug, personal mug, cup again, coffee mug, toddler cup, styrofoam cup, my mug, 11: reading glasses, glasses, eye glasses, 12: slippers, nike trainers, my shoes, boot, trainers, trainer shoe, slipper, 13: watch, wrist watch, apple watch, apple wath, risk watch, my apple watch, 14: backpack, my work backpack, work bag, bag, shoulder bag, bum bag, back pack, work backpack, 15: guide dog play cola, dogs lead, guide dog harness, retractable dog lead, leash, guide dog harness, dog lead, dogs, 16: front door, back door, house door, front door to house, door, my front door, shed door, 17: blue tooth keyboard, bluetooth keyboard, my keyboard, apple wireless keyboard, mini bluetooth keyboard, portable keyboard, keyboard, 18: face mask, covid mask, my purple mask, blue facemask, 19: baseball cap, cap, orange skullcap, my tilly hat, hat, my tilly hat upside down, 20: hairbrush, comb, hair brush, 21: victor stream book reader, orbit braille reader and notetaker, victor reader stream, orbit reader 20 braille display, braillepen slim braille keyboard, braille orbit reader, solo audiobook player, braille note, my braille displat, 22: my water bottle, bottle, pop bottle, green water bottle, water bottle, 23: digital recorder, dictaphone, pen friend, handheld police scanner, 24: memory stick, usb c dongle, usb stick, usb, 25: sunglasses case, glasses case, eyewear case, 26: fridge, fridge freezer indicator, 27: adaptive washing machine, washing machine, tumble dryer, adaptive dryer, 28: pint glass, wine glass, glass, 29: airpod pro, single airpod, apple airpods, pro, 30: wireless earphones, bone conducting headset, bose earpods, my muse s headband, 31: mouse, 32: apple phone charger, phone charger, 33: scissors, secateurs, 34: local post box, post box, 35: wheely bin, wheelie bin, black bin, recycling bin, bin, 36: black strappy vest, t-shirts, 37: lipstick, chap stick, lip balm, 38: alcohol wipe, skip prep, 39: my hearing aid, hearing aid, 40: cooker, 41: microwave, toaster, kettle, one cup kettle, 42: bottle opener, corkscrew, 43: large sewing needle, knitting needle, 44: prosecco, jd whisky bottle, bottle of alcoholic drink, vagabond ale bottle, 45: lighter, vape pen, cannabis vape battery, 46: mayonnaise jar, mustard, 47: rice, chicken instant noodles, 48: tea, cranberry cream tea, 49: ottawa bus stop, bus stop sign, 50: money clip, 51: headphone case, 52: stylus, apple pencil, 53: my laptop, 54: compact disc, 55: digital dab radio, dab radio, speaker, 56: phone stand, nobile phone stand, ipod stand, iphone stand, 57: litter and dog waste bin, 58: socks, sock, 59: gloves, winter gloves, 60: sleep mask, 61: glasses cleaning wipe, 62: tissue box, 63: hand gel, lotion bottle, 64: make up, 65: clear nail varnish, 66: hair clip, headband, 67: toothbrush, 68: inhaler, insulin pen, 69: eye drops, 70: necklace, brown leather bracelet, ladies silver bracelet, favourite earings, 71: magnifier, pocket magnifying glass, 72: screwdriver, small space screwdriver, small screwdriver, spanner, wood phillips head, pex plumbers pliers, 73: hand saw, 74: battery drill, 75: tape measure, ruler, 76: electric sanding disc, miter saw, 77: tv unit, flat screen television, tv, smarttv, 78: bed, 79: sofa, 80: dining table setup, desk, garden table, 81: mirror, 82: tred mill, 83: exercise bench, 84: skipping rope, 85: stairgate, back patio gate, 86: liquid level indicator, water level sensor, 87: sharp knife, 88: embroidery thread cone, 89: hole punch, 90: pencil case, 91: baked bean tin, 92: banana, 93: pepper shaker, pink himalayan salt, mediterranean sea salt, 94: mountain dew can, 95: fish food, 96: dog poo, 97: dog streetball, dog toy, adaptive tennis ball, 98: av tambourine, 99: grinder, 100: journal

**Clusters for Figure A.7d** [6 validation users, 50 objects, 36 clusters] 1: keys, house keys, radar key, my keys, front door keys, keychain, door keys, key, set of keys, 2: black small wallet, my purse, my wallet, ladies purse, money pouch, coin purse, wallet for bus pass cards and money, i d wallet, ipod in wallet, walletv, wallet, purse, 3: blue headphones, my headphones, headphone, bose wireless headphones, headset, my sennheiser pxc 350-2, headphones, 4: front door, back door, house door, front door to house, door, my front door, shed door, 5: remote, tv remote, apple tv remote, amazon remote control, television remote control, remote control, virgin remote control, tv remote control, samsung tv remote control, presentation remote, sky q remote, clickr, fire stick remote, control, television control, remote tv, 6: blue tooth keyboard, bluetooth keyboard, my keyboard, apple wireless keyboard, mini bluetooth keyboard, portable keyboard, keyboard, 7: phone, mobile phone, iphone 6, apple mobile phone, i phone 11 pro, iphone in case, iphone, cell phone, work phone, phone case, 8: watch, wrist watch, apple watch, apple wath, risk watch, my apple watch, 9: sunglasses, my wraparound sunglasses, dark glasses, 10: necklace, brown leather bracelet, ladies silver bracelet, favourite earings, 11: tv unit, flat screen television, tv, smarttv, 12: apple phone charger, phone charger, 13: car, 14: slippers, nike trainers, my shoes, boot, trainers, trainer shoe, slipper, 15: socks, sock, 16: my clock, 17: condom box, 18: baseball cap, cap, orange skullcap, my tilly hat, hat, my tilly hat upside down, 19: hairbrush, comb, hair brush, 20: perfume, deodorant, 21: tweezers, finger nail clipper, 22: hair clip, headband, 23: backpack, my work backpack, work bag, bag, shoulder bag, bum bag, back pack, work backpack, 24: symbol cane, black mobility cane, my cane, white mobility came, folded white cane, long cane, cane, folded cane, folded long guide cane, p939411 white cane, white came, visibility stick, mobility, mobility cane white cane, 25: victor

stream book reader, orbit braille reader and notetaker, victor reader stream, orbit reader 20 braille display, braillepen slim braille keyboard, braille orbit reader, solo audiobook player, braille note, my braille displat, 26: magnifier, pocket magnifying glass, 27: guide dog play cola, dogs lead, guide dog harness, retractable dog lead, leash, guide dog harness, dog lead, dogs, 28: shelf unit with things, wardrobe, 29: dining table setup, desk, garden table, 30: vase with flowers, 31: mug, personal mug, cup again, coffee mug, toddler cup, styrofoam cup, my mug, 32: wall plug, 33: paperclips, 34: prosecco, jd whisky bottle, bottle of alcoholic drink, vagabond ale bottle, 35: pinesol cleaner, 36: my slate

Clusters for Figure A.7e [17 test users, 158 objects, 67 clusters] 1: black small wallet, my purse, my wallet, ladies purse, money pouch, coin purse, wallet for bus pass cards and money, i d wallet, ipod in wallet, wallety, wallet, purse, 2: keys, house keys, radar key, my keys, front door keys, keychain, door keys, key, set of keys, 3: remote, tv remote, apple tv remote, amazon remote control, television remote control, remote control, virgin remote control, tv remote control, samsung tv remote control, presentation remote, sky q remote, clickr, fire stick remote, control, television control, remote tv, 4: front door, back door, house door, front door to house, door, my front door, shed door, 5: earphones, apple headphones, apple earpods, airpods, my airpod pros case, iphone air pods, airpods pro charging case, airpod case, airpods, my airpods, airpods case, earpods, 6: victor stream book reader, orbit braille reader and notetaker, victor reader stream, orbit reader 20 braille display, braillepen slim braille keyboard, braille orbit reader, solo audiobook player, braille note, my braille displat, 7: 12 measuring cup, 13 measuring cup, 14 measuring cup, 1 cup measuring cup, measuring spoon, 8: lime green marker, yellow marker, pink marker, migenta marker, reptile green marker, 9: watch, wrist watch, apple watch, apple wath, risk watch, my apple watch, 10: symbol cane, black mobility cane, my cane, white mobility came, folded white cane, long cane, cane, folded cane, folded long guide cane, p939411 white cane, white came, visibility stick, mobility, mobility cane, white cane, 11: screwdriver, small space screwdriver, small screwdriver, spanner, wood phillips head, pex plumbers pliers, 12: mug, personal mug, cup again, coffee mug, toddler cup, styrofoam cup, my mug, 13: blue headphones, my headphones, headphone, bose wireless headphones, headset, my sennheiser pxc 350-2, headphones, 14: blue tooth keyboard, bluetooth keyboard, my keyboard, apple wireless keyboard, mini bluetooth keyboard, portable keyboard, keyboard, 15: phone, mobile phone, iphone 6, apple mobile phone, i phone 11 pro, iphone in case, iphone, cell phone, work phone, phone case, 16: phone stand, nobile phone stand, ipod stand, iphone stand, 17: wheely bin, wheelie bin, black bin, recycling bin, bin, 18: backpack, my work backpack, work bag, bag, shoulder bag, bum bag, back pack, work backpack, 19: adaptive washing machine, washing machine, tumble dryer, adaptive dryer, 20: my water bottle, bottle, pop bottle, green water bottle, water bottle, 21: airpod pro, single airpod, apple airpods, pro, 22: wireless earphones, bone conducting headset, bose earpods, my muse s headband, 23: digital recorder, dictaphone, pen friend, handheld police scanner, 24: digital dab radio, dab radio, speaker, 25: scissors, secateurs, 26: slippers, nike trainers, my shoes, boot, trainers, trainer shoe, slipper, 27: socks, sock, 28: sunglasses, my wraparound sunglasses, dark glasses, 29: baseball cap, cap, orange skullcap, my tilly hat, hat, my tilly hat upside down, 30: microwave, toaster, kettle, one cup kettle, 31: knitting basket, washing basket, basket, 32: liquid level indicator, water level sensor, 33: pepper shaker, pink himalayan salt, mediterranean sea salt, 34: dog streetball, dog toy, adaptive tennis ball, 35: stylus, apple pencil, 36: ps4 controller, 37: ipad, 38: memory stick, usb c dongle, usb stick, usb, 39: car, 40: garden wall, 41: garden shed, 42: gloves, winter gloves, 43: face mask, covid mask, my purple mask, blue facemask, 44: reading glasses, glasses, eye glasses, 45: hand gel, lotion bottle, 46: hairbrush, comb, hair brush, 47: perfume, deodorant, 48: lipstick, chap stick, lip balm, 49: tweezers, finger nail clipper, 50: medication, blister pack of painkillers, tablets, money, aspirin vs tylenol, aspirin, pill bottle, my pill dosette, buckleys, 51: inhaler, insulin pen, 52: necklace, brown leather bracelet, ladies silver bracelet, favourite earings, 53: guide dog play cola, dogs lead, guide dog harness, retractable dog lead, leash, guide dog harness, dog lead, dogs, 54: hand saw, 55: tape measure, ruler, 56: electric sanding disc, miter saw, 57: fridge, fridge freezer indicator, 58: cushion, 59: shelf unit with things, wardrobe, 60: dining table setup, desk, garden table, 61: table fan, 62: stairgate, back patio gate, 63: watering can, watering, 64: veg peeler, 65: bottle opener, corkscrew, 66: prosecco, jd whisky bottle, bottle of alcoholic drink, vagabond ale bottle, 67: lighter, vape pen, cannabis vape battery

# E. Datasheet for the ORBIT Dataset

Here we include a datasheet [10] for the ORBIT dataset.

## E.1. Motivation for dataset creation

Why was the dataset created? (e.g., was there a specific task in mind?; was there a specific gap that needed to be filled?) The ORBIT dataset was created with two aims in mind: 1) to drive research in few-shot object recognition from high-variation videos — compared to existing few-shot learning datasets which contain largely curated images and pose structured benchmark tasks, the ORBIT dataset contains realistic, high-variation videos captured by people who are blind/low-vision on their mobile phones in real-world settings; and 2) to enable the real-world application of teachable object recognisers (TORs) for people who are blind/low-vision.

Has the dataset been used already? If so, where are the results so others can compare (e.g., links to published papers)? No, this is a new dataset. This paper presents the first results on a few-shot learning benchmark task (see Section 4) that was developed on this dataset.

What (other) tasks could the dataset be used for? The dataset could be used to explore a wide range of frame- and video-level recognition tasks including classification (this paper), object detection, object tracking, semantic segmentation, and instance segmentation. Within each of these tasks, the dataset lends itself to exploring the robustness of models (e.g. performance, quantification of uncertainty) in the face of few high-variation examples. Beyond a few-shot setting, it could also be used to explore other learning paradigms, for example, continual learning where new objects are incrementally added over time.

In addition, this dataset supports research from a human and accessibility perspective. It provides evidence of how people who are blind/low-vision take videos, what objects are of interest to them, and how they could be further supported in taking videos. This work could be directly applied in assistive technology to support object recognition, and be extended to other application areas, such as navigation, question-answering, etc.

Who funded the creation of the dataset? The dataset was funded by the Microsoft AI for Accessibility program.

## Any other comment? None.

# E.2. Dataset composition

What are the instances? (that is, examples; e.g. documents, images, people, countries) Are there multiple types of instances? (e.g. movies, users, ratings; people, interactions between them; nodes, edges) The ORBIT dataset is a collection of videos of objects recorded by people who are blind/low-vision on their mobile phones. Each collector recorded 2 types of videos for each object: clean videos which show the object in isolation on a clear surface, and clutter videos which show the object in a realistic, multi-object scene. The clutter videos were recorded with two different techniques: a zoom-out and a panning motion (see Appendix A.1). The dataset is arranged into two parts: the unfiltered ORBIT dataset which contains all videos contributed by collectors, and the ORBIT benchmark dataset which is a subset of the unfiltered dataset and is used to run the benchmark evaluation (see Section 4.2).

How many instances are there in total (of each type, if appropriate)? The unfiltered dataset has 4,733 videos (3,161,718 frames, 97GB) of 588 objects. 3,356 of these videos are clean videos and 1,377 are clutter videos. Of the clutter videos, 873 were recorded with a zoom-out technique, and 504 with a panning technique. The benchmark dataset has 3,822 videos (2,687,934 frames, 83GB) of 486 objects.

2996 of these videos are clean videos and 826 videos are clutter videos. Here, the clutter videos were all recorded with the zoom-out technique. For detailed statistics of the unfiltered and benchmark datasets, see Table A.1 and Table 2, respectively.

What data does each instance consist of? "Raw" data (e.g. unprocessed text or images)? Features/attributes? If the instances related to people, are sub-populations identified (e.g. by age, gender, etc.) and what is their distribution? Each instance is an unprocessed video frame (i.e. a JPEG). Frames are grouped into folders by video. The video folders are then grouped by video type (clean/clutter), then by object, and finally by the (anonymized) collector who contributed them (see file hierarchy in Appendix B.1). Sub-populations of collectors were not identified.

Is there a label or target associated with each instance? If so, please provide a description. All frames from a video are labeled with the (video-level) object name entered by the collector prior to recording the video. All clutter videos were further annotated (offline) with per-frame bounding boxes around the target object.

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 instances (i.e. frames), however, in the collection process, whole videos were removed by a human annotator if: i) the object was not present in the video, ii) the video contained personally identifiable information (PII), or iii) the object or object name were inappropriate. Note, if PII was present in a small number of frames in a video, only these frames were removed.

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. Yes, frames are organized by video, video type, object and finally the (anonymized) collector who contributed them. These relationships are made explicit in the dataset's file structure (see Appendix B.1).

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 dataset was collected by people who are blind/low-vision. Although participation was open to the public, most collectors were recruited via blind/low-vision charities in Level 4<sup>5</sup> countries (primarily UK, Canada, USA, Australia). In addition, the data collection task required an iPhone with iOS  $\geq$  13.2, with sufficient space to store the videos before upload. To upload the videos to the server, collectors required a stable internet connection, and power supply.

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. Yes, the recommended splits for the benchmark dataset are described in Appendix B.3. The unfiltered dataset has no recommended splits as it is not explicitly used for the benchmark evaluation.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description. By design, the dataset contains high-variation videos of objects. Variation comes from objects being fully/partially out-offrame, blurred (due to camera or object motion), too close or far from the camera, obstructed by hands or other objects, or under/overexposed. All videos are unique, but there is always at least one video per unique object. Some objects occur multiple times across different collectors, but each collector never has duplicate objects.

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. The dataset and its additional annotations are self-contained and do not rely on any external resources.

Any other comments? None.

#### **E.3.** Collection process

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? The procedures used to collect the ORBIT data through manual curation are described in Appendix A.1. We validated these procedures through user testing and user feedback. The app

<sup>&</sup>lt;sup>5</sup>https://www.gapminder.org/fw/income-levels

used to collect the data is described in Appendix A.2. It followed Apple's accessibility guidelines and was tested with blind/low-vision users.

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. Videos were recorded and contributed manually by blind/low-vision collectors via a custom-built iOS app (see Appendix A.2).

If the dataset is a sample from a larger set, what was the sampling strategy (e.g. deterministic, probabilistic with specific sampling probabilities)? The dataset was not sampled from a larger set.

Over what time-frame was the data collected? Does this time-frame match the creation time-frame of the data associated with the instances (e.g. recent crawl of old news articles)? If not, please describe the time-frame in which the data associated with the instances was created. The dataset was collected in two phases. The first phase ran from June to August 2020 and involved 49 collectors who were blind/low-vision. The second phase ran from November 2020 to January 2021 and involved 49 collectors who were blind/low-vision. The time-frames for the dataset creation and the data itself overlap.

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 dataset was collected by people who are blind/lowvision. In the first phase, collectors were based in the UK and were incentivized with a £50 voucher for their participation. In the second phase, collectors were based globally and a donation of £25 (or the local currency equivalent) was made to one of five charities chosen by the collector.

### E.4. Data pre-processing

Was any pre-processing/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. During the data collection process, videos were expunged from the server if i) the object was not present in the video, ii) the video contained PII, or iii) the object or object name were inappropriate. All remaining videos formed the unfiltered ORBIT dataset, and were not pre-processed in any other way.

The ORBIT benchmark dataset was constructed from the unfiltered dataset by removing i) clutter videos recorded

with the panning technique, ii) videos shorter than 30 frames, and iii) objects without at least 2 clean videos and 1 clutter video (see further details in Appendix B). Finally, to run the baseline models on the benchmark task, video frames were resized from  $1080 \times 1080$  to  $84 \times 84$  pixels.

Regarding labeling, in addition to the object labels provided by collectors for each video, the clutter videos were further annotated with bounding boxes around the target object. These were provided by a private company with whom we worked closely to ensure consistent and high-quality annotations across the dataset.

Was the "raw" data saved in addition to the pre-processed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the "raw" data. Both the unfiltered and benchmark datasets are available at https://doi.org/10.25383/city.14294597. For download details see the code repository at https://github.com/microsoft/ORBIT-Dataset. Note, videos expunged from the server during the data collection process were not saved.

Is the software used to pre-process/clean/label the instances available? If so, please provide a link or other access point. Pre-processing/cleaning involved removing videos that did not meet specific criteria from the unfiltered dataset to obtain the benchmark dataset (see Appendix B.2). While scripts to do this are not explicitly provided, scripts to download (and resize frames) in the unfiltered and benchmark datasets are provided in the code repository. Annotations beyond the video labels (e.g. bounding boxes) were collected by a private annotation company using proprietary software.

Does this dataset collection/processing procedure achieve the motivation for creating the dataset stated in the first section of this datasheet? If not, what are the limitations? Yes, by collecting videos from blind/low-vision collectors "in situ" and doing minimal pre-processing/cleaning of these videos has yielded a dataset that is highly representative of the high-variation real-world scenarios they encounter. As a result, the unfiltered and benchmark datasets can be used to 1) drive research in few-shot object recognition from high-variation videos, and 2) realize TORs for people who are blind/low-vision.

Any other comments None.

#### E.5. Dataset distribution

How will the dataset be distributed? (e.g., tarball on website, API, GitHub; does the data have a DOI and is it archived redundantly?) The ORBIT dataset is available as 4 zip files (train, validation, test, other) at DOI: https://doi.org/10.25383/city.14294597. We provide versions of each zip file with frames resized to  $224 \times 224$  pixels.

When will the dataset be released/first distributed? What license (if any) is it distributed under? The dataset was released on 7 April 2021 under an CC4.0 license. The code repository is released under an MIT license.

Are there any copyrights on the data? Under the CC4.0 license, the data can be shared and adapted, even commercially. The data must be correctly attributed.

Are there any fees or access/export restrictions? No, the dataset and code are open-source.

Any other comments? None.

### E.6. Dataset maintenance

Who is supporting/hosting/maintaining the dataset? The ORBIT dataset (unfiltered and benchmark) are hosted, and maintained by City, University of London.

Will the dataset be updated? If so, how often and by whom? The dataset may be updated in the future with more videos collected by people who are blind/low-vision.

How will updates be communicated? (e.g. mailing list, GitHub) All updates will be communicated via the dataset's DOI (https://doi.org/10.25383/city.14294597), code repository (https://github.com/microsoft/ORBIT-Dataset), and website (https://orbit.city.ac.uk), including if the dataset becomes obsolete.

Is there a repository to link to any/all papers/systems that use this dataset? No explicit repository has been built for this purpose, however the usage of this dataset can be tracked via the 'cited by' feature available in most referencing tools.

If others want to extend/augment/build on this dataset, is there a mechanism for them to do so? If so, is there a process for tracking/assessing the quality of those contributions. What is the process for communicating/distributing these contributions to users? Individuals are free to extend/augment/build on the dataset. This may include collecting more data/annotations, developing new models, or contributing to code in the repository. For any questions, individuals can reach out to info@orbit.city.ac.uk or use the issue tracker in the code repository. Any future updates will be communicated/distributed via the DOI, code repository, and website.

### E.7. Legal and ethical considerations

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. The pilot study for the ORBIT data collection was approved by the City, University of London, UK Computer Science and Library & Information Science Research Ethics Committee (ETH1920-0331). The first and second phases of the data collection were approved by the City, University of London, UK Computer Science and Library & Information Science Research Ethics Committee (ETH1920-1126 and ETH2021-0032, respectively).

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, all videos were checked by a human annotator during the data collection process and were expunged from the server if they contained any PII or confidential content.

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, all videos were checked by a human annotator during the data collection process and were expunged from the server if they contained any offensive or inappropriate content.

**Does the dataset relate to people? If not, you may skip the remaining questions in this section.** Yes, the dataset was collected by people who are blind/low-vision with the goal of realizing a real-world application, namely TORs.

**Does the dataset identify any sub-populations (e.g., by age, gender)?** If so, please describe how these sub-populations are identified and provide a description of their respective distributions within the dataset. No, no sub-populations are identified in the dataset.

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, collection was completely anonymized. No meta-data or PII was collected that would allow individuals to be identified indirectly.

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. No, the dataset contains no sensitive data.

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

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, collectors were invited to take part in the data collection through advertisements on social media, mailing lists, blogs, and podcasts.

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. Collectors were required to provide informed consent as part of the sign-up procedure and app download. Informed consent covered the following points:

- I confirm that I have read and understood the participant information dated [INSERT DATE AND VERSION NUMBER] for the above study. I have had the opportunity to consider the information and ask questions which have been answered satisfactorily.
- I understand that, if approved, the videos I take with the iPhone app will become part of an open dataset.
- I understand that the researcher might choose not to publish one or more of my videos in the open dataset.
- I understand that my participation is voluntary and that I am free to withdraw without giving a reason, and without being penalized or disadvantaged.
- I agree that, unless I decide to withdraw, the data collected up to this point will be used in the study.
- I agree to City, University of London recording and processing this information about me. I understand that this information will be used only for the purposes explained in the participant information, and my consent is conditional on City, University of London complying with its duties and obligations under the General Data Protection Regulation (GDPR).

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). Collectors were able to delete data through the iOS app before the data collection period ended. Collectors were free to withdraw their consent at any time, and email the researchers to delete their data from the server.

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. A Data Protection Impact Analysis was carried out as part of the Ethics approval process, and the public release of the dataset. It was determined low risk.

#### Any other comments? None.

#### **F.** Episodic & non-episodic learning for TORS

In the (2) **personalize** step, a few-shot recognition model, or a TOR, aims to learn to recognize completely new objects from only a few examples (in our case videos) captured by the user themselves. The recognition model can be trained to do this in a episodic or non-episodic manner (i.e. the (1) **train** step). Here, we describe both classes of approach.

For a user  $\kappa$ , we denote that each of their personal objects  $p \in \mathcal{P}^{\kappa}$  has  $N_p^{\kappa}$  context videos and  $M_p^{\kappa}$  target videos:

$$\overline{\mathcal{V}}_{p}^{\kappa} = \{(\overline{\boldsymbol{v}}, p)_{i}\}_{i=1}^{N_{p}^{\kappa}} \qquad \mathcal{V}_{p}^{\kappa} = \{(\boldsymbol{v}, p)_{i}\}_{i=1}^{M_{p}^{\kappa}} \quad (1)$$

where  $\bar{v}, v$  are context and target videos, respectively, and p is the object label. We group the user's context and target videos over all their  $|\mathcal{P}^{\kappa}|$  objects as:

$$\overline{\mathcal{V}}^{\kappa} = \overline{\mathcal{V}}_{1}^{\kappa} \cup \dots \cup \overline{\mathcal{V}}_{|\mathcal{P}^{\kappa}|}^{\kappa} \qquad \mathcal{V}^{\kappa} = \mathcal{V}_{1}^{\kappa} \cup \dots \cup \mathcal{V}_{|\mathcal{P}^{\kappa}|}^{\kappa}.$$
(2)

## F.1. Episodic few-shot learning

An episodic approach [40, 49, 9, 52, 38] is characterized by training the model in "episodes" or tasks. For a train user  $\kappa \sim \mathcal{K}^{\text{train}}$ , a task is a random sample of their context videos  $\mathcal{C}^{\kappa}$  and target videos  $\mathcal{T}^{\kappa}$ :

$$\mathcal{C}^{\kappa} = \{ (\bar{\boldsymbol{v}}, p)_i \}_{i=1}^N \qquad \mathcal{T}^{\kappa} = \{ (\boldsymbol{v}, p)_i \}_{i=1}^M \qquad (3)$$

where  $\bar{v}, v$  are context and target videos, respectively, and  $p \in \mathcal{P}^{\kappa}$  is the object label. Note,  $\mathcal{C}^{\kappa} \sim \overline{\mathcal{V}}^{\kappa}$  and  $\mathcal{T}^{\kappa} \sim \mathcal{V}^{\kappa}$ , where  $\overline{\mathcal{V}}^{\kappa}$  and  $\mathcal{V}^{\kappa}$  are the total set of the user's context and target videos, respectively (see Appendix G.2 for further task sampling details).

A few-shot recognition model f is trained over many randomly sampled tasks per train user  $\kappa \sim \mathcal{K}^{\text{train}}$ , typically by maximizing the likelihood of predicting the user's correct personal object in each frame in a target video  $v_f \in v$  after seeing only that user's context videos  $C^{\kappa}$ :

$$\theta^{*} = \arg \max_{\theta} \mathbb{E}_{\kappa \sim \mathcal{K}^{\text{train}}} \left[ \mathbb{E}_{\substack{\mathcal{C}^{\kappa} \sim \overline{\mathcal{V}^{\kappa}} \\ \mathcal{T}^{\kappa} \sim \mathcal{V}^{\kappa}}} \left[ \mathbb{E}_{(\boldsymbol{v}, p) \sim \mathcal{T}^{\kappa}} \left[ \sum_{v_{f} \in \boldsymbol{v}} \log p_{\theta^{\kappa}}(y_{f} \mid v_{f}, \mathcal{C}^{\kappa}) \right] \right] \right]$$
(4)

where  $p_{\theta^{\kappa}}(y_f \mid v_f, C^{\kappa})$  is the probability the user's personalized recognition model  $f_{\theta^{\kappa}}$  assigns to the ground-truth object label for frame  $v_f$ . The parameters are typically updated with stochastic gradient descent (SGD) and a cross entropy loss function.

The resulting model  $f_{\theta^*}$  can be thought of as knowing *how* to personalize to a user. That is, for a completely new test user  $\kappa \sim \mathcal{K}^{\text{test}}$  in the real-world,  $\theta^*$  can rapidly be updated to the user's personalized parameters  $\theta^{\kappa}$  with only the context videos  $\mathcal{C}^{\kappa}$  of that user's personal objects. The user can then scan  $f_{\theta^{\kappa}}$  around any new target scenario to identify their objects using Eq. (1).

#### F.2. Non-episodic few-shot learning

A non-episodic approach [51, 46, 5] is characterized by training the model in a standard batch-wise manner. For a train user  $\kappa \sim \mathcal{K}^{\text{train}}$ , a batch  $\mathcal{B}^{\kappa}$  is the union of the user's context and target videos from Eq. (3):

$$\mathcal{B}^{\kappa} = \mathcal{C}^{\kappa} \cup \mathcal{T}^{\kappa} = \{(\boldsymbol{v}, p)_i\}_{i=1}^{N+M}$$
(5)

where v is a context or target video,  $p \in \mathcal{P}^{\text{all}}$  is the object label, and  $\mathcal{P}^{\text{all}}$  is the set of personal objects pooled across all users in  $\mathcal{K}^{\text{train}}$ .

A recognition model f is trained over many randomly sampled batches per train user  $\kappa \sim \mathcal{K}^{\text{train}}$ , typically by maximizing the likelihood of predicting the user's correct personal object in each frame in a target video  $v_f \in v$ :

$$\theta^{*} = \arg \max_{\theta} \mathbb{E}_{\kappa \sim \mathcal{K}^{\text{train}}} \left[ \mathbb{E}_{\substack{\mathcal{C}^{\kappa} \sim \overline{\mathcal{V}}^{\kappa}} \\ \mathcal{T}^{\kappa} \sim \mathcal{V}^{\kappa}} \left[ \mathbb{E}_{(\boldsymbol{v}, p) \sim \mathcal{B}^{\kappa}} \right[ \sum_{v_{f} \in \boldsymbol{v}} \log p_{\theta}(y_{f} \mid v_{f})] \right] \right]$$
(6)

where  $p_{\theta}(y_f \mid v_f)$  is the probability the (generic) model  $f_{\theta}$  assigns to the ground-truth label  $y_f \in \mathcal{P}^{\text{all}}$  for frame  $v_f$ . As above, typically SGD and a cross entropy loss are used.

Unlike an episodic approach, the resulting  $f_{\theta^*}$  is a generic recognizer for all the objects in  $\mathcal{P}^{\text{all}}$  and does not explicitly know how to personalize. Instead, the generic model is personalized to a new test user  $\kappa \sim \mathcal{K}^{\text{test}}$  by fine-tuning  $f_{\theta^*}$  at test-time with the user's context videos  $\mathcal{C}^{\kappa}$  of their personal objects. This is equivalent to transfer learning where the (generic) feature extractor is typically frozen and only a new linear classification layer is learned per user for their objects – together giving the user's personalized parameters  $\theta^{\kappa}$ . As above, the user can then scan  $f_{\theta^{\kappa}}$  around to identify their objects in target scenarios using Eq. (1).

## **G.** Implementation Details

### **G.1.** Baselines

**ProtoNets [40].** ProtoNets is a metric-based few-shot learner which classifies a target frame in a task based on the shortest distance (typically squared Euclidean) between its embedding and the mean embedding of each object class computed using the task's context set. We use a ResNet18 [15] feature extractor pre-trained on ILSVRC [39]. We then train the model episodically on  $T^{\text{train}}$  tasks per train user (drawn from  $\mathcal{K}^{\text{train}}$ ) for 20 epochs with a learning rate of  $10^{-4}$  (reduced by 0.1 for the feature extractor). Note, each epoch samples  $T^{\text{train}}$  new tasks per train user. We evaluate the trained model on all target videos per test user in  $\mathcal{K}^{\text{test}}$  (repeated for  $T^{\text{test}}$  tasks per user), and report the metrics over  $\mathcal{T}^{\text{all}}$ , the flattened set of all target videos across all tasks for all test users.

**CNAPs [38].** CNAPs is an amortization-based few-shot learner which classifies a target frame in a task with a model that has been *generated* by forward passing the task's context set through a hypernet. In practice, the hypernet generates only a subset of the model's parameters following [38] – the task's classifier, and *FiLM* adapters [35] which are placed after convolution blocks in the feature extractor. We use a ResNet18 [15] feature extractor that is pre-trained on ILSVRC [39] and kept frozen. We train the hypernet episodically on  $T^{\text{train}}$  tasks per train user (drawn from  $\mathcal{K}^{\text{train}}$ ) for 10 epochs with a learning rate of  $10^{-4}$ . Note, each epoch samples  $T^{\text{train}}$  new tasks per train user. We evaluate the trained model on all target videos per test user in  $\mathcal{K}^{\text{test}}$  (repeated for  $T^{\text{test}}$  tasks per user), and report the metrics over  $\mathcal{T}^{\text{all}}$ .

MAML [9]. MAML is an optimization-based few-shot learner which classifies a target frame in a task with a model after it has taken G gradient steps on the task's context set. We use a ResNet18 [15] feature extractor that is pre-trained on ILSVRC [39]. We then train the model episodically on  $T^{\text{train}}$  tasks per train user (drawn from  $\mathcal{K}^{\text{train}}$ ) for 10 epochs with the Adam optimizer [20] and a learning rate of  $10^{-5}$  for the outer loop, and the SGD optimizer and a learning rate of  $10^{-3}$  for the inner loop (rates reduced by 0.1 for the feature extractor in both loops). Note, each epoch samples  $T^{\text{train}}$ new tasks per train user. We evaluate the trained model on all target videos per test user in  $\mathcal{K}^{\text{test}}$  (repeated for  $T^{\text{test}}$  tasks per user), and report the metrics over  $\mathcal{T}^{\text{all}}$ . Note that since each (train and test) task can have a different number of objects, we append a new classifier per task, initialized with zeroes following [48]. The classifier is therefore not learned but instead the result after G gradient steps. At both train and test time, G = 15.

FineTuner [46]. FineTuner is a non-episodic few-shot learner which classifies a target frame in a task with a model that has been finetuned using the task's context set, typically with its feature extractor frozen. This is equivalent to a transfer learning approach. We use a ResNet18 [15] feature extractor pre-trained on ILSVRC [39]. Following [46], we first train the extractor using  $\mathcal{K}^{\text{train}}$ : we pool the personal objects across *all* train users in  $\mathcal{K}^{\text{train}}$  and train a standard classification model<sup>6</sup> for 10 epochs using  $T^{\text{train}}$  tasks per train user. Note, each epoch samples T<sup>train</sup> new tasks per train user. Because the model is non-episodic, the context and target videos of each task are pooled and treated as a standard batch, following Eq. (5). We use the Adam optimizer [20] and a learning rate of  $10^{-4}$  (reduced by 0.1 for the feature extractor). At test time, this feature extractor is frozen, and for each test user drawn from  $\mathcal{K}^{\text{test}}$ , a new classifier is appended, initialized with zeroes following [48] and tuned for G = 50 gradient steps on the user's context videos, using SGD and a learning rate of 0.1. We evaluate the trained

<sup>&</sup>lt;sup>6</sup>Because there are similar objects across the train users, we use the objects' cluster labels (i.e. 100-way classification task), as in Figure A.7c

model on all target videos per test user in  $\mathcal{K}^{\text{test}}$  (repeated for  $T^{\text{test}}$  tasks per user), and report the metrics over  $\mathcal{T}^{\text{all}}$ .

#### G.2. Sampling tasks

The only strict requirement of the ORBIT evaluation protocol is that a model outputs a prediction for every frame in every target video for every test user. How tasks are sampled from context/target videos during training, and context videos during testing is otherwise flexible. In this section, we detail how we sample tasks, but acknowledge that other strategies could be chosen.

**Train user tasks** We sample  $T^{\text{train}}$  tasks (per epoch) for each train user  $\kappa \sim \mathcal{K}^{\text{train}}$  as follows:

- 1. Way: sample a set of the user's objects  $\ddot{\mathcal{P}}^{\kappa} \sim \mathcal{P}^{\kappa}$
- 2. Shot: for each object  $p \in \ddot{\mathcal{P}}^{\kappa}$ , sample  $n_p \sim [1, \ldots, N_n^{\kappa}]$ and  $m_p \sim [1, \ldots, M_p^{\kappa}]^7$
- 3. Construct context set  $\mathcal{C}^{\kappa} = \{(\bar{\boldsymbol{v}}, p) \stackrel{n_p}{\sim} \overline{\mathcal{V}}_p^{\kappa} \mid \forall p \in \ddot{\mathcal{P}}^{\kappa}\}$ 4. Construct target set  $\mathcal{T}^{\kappa} = \{(\boldsymbol{v}, p) \stackrel{m_p}{\sim} \mathcal{V}_p^{\kappa} \mid \forall p \in \ddot{\mathcal{P}}^{\kappa}\}$

For memory reasons, during training we impose a video cap per object of  $n_p = 5$  and  $m_p = 4$  if  $|\ddot{\mathcal{P}}^{\kappa}| \ge 6$  (otherwise the caps are doubled) and an object cap of  $|\ddot{\mathcal{P}}^{\kappa}| = 10$ .

**Test user tasks** We sample  $T^{\text{test}}$  tasks (per epoch) for each test user  $\kappa \sim \mathcal{K}^{\text{test}}$  as follows:

- 1. Way: select *all* the user's objects,  $\ddot{\mathcal{P}}^{\kappa} = \mathcal{P}^{\kappa}$  (i.e. test personalization for all objects)
- 2. Shot: for each object  $p \in \ddot{\mathcal{P}}^{\kappa}$ ,  $n_p = N_p^{\kappa}$  and  $m_p = M_p^{\kappa}$ (i.e. use all context and target videos)
- 3. Construct context set  $\mathcal{C}^{\kappa} = \{ (\bar{\boldsymbol{v}}, p) \stackrel{n_p}{\sim} \overline{\mathcal{V}}_p^{\kappa} \mid \forall p \in \ddot{\mathcal{P}}^{\kappa} \}$
- 4. Construct target set  $\mathcal{T}^{\kappa} = \{(\boldsymbol{v}, p) \stackrel{m_p}{\sim} \mathcal{V}_p^{\kappa} \mid \forall p \in \mathcal{P}^{\kappa}\}$

For testing, no caps are imposed on the number of videos per object, or number of objects.

## G.3. Task sampling hyper-parameters

The number of tasks sampled per train user (per epoch) was selected over a range of 5 to 500 (see Table A.7). The number of tasks sampled per test (and validation) user was selected to reduce variance in the reported metrics. The number of clips per video and the clip length were selected based on available GPU memory during training/testing  $(2 \times$ Dell NVIDIA Tesla V100 32GB GPUs).

- Tasks per train user,  $T^{\text{train}} = 50$
- Tasks per test (and validation) user,  $T^{\text{test}} = 5$
- Clip length, L = 8 frames
- Clips per context video in a train task,  $S_{\mathcal{C}}^{\text{train}} = 4$
- Clips per target video in a train task,  $S_T^{\text{train}} = 4$
- Clips per context video in a test task,  $S_C^{\text{test}} = 8$

#### G.4. Optimization hyper-parameters

The learning rates were selected with a grid search over the range  $10^{-5}$  to  $10^{-1}$ . Unless otherwise specified, for all baselines we use the Adam optimizer [20] with default parameters  $\beta = (0.9, 0.999)$ ,  $\epsilon = 10^{-8}$ , and no weight decay.

# H. Extended analyses

Here we extend analyses from Section 5.2 in the main paper.

Meta-training on other few-shot learning datasets. Table A.3 shows the results when meta-training on Meta-Dataset and meta-testing on the ORBIT test users, for each of the baseline models (extends Table 5 in the main paper).

User-centric versus user-agnostic training. The baselines are trained in a user-centric manner whereby each task is sampled from only one train user's objects. We compare this to a user-agnostic training regime in Table A.4 where we pool objects across *all* train users in  $\mathcal{K}^{\text{train}}$  and construct train tasks by randomly sampling objects from this pool (keeping all other sampling procedures the same). This is akin to typical few-shot learning pipelines. The performance of both regimes is overall equivalent, suggesting that imposing a user-centric hierarchy on task sampling does not negatively impact performance, relative to imposing no hierarchy.

Per-user performance. Extending Figure 4 in the main paper, we plot CLU-VE performance across the full suite of metrics (frame accuracy, frames-to-recognition and video accuracy) per user across all baseline models in Figure A.10. Here, we see the same trend as the main paper - personalization performance varies widely across users.

Train task composition. In Figures 5 and 6 in the main paper, we investigate the impact of the number of context videos per object, and the number of objects per user, sampled in train tasks on CLU-VE frame accuracy. We extend these plots to the frames-to-recognition and video accuracy metrics in Figure A.8 and Figure A.9 and complement them with their corresponding tables in Tables A.5 and A.6.

Number of tasks per train user. In Table A.7), we investigate the impact of the number of tasks sampled per train user (per epoch), T<sup>train</sup> during meta-training, and observe that CLU-VE performance increases with more tasks, but levels off at around 50 tasks per user. We expect that bigger gains could be achieved by accounting for the informativeness of frames, instead of uniformly sampling frames from videos. This is likely because clean and clutter videos are noisy and contain many redundant frames.

<sup>&</sup>lt;sup>7</sup>Note, N and M in Eq. (3) and Eq. (5) are  $N = \sum_{p \in \vec{\mathcal{P}^{\kappa}}} n_p$  and  $M = \sum_{p \in \vec{\mathcal{P}^{\kappa}}} m_p$ , respectively.

	<b>Clean Video Evaluation (CLE-VE)</b>			<b>Clutter Video Evaluation (CLU-VE)</b>		
MODEL	FRAME ACC	FTR	VIDEO ACC	FRAME ACC	FTR	VIDEO ACC
ProtoNets [40]	58.98 (2.23)	11.55 (1.79)	69.17 (3.01)	46.97 (1.84)	20.42 (1.71)	52.80 (2.53)
CNAPs [38]	51.86 (2.49)	20.81 (2.33)	60.77 (3.18)	41.59 (1.94)	30.72 (2.13)	43.00 (2.51)
MAML [9]	42.55 (2.67)	37.28 (2.99)	46.96 (3.25)	24.35 (1.83)	62.30 (2.34)	25.73 (2.21)
FineTuner [46]	61.01 (2.24)	11.53 (1.82)	72.60 (2.91)	48.45 (1.86)	19.13 (1.69)	54.07 (2.52)

Table A.3: Models perform poorly when meta-trained on Meta-Dataset and meta-tested on ORBIT test users, suggesting that existing few-shot datasets may be insufficient for real-world adaptation. Results are reported as the average (95% confidence interval) over all target videos pooled from 85 test tasks (5 tasks per test user, 17 test users).

	<b>Clean Video Evaluation</b> (CLE-VE)			Clutter Video Evaluation (CLU-VE)		
MODEL	FRAME ACC	FTR	VIDEO ACC	FRAME ACC	FTR	VIDEO ACC
ProtoNets [40]	65.31 (2.10)	7.44 (1.42)	78.67 (2.67)	49.95 (1.76)	14.87 (1.50)	59.07 (2.49)
CNAPs [38]	67.21 (2.14)	9.30 (1.59)	79.56 (2.63)	51.26 (1.76)	15.36 (1.54)	60.13 (2.48)
MAML [9]	68.98 (2.16)	9.54 (1.68)	79.45 (2.63)	50.95 (1.92)	21.67 (1.89)	55.00 (2.52)
FineTuner [46]	69.01 (2.19)	7.56 (1.48)	78.45 (2.68)	53.17 (1.85)	15.32 (1.57)	62.73 (2.45)

Table A.4: User-agnostic training on the ORBIT dataset (i.e. a train task is sampled from objects pooled across *all* users in  $\mathcal{K}^{\text{train}}$ ) performs as well as user-centric training (i.e. a train task is sampled from only one user at a time). Results are reported as the average (95% confidence interval) over all target videos pooled from 85 test tasks (5 tasks per test user, 17 test users).



Figure A.8: Meta-training with more context videos per object leads to better CLU-VE performance (frame accuracy - left; frames-to-recognition - center; video accuracy - right). Frames are sampled from an increasing number of clean videos per object using the number of clips per video ( $S^{train}$ ) to keep the total number of context frames fixed per train task. Extends Figure 5 in the main paper.



Figure A.9: Meta-training and -testing with more objects per user poses a harder recognition problem (solid line), however, meta-training with fewer objects than encountered at meta-testing (dashed line) shows only a small CLU-VE performance drop compared to Table 4, suggesting that models may be able to adapt to more objects in the real-world (frame accuracy - left; frames-to-recognition - center; video accuracy - right). Extends Figure 6 in the main paper.





(a) Results with CNAPs [38].



0.35 0.3 0.3 0.25 0.

Accuracy of the second second



accuracy

rame

(b) Results with ProtoNets [40].



(c) Results with MAML [9].







(d) Results with FineTuner [46].

Figure A.10: CLU-VE performance varies widely across test users, for all metrics, across all baseline models. Error bars are 95% confidence intervals.

$N_p^{\kappa}$	$S_{\mathcal{C}}^{\mathrm{train}}$	L	FRAME ACC	FTR	VIDEO ACC
1	12	8	40.53 (1.84)	27.49 (2.00)	44.27 (2.51)
2	6	8	46.19 (1.89)	22.59 (1.83)	52.40 (2.53)
4	3	8	49.23 (1.88)	20.69 (1.80)	55.40 (2.52)
6	2	8	49.44 (1.85)	20.83 (1.78)	57.00 (2.51)

(a) Results with CNAPs [38].

$N_p^{\kappa}$	$S_{\mathcal{C}}^{\mathrm{train}}$	L	FRAME ACC	FTR	VIDEO ACC
1	12	8	37.79 (1.72)	26.44 (1.91)	42.13 (2.50)
2	6	8	42.66 (1.78)	23.00 (1.82)	49.07 (2.53)
4	3	8	46.87 (1.76)	20.17 (1.70)	55.53 (2.51)
6	2	8	46.92 (1.79)	19.36 (1.67)	54.47 (2.52)

(b) Results with ProtoNets [40].

$N_p^{\kappa}$	$S_{\mathcal{C}}^{\mathrm{train}}$	L	FRAME ACC	FTR	VIDEO ACC
1	12	8	45.37 (1.94)	28.03 (2.02)	50.07 (2.58)
2	6	8	49.04 (1.90)	23.62 (1.91)	55.21 (2.57)
4	3	8	50.32 (1.88)	20.42 (1.81)	55.83 (2.56)
6	2	8	47.91 (1.95)	22.72 (1.92)	52.29 (2.58)

(c) Results with MAML [9].

$N_p^{\kappa}$	$S_{\mathcal{C}}^{\mathrm{train}}$	L	FRAME ACC	FTR	VIDEO ACC
1	12	8	38.28 (1.98)	35.76 (2.25)	40.62 (2.54)
2	6	8	44.95 (1.97)	26.41 (2.06)	48.47 (2.58)
4	3	8	49.03 (1.92)	20.12 (1.83)	53.75 (2.58)
6	2	8	51.71 (1.89)	19.18 (1.78)	57.71 (2.55)

(d) Results with FineTuner [46].

Table A.5: Meta-training with more context videos per object leads to better CLU-VE performance. Frames are sampled from an increasing number of clean videos per object  $(N_p^{\kappa})$  using the number of clips per video  $(S_c^{\text{train}})$  to keep the total number of context frames fixed per train task. Corresponds to Figure 5 in the main paper.

#objs/ train user	#objs/ test user	FRAME ACC	FTR	VIDEO ACC
2	2	77.08 (2.98)	4.10 (1.41)	84.17 (3.77)
4	4	63.85 (2.45)	9.61 (1.61)	72.17 (3.34)
6	6	57.10 (2.14)	11.57 (1.57)	63.72 (2.99)
8	8	54.00 (1.96)	15.97 (1.69)	62.84 (2.62)
2	no-cap	45.38 (1.94)	24.08 (1.97)	49.51 (2.58)
4	no-cap	50.91 (1.85)	17.24 (1.62)	59.51 (2.54)
6	no-cap	50.57 (1.89)	17.38 (1.69)	57.29 (2.55)
8	no-cap	53.68 (1.90)	16.29 (1.63)	62.78 (2.50)

#### (a) Results with CNAPs [38].

#objs/ train user	#objs/ test user	FRAME ACC	FTR	VIDEO ACC
2	2	78.38 (2.45)	4.10 (1.08)	86.94 (3.48)
4	4	56.36 (2.53)	16.44 (2.22)	64.78 (3.56)
6	6	54.28 (2.09)	13.94 (1.71)	61.91 (3.02)
8	8	48.34 (1.87)	17.07 (1.64)	53.79 (2.71)
2	no-cap	45.24 (1.75)	19.61 (1.66)	51.67 (2.58)
4	no-cap	46.09 (1.75)	20.03 (1.69)	54.51 (2.57)
6	no-cap	48.57 (1.76)	17.59 (1.61)	56.46 (2.56)
8	no-cap	48.62 (1.77)	17.67 (1.61)	55.14 (2.57)

(b) Results with ProtoNets [40].

#objs/ train user	#objs/ test user	FRAME ACC	FTR	VIDEO ACC
2	2	78.98 (2.69)	6.74 (1.94)	84.72 (3.72)
4	4	61.82 (2.59)	13.55 (2.11)	66.81 (3.51)
6	6	58.64 (2.27)	16.28 (1.96)	64.32 (2.98)
8	8	51.01 (1.99)	19.89 (1.87)	55.63 (2.70)
2	no-cap	48.34 (1.84)	21.27 (1.79)	52.33 (2.53)
4	no-cap	51.55 (1.89)	18.73 (1.73)	56.87 (2.51)
6	no-cap	51.59 (1.87)	19.12 (1.75)	56.33 (2.51)
8	no-cap	51.75 (1.87)	19.81 (1.78)	56.80 (2.51)

#### (c) Results with MAML [9].

#objs/ train user	#objs/ test user	FRAME ACC	FTR	VIDEO ACC
2	2	79.68 (2.97)	3.65 (1.73)	81.11 (4.04)
4	4	62.47 (2.61)	13.77 (2.27)	69.42 (3.44)
6	6	58.86 (2.31)	16.97 (2.06)	63.72 (2.99)
8	8	52.34 (2.00)	18.46 (1.87)	56.25 (2.69)
2	no-cap	50.62 (1.90)	20.69 (1.85)	55.40 (2.52)
4	no-cap	52.39 (1.92)	19.27 (1.80)	57.67 (2.50)
6	no-cap	51.33 (1.88)	18.94 (1.78)	55.73 (2.51)
8	no-cap	52.14 (1.90)	19.10 (1.78)	56.47 (2.51)

(d) Results with FineTuner [46].

Table A.6: Meta-training and -testing with more objects per user poses a harder recognition problem, however, meta-training with fewer objects than encountered at meta-testing shows only a small CLU-VE performance drop compared to Table 4, suggesting that models may be able to adapt to more objects in the real-world. Corresponds to Figure 6 in the main paper.

$T^{\text{train}}$	FRAME ACC	FTR	VIDEO ACC
5	48.08 (1.85)	22.21 (1.84)	55.07 (2.52)
25	51.85 (1.82)	16.58 (1.59)	60.47 (2.47)
50	51.47 (1.81)	17.87 (1.69)	59.53 (2.48)
100	51.24 (1.86)	18.96 (1.71)	59.53 (2.48)
200	50.83 (1.82)	20.25 (1.72)	60.20 (2.48)
500	49.72 (1.81)	18.83 (1.69)	56.07 (2.51)

(a) Results with CNAPs [38].

$T^{\text{train}}$	FRAME ACC	FTR	VIDEO ACC
5	43.90 (1.71)	19.42 (1.63)	49.07 (2.53)
25	46.89 (1.76)	20.40 (1.72)	54.00 (2.52)
50	50.34 (1.74)	14.93 (1.52)	59.93 (2.48)
100	49.59 (1.79)	19.38 (1.64)	59.33 (2.49)
200	49.78 (1.74)	15.51 (1.46)	58.20 (2.50)
500	48.70 (1.79)	18.28 (1.61)	54.67 (2.52)

(b) Results with ProtoNets [40].

$T^{\text{train}}$	FRAME ACC	FTR	VIDEO ACC
5	44.33 (2.05)	32.60 (2.20)	47.27 (2.53)
25	51.52 (1.89)	19.16 (1.75)	57.33 (2.50)
50	51.67 (1.86)	19.74 (1.76)	57.47 (2.50)
100	51.64 (1.86)	20.09 (1.80)	56.87 (2.51)
200	50.61 (1.92)	21.64 (1.85)	56.13 (2.51)
500	49.78 (1.89)	21.86 (1.84)	56.27 (2.51)

(c) Results with MAML [9].

$T^{\mathrm{train}}$	FRAME ACC	FTR	VIDEO ACC
5	48.99 (1.82)	19.32 (1.75)	56.47 (2.51)
25	51.90 (1.87)	19.71 (1.81)	58.13 (2.50)
50	52.52 (1.88)	18.60 (1.76)	58.00 (2.50)
100	52.20 (1.89)	17.99 (1.72)	57.13 (2.50)
200	54.79 (1.87)	16.86 (1.67)	61.13 (2.47)
500	52.90 (1.88)	18.18 (1.74)	57.87 (2.50)

(d) Results with FineTuner [46].

Table A.7: Sampling more tasks per train user has limited benefits on CLU-VE performance beyond  $T^{\text{train}} = 50$  suggesting that sampling techniques that account for frame informativeness may be required.