

# Audio classification of the content of food containers and drinking glasses

Santiago Donaher  
Centre for Intelligent Sensing  
Queen Mary University of London  
London, UK  
s.donaher@qmul.ac.uk

Alessio Xompero  
Centre for Intelligent Sensing  
Queen Mary University of London  
London, UK  
a.xompero@qmul.ac.uk

Andrea Cavallaro  
Centre for Intelligent Sensing  
Queen Mary University of London  
London, UK  
a.cavallaro@qmul.ac.uk

**Abstract**—Food containers, drinking glasses and cups handled by a person generate sounds that vary with the type and amount of their content. In this paper, we propose a new model for sound-based classification of the type and amount of content in a container. The proposed model is based on the decomposition of the problem into two steps, namely action recognition and content classification. We use the scenario of the recent CORSMAL Containers Manipulation dataset and consider two actions (shaking and pouring), and seven combinations of material and filling level. The first step identifies the action performed by a person with the container. The second step determines the amount and type of content using an action-specific classifier. Experiments show that the proposed model achieves 76.02, 78.24, and 41.89 weighted average  $F_1$  score on the three test sets, respectively, and outperforms baselines and existing approaches that classify the content amount and type either independently or jointly.

**Index Terms**—Audio classification, Deep Neural Networks, Object properties.

## I. INTRODUCTION

The contactless estimation of the weight of a food container or a drinking glass handled by a person is important to support human-robot collaborations when prior knowledge of its content is unavailable [1]–[3]. The object weight (and hence the force to be applied by the robot hand or gripper for a human-to-robot handover) varies with the amount and type of content. We are interested in inferring this information from the sound generated by the contact of the content with the surfaces of the container during a pouring or a shaking action [4]–[7].

Previous work on audio regression of content measured the air volume variations within a container [7], the vibrations of liquid poured from a bottle [5] and the vibrations in the acoustic frequency generated by a robot shaking and pouring granular materials from and into a tub [6]. Clarke *et al.* [6] evaluated multiple models, such as a linear regressor, a convolutional neural network, and recurrent neural networks based on either a Long Short-Term Memory (LSTM) module or a Gated Recurrent Unit (GRU) module, to regress the amount of content shaken or poured, using spectrograms as input. Liang *et al.* [7] regressed the level of the remaining air

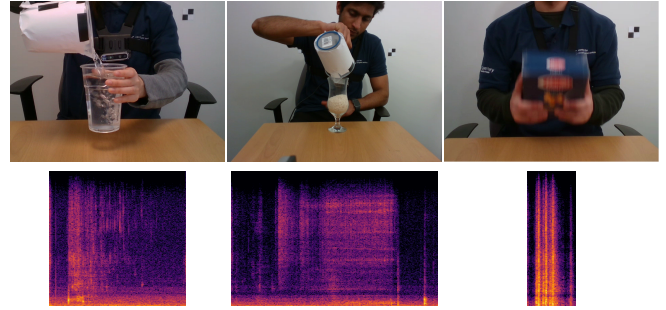


Fig. 1. Sample spectrograms (bottom) of the sound produced by the action of pouring content into a drinking glass (left, centre) and shaking a food container (right).

column in a container (e.g. glass, cup, thermo, mug) during the pouring of a liquid (water, juice) with a LSTM-based recurrent neural network. To handle noisy scenarios, Liang *et al.* [8] regressed the liquid level by conditioning on both audio and force/torque data of three containers, and the trained model generalises to different containers. However, these works specifically addressed the pouring action of containers standing on surfaces. We are instead interested in the problem of classifying the content of a range of food containers and drinking glasses *handled by a person* (see Fig. 1).

In this paper, we propose an audio classification model to recognise the properties of the content of containers manipulated by a person. We address the problem by first recognising the action occurring and then identifying an action-specific classifier to determine a set of seven combinations of content type and level. We validate the proposed model on the CORSMAL Containers Manipulation (CCM) dataset [9], the only publicly available dataset for the recognition of the properties of an object manipulated by a person. In addition to CCM, we evaluate the proposed model on another set of audio recordings acquired in a different environment and with a different setup<sup>1</sup>.

## II. PROPOSED METHOD


Let us consider a person that interacts with a container by, for example, shaking a food box or pouring content into a

This work is supported by the CHIST-ERA programme through the project CORSMAL, under UK EPSRC grant EP/S031715/1.

<sup>1</sup>Code, pre-trained models and data are available at: [http://corsmal.eecs.qmul.ac.uk/audio\\_classification.html](http://corsmal.eecs.qmul.ac.uk/audio_classification.html).

TABLE I

THE MASS OF OBJECTS (CONTAINER AND FILLING) IN THE TRAINING SPLIT OF THE CCM DATASET, AND THE CAPACITY (CAP.) OF EACH CONTAINER. THE CLASS EMPTY CORRESPONDS TO THE MASS OF THE CONTAINER, WHICH IS KNOWN. KEY – CX: CONTENT TYPE (C) AND LEVEL (X), WHERE C CAN BE PASTA (P), RICE (R), OR WATER (W); AND X CAN BE HALF-FULL (5) OR FULL (9).



Empty	2	3	10	15	86	134	15	31	59
P5	37	39	119	37	131	205	104	297	776
P9	65	69	205	54	175	262	174	509	1350
R5	78	85	232	68	208	280	208	633	1548
W5	95	104	270	79	234	316			
R9	139	151	409	110	305	396	362	1115	2739
W9	169	185	478	130	352	461			
Cap. (mL)	185	202	520	128	296	363	472	1240	3209

cup or drinking glass. During the interaction, the contact of the content with the container and other content generates different acoustic phenomena. The CORSMAL Containers Manipulation [9] dataset covers this scenario with 1,140 audio recordings of 12 people manipulating 15 containers (5 cups, 5 drinking glasses, 5 food boxes) filled with pasta, rice, or water (note that water is not used with food boxes).

While each property (content type, content level) could be classified independently [10]–[12], the combination of the two predictions can result in a wrong classification, if either is incorrect. We thus define a set of seven classes that combine content types and levels,  $\mathcal{C} = \{\text{empty}, \text{pasta-half-full}, \text{pasta-full}, \text{rice-half-full}, \text{rice-full}, \text{water-half-full}, \text{water-full}\}$  (see Tab. I). We tackle the task of predicting class  $\tilde{c} \in \mathcal{C}$  with a model consisting of three deep neural network classifiers (see Fig. 2). We named the proposed model Audio Classification of Content, or ACC. As pre-processing, we down-sample the signal to 22,050 Hz and, to eliminate relative gain variations, we normalise it to  $[-1, 1]$ . We perform onset [13] detection to identify the beginning of an action, and transform a temporal window of duration  $L$  from the onset to a spectrogram  $\mathbf{S}$ .

The first classifier,  $f(\cdot)$ , is an action model that estimates the type of container manipulation. We also consider an *unknown* class that covers cases with no content. Thus we define the set  $\mathcal{A} = \{\text{pouring}, \text{shaking}, \text{unknown}\}$  with the classes for the action model,  $f(\cdot)$ , which is trained to classify the action  $\pi \in \{0, 1\}^{|\mathcal{A}|}$ , where  $|\cdot|$  is the cardinality of a set. The action model is a deep neural network that consists of 4 convolutional, 2 max-pooling, and 3 fully connected layers.

The second classifier is selected based on the action:  $g(\cdot)$  for pouring and  $h(\cdot)$  for shaking. Each specialised classifier estimates the content type and level, and is a deep neural network consisting of 6 convolutional, 3 max-pooling, and 3 fully connected layers for the pouring model; and 4 convolutional, 2 max-pooling, and 2 fully connected layers for the shaking model. For the three models, we use ReLU as activation function after each convolutional and fully connected layer, except for the last layer (softmax); max-pooling layers after

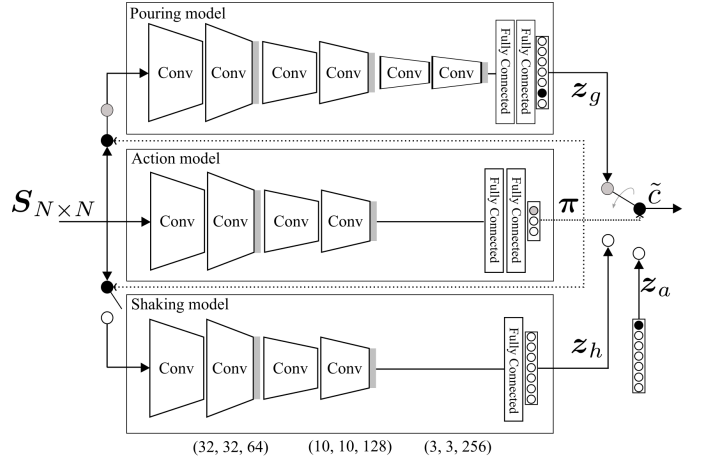


Fig. 2. The proposed model for the classification of the content type and level,  $\tilde{c}$ , of a container manipulated by a person. The model first understands the action  $\pi$  based on an  $N \times N$  spectrogram  $\mathbf{S}$  and then selects the appropriate branch for pouring or shaking. The size of the feature maps after each pooling layer are shown in brackets at the bottom of the diagram. Legend: Conv: convolutional layer,  $\equiv$  max-pooling layer.

every other convolutional layer; and set the kernel size for convolutional and pooling layers to  $3 \times 3$ . The pouring model implements dropout [14] in each convolutional layer before the pooling layer to prevent overfitting.

Each specialised classifier outputs a one-hot vector  $\mathbf{z} \in \{0, 1\}^{|\mathcal{C}|}$ , i.e.  $\mathbf{z}_g = g(\mathbf{S})$  and  $\mathbf{z}_h = h(\mathbf{S})$ . For pouring, the classifier  $g(\cdot)$  predicts the classes  $\mathcal{C} \setminus \{\text{empty}\}$ , whereas the classifier for shaking,  $h(\cdot)$  predicts the classes  $\mathcal{C} \setminus \{\text{empty}, \text{water-half-full}, \text{water-full}\}$ . As the two specialised classifiers cover situations with content only, we define an additional 7-dimensional vector  $\mathbf{z}_a \in \{0, 1\}^{|\mathcal{C}|}$  specifically for the *unknown* cases when classifying the content properties, and that has only the first element set to 1, i.e. *empty* container for the CCM dataset. As the decision of the action classifier conditions the choice of the classifier for the recognition of the content property (or class *unknown*), we compute the final estimation as

$$\tilde{c} = [\pi_1 \mathbf{z}_a + \pi_2 \mathbf{z}_g + \pi_3 \mathbf{z}_h]. \quad (1)$$

### III. VALIDATION

#### A. Methods under comparison

We compare ACC with 9 alternative approaches, namely ResNet (18 layers) [15], a simpler ResNet variant (ResNet\*), a ResNet (18 layers) pre-trained on ImageNet (ResNet-18<sup>T</sup>) [16] and fine-tuned on the training split of CCM, VGG (11 layers) [17], Support Vector Machine (SVM) [18], Random Forest [19], K-Nearest Neighbours (kNN) [20], and the top-2 submissions of the 2020 CORSMAL Challenge<sup>2</sup>, namely Because It's Tactile (BIT) [10], and HVRL [11]. SVM, kNN, Random Forest, VGG, and ResNet-based classifiers perform direct classification as a single model. ResNet\* is a lightweight version of ResNet, consisting of 15 convolutional

<sup>2</sup><https://corsmal.eecs.qmul.ac.uk/challenge.html>

layers grouped into 3 blocks, with 16, 32, and 64 filters, respectively, and a fully connected layer for the 7 classes; and we also removed the first pooling layer.

BIT [10] classifies content level and type independently with a late fusion of audio and RGB data. Audio features provided as input to a Random Forest classifier are Mel-frequency cepstrum coefficients, spectral characteristics, zero crossing rate, chroma vector and deviation; and audio features generated by a modified VGG architecture are provided as input to a GRU-based recurrent neural network. For content-type classification, BIT averages the estimation of the two classifiers. For content-level classification, BIT retrains the two classifiers and includes an additional classifier that uses spatio-temporal features extracted from the RGB frames of the four camera views in CCM. BIT then fuses the decision of the three classifiers.

HVRL [11] uses a convolutional neural network based on VGG [17] to classify the content type from multiple sequential audio frames converted to spectrogram, followed by majority voting. For each audio frame, the feature map prior to the fully connected layers is processed by an LSTM-based recurrent neural network, followed by a fully connected layer, to temporally estimate the content level. The decision of the last frame is the final classification result.

#### B. Datasets and performance measure

The 1,140 audio recordings of the CCM dataset [9] are split into three sets that are evenly distributed among the container types: training (684), public testing (228) and private testing (228). We randomly split the recordings of the training set into training and validation with a 80/20 ratio.

The audio of the CCM dataset was recorded with a circular array of 8 Boya BY-M1 omnidirectional Lavalier microphones and sampling rate 44.1 kHz. The array was placed on the table and the manipulation occurred at a distance between 0.5 to 3 m from the microphones. The multi-channel sound signal is converted into mono channel by averaging samples across channels [13].

We also consider 21 additional recordings acquired in a different room with a Blue Yeti Studio microphone and a sampling rate of 16 bit/48 kHz. We use 4 different containers (two drinking glasses, one cup, one box) and perform a total of 19 pouring and 2 shaking actions with a different type of pasta and rice but the same filling levels as CCM. All recordings were acquired with the microphone placed on a table at a 12 cm from the container for the pouring action and about 20 cm for the shaking action.

To handle class-unbalanced data, we use the weighted average  $F_1$  score,  $\bar{F}_1$ , as performance measure across the  $|C|$  classes, each weighted by the number of audio recordings,  $R_n$ , in each class  $n$ :

$$\bar{F}_1 = \frac{1}{R} \sum_{n=1}^{|C|} R_n F_n, \quad (2)$$

where  $R = \sum_{n=1}^{|C|} R_n$  is the total number of recordings and  $F_n$  is the  $F_1$  score of class  $n$ , i.e. the harmonic mean of precision

TABLE II  
COMPARISON OF THE WEIGHTED AVERAGE  $F_1$  SCORE ( $\bar{F}_1$ ) OF THE COMBINED CONTENT TYPE AND LEVEL CLASSES ON THE TEST SPLITS OF CCM [9] AND COMPLEXITY OF THE MODELS. NOTE THAT BIT AND HVRL ESTIMATE CONTENT TYPE AND LEVEL AS TWO INDEPENDENT PROPERTIES. BEST RESULTS ARE IN BOLD, SECOND BEST IN ITALIC. KEY – # PARAMS.: NUMBER OF PARAMETERS.

	Model	# Params. × 1000	Storage MB	CCM test sets	
				Public	Private
	Random	–	–	9.57	11.22
[10]	BIT	–	–	75.00	77.86
[11]	HVRL	6,839	82.1	<b>82.14</b>	73.40
[20]	kNN	–	63.5	70.33	62.99
[18]	SVM	–	34.8	69.81	73.43
[19]	Random Forest	–	1.7	73.59	70.48
[17]	VGG-11	44,907	175.4	74.03	72.67
[15]	ResNet*	179	0.8	70.45	71.56
[15]	ResNet-18	11,692	45.8	62.28	55.28
[15]	ResNet-18 <sup>T</sup>	11,692	45.8	59.69	63.01
	<b>ACC</b>	16,482	64.3	76.02	<b>78.24</b>

(the number of true positives over the total number of true positives and false positives) and recall (the number of true positives over the total number of true positives and false negatives) for that class. We report  $\bar{F}_1$  as a percentage for the discussion of the results.

#### C. Implementation details

We train each classifier independently for 100 epochs with categorical cross-entropy loss and select the model with the minimum validation loss. During training, we use the Adam [21] optimiser with a learning rate of 0.001 and a batch size of 16. Since the classifiers need spectrograms of  $N \times N$  resolution as input, we resize the spectrograms to  $N = 96$ . For the temporal segmentation, we also set  $L = 10$  s based on the average duration of the actions. Note that we complement sound signals with duration shorter than  $L$  with zero padding, as other strategies, such as repeating or interpolation, affect the prediction of the content level and worsen the accuracy in validation.

All the deep neural networks that perform direct classification are re-implemented or provided in TensorFlow. These models receive the same input as ACC. SVM, kNN, and Random Forest are provided in *scikit-learn* and receive a spectrogram reshaped into a 1D vector.

#### D. Discussion

Tab. II compares the weighted average  $F_1$  scores of the models under analysis on the CCM test splits, and their complexity in terms of number of parameters and storage size. ACC achieves  $\bar{F}_1$  of 76.02 and 78.24 for the two test splits, respectively. ACC outperforms all the models that directly classify on both test splits. HVRL is the best performing on the public test split, but its accuracy decrease on the private test split shows that independently estimating the two content properties does not always guarantee consistency of the two predictions. BIT classifies content types and levels independently, and shows a similar behaviour of HVRL but

TABLE III

COMPARISON OF THE WEIGHTED AVERAGE  $F_1$  SCORE ( $\bar{F}_1$ ) ON THE TESTING SET WITH 21 AUDIO RECORDINGS OF NOVEL CONTAINERS USED IN A DIFFERENT ENVIRONMENT AND SETUP. BEST RESULTS ARE IN BOLD, SECOND BEST IN ITALIC.

	Model	Level	Type	Type & Level
	Random	43.98	17.05	15.24
[20]	kNN	37.28	37.06	19.09
[18]	SVM	40.77	47.39	26.98
[19]	Random Forest	47.62	44.49	20.07
[17]	VGG-11	<b>58.53</b>	25.95	15.24
[15]	ResNet*	52.04	<i>56.11</i>	<i>34.95</i>
[15]	ResNet-18	50.40	45.31	30.87
[15]	ResNet-18 <sup>†</sup>	50.30	<b>61.50</b>	32.02
	ACC	58.34	53.85	<b>41.89</b>

achieving higher weighted average  $F_1$  score in the private test split and a lower one in the public test split.

The architecture of ACC is based on that of VGG but with almost 3 times fewer parameters (16,482,637 versus 44,907,000, respectively), even though three independent classifiers are used: the action classifier has 7,077,315 parameters, the classifier specific for pouring has 2,590,150 parameters, and the classifier specific for shaking has 6,815,172 parameters. Moreover, ACC has 5,000,000 more parameters than the two ResNet-18 models, whereas HVRL has a convolutional neural network of 4,472,580 and an LSTM of 2,366,211 parameters, respectively, that requires more storage (82.1 MB). However, the modified ResNet has only 179,000 parameters, less than 1 MB in storage, and achieves higher performance than ResNet-18. Note also how the use of the resized spectrogram as input during training makes the storage of the trained kNN and SVM classifiers comparable to the deep neural network classifiers, while Random Forest better copes with the large input and needs only 1.7 MB to store the trained model.

Fig. 3 compares the confusion matrices for the models performing classification of the combined content type and level on the combined CCM test splits. ResNet-18 and ResNet-18<sup>†</sup> wrongly classifies *pasta-half-full* with *pasta-full*, and transfer learning helps to reduce the wrong classification of *water-half-full* with *water-full*. Reducing the number of layers from 18 to 15 helps to increase the correct classification across the different classes, while decreasing the overfitting on this small-size dataset. Overall, the methods tend to wrongly classify the level *half-full* with *full* and the opposite, while the *empty (unknown)* case is correctly classified by all methods. ACC reduces the wrong classifications and achieves a higher number of correct classifications along the diagonal compared to the other models.

Tab. III compares the results obtained from the recordings of the second, additional setup. ACC achieves  $\bar{F}_1 = 58.34$  in correctly predicting only the content level, outperforming the other approaches except VGG-11 ( $\bar{F}_1 = 58.53$ ). For content type, ACC achieves  $\bar{F}_1 = 53.85$  and only ResNet-18 and ResNet\* achieves higher accuracy ( $\bar{F}_1 = 61.50$  and  $56.11$ ,

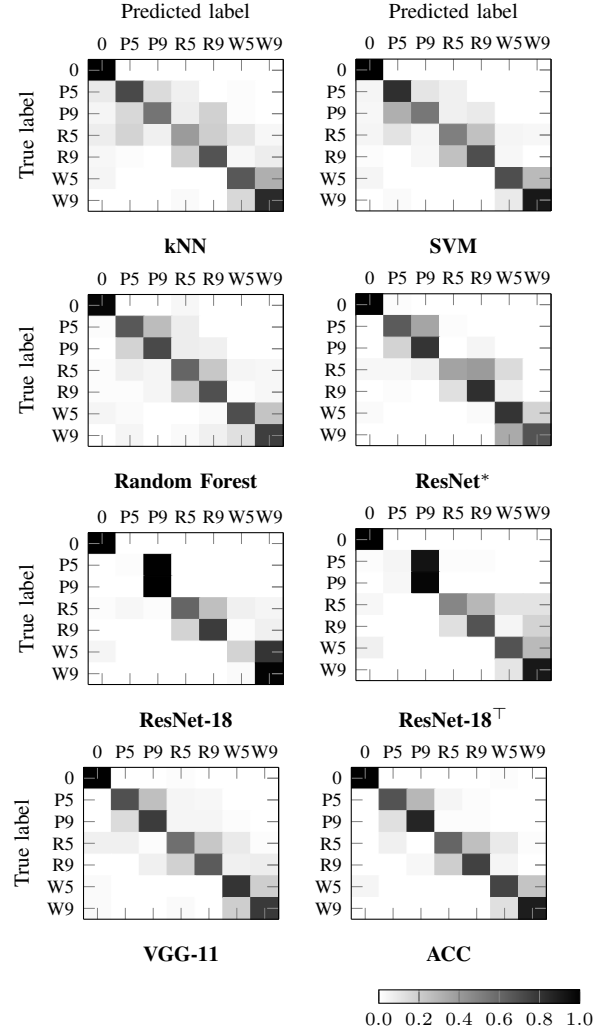


Fig. 3. Confusion matrices for all methods across all the containers of the public and private testing splits of CCM [9]. KEY – 0: empty; CX: content type (C) and level (X), where C represents pasta (P), rice (R), or water (W); and X is half-full (5) or full (9).

respectively). Despite this difference between the accuracy for content type and content level, ACC outperforms other approaches when validating the combination of both content type and level. ACC achieves  $\bar{F}_1 = 41.89$ , while the second best, ResNet\*, achieves  $\bar{F}_1 = 34.95$ .

Fig. 4 shows the confusion matrices of the action classifier for the combined test sets in CCM and the test set acquired in a different environment and setup. The action classifier of ACC can correctly distinguish between the two actions of pouring and shaking, as well as the *unknown* case, for unknown containers in the same setup and environment (CCM). The prediction is instead more challenging in the novel setup and environment, as *unknown* is confused with *pouring* and some recordings with *pouring* are confused with *shaking*. The higher sensitivity and closer position of the Blue Yeti Studio microphone compared to the Boya BY-M1 omnidirectional Lavelier microphone are possible causes of the lower performance in addition to the type of content.

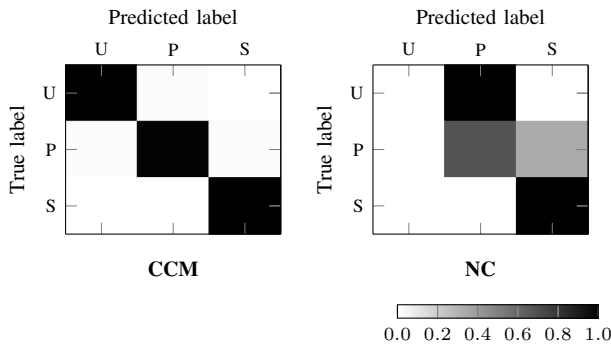


Fig. 4. Confusion matrices for the action classifier of ACC across all the containers of the testing split in CCM (left) and the novel containers (NC) used in the second setup (right). KEY – U: unknown, P: pouring, S: shaking.

#### IV. CONCLUSION

We presented an audio classification model that recognises the content type and level in a container that is manipulated by a person who shakes it or pours content into it. The proposed model works with both liquid and granular materials and classifies the properties of a range of unknown containers with different materials, shapes and sizes. The proposed model outperforms existing solutions for the estimation of content type and level independently, as well as baselines based on SVM, kNN, random forest and deep neural networks that directly classify content type and level. As future work, we will extend the proposed model to generalise to different setups and environments, and to cope with multi-modal data [7].

#### REFERENCES

- [1] R. Sanchez-Matilla, K. Chatzilygeroudis, A. Modas, N. Ferreira Duarte, A. Xompero, P. Frossard, A. Billard, and A. Cavallaro, "Benchmark for human-to-robot handovers of unseen containers with unknown filling," *IEEE Robotics Autom. Lett.*, vol. 5, no. 2, Apr. 2020.
- [2] J. R. Medina, F. Duvallet, M. Karnam, and A. Billard, "A human-inspired controller for fluid human-robot handovers," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, Cancun, Mexico, 15–17 Nov. 2016.
- [3] P. Rosenberger, A. Cosgun, R. Newbury, J. Kwan, V. Ortenzi, P. Corke, and M. Grafinger, "Object-independent human-to-robot handovers using real time robotic vision," *IEEE Robotics Autom. Lett.*, vol. 6, no. 1, pp. 17–23, Jan. 2021.
- [4] S. Griffith, V. Sukhoy, T. Wegter, and A. Stoytchev, "Object categorization in the sink : Learning behavior – grounded object categories with water," in *Proc. IEEE Int. Conf. Robotics Autom.*, Minneapolis, MN, USA, 14–18 May 2012.
- [5] S. Ikeno, R. Watanabe, R. Okazaki, T. Hachisu, M. Sato, and H. Kajimoto, "Change in the amount poured as a result of vibration when pouring a liquid," in *Int. AsiaHaptics Conf.*, Tsukuba, Japan, 18–20 Nov. 2014.
- [6] S. Clarke, T. Rhodes, C. Atkeson, and O. Kroemer, "Learning audio feedback for estimating amount and flow of granular material," in *Proc. Conf. Robot Learn.*, Zürich, Switzerland, 29–31 Oct. 2018.
- [7] H. Liang, S. Li, X. Ma, N. Hendrich, T. Gerkmann, F. Sun, and J. Zhang, "Making sense of audio vibration for liquid height estimation in robotic pouring," in *Proc. IEEE Int. Conf. Intell. Robot Syst.*, Macau, China, Nov. 2019.
- [8] H. Liang, C. Zhou, S. Li, X. Ma, N. Hendrich, T. Gerkmann, F.-C. Sun, and J. Zhang, "Robust robotic pouring using audition and haptics," in *Proc. IEEE Int. Conf. Intell. Robot Syst.*, Las Vegas, NV, USA, 24 Oct. 2020–24 Jan. 2021.
- [9] A. Xompero, R. Sanchez-Matilla, R. Mazzon, and A. Cavallaro, "CORSMAL Containers Manipulation," 2020, (1.0) [Data set]. Queen Mary University of London. <https://doi.org/10.17636/101CORSMAL1>.

- [10] V. Iashin, F. Palermo, G. Solak, and C. Coppola, "Top-1 CORSMAL Challenge 2020 submission: Filling mass estimation using multi-modal observations of human-robot handovers," in *Proc. IEEE Conf. Pattern Recognit.*, Online, 10–15 Jan. 2021.
- [11] R. Ishikawa, Y. Nagao, R. Hachiuma, and H. Saito, "Audio-visual hybrid approach for filling mass estimation," in *Proc. IEEE Conf. Pattern Recognit.*, Online, 10–15 Jan. 2021.
- [12] Q. Liu, F. Feng, C. Lan, and R. H. M. Chan, "Va2mass: Towards the fluid filling mass estimation via integration of vision & audio learning," in *Proc. IEEE Conf. Pattern Recognit.*, Online, 10–15 Jan. 2021.
- [13] B. McFee, C. Raffel, D. Liang, D. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and music signal analysis in python," in *Proc. Python in Science Conf.*, 2015.
- [14] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, pp. 1929–1958, 2014.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Las Vegas, NV, USA, 27–30 June 2016.
- [16] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. Berg, and L. Fei-Fei, "Imagenet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, pp. 211–252, 2015.
- [17] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. on Learning Represent.*, Banff, Canada, 14–16 Apr. 2014.
- [18] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, pp. 273–297, 1995.
- [19] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5–32, 2004.
- [20] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Trans. Inf. Theory*, vol. 13, pp. 21–27, 1967.
- [21] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. on Learning Represent.*, San Diego, CA, USA, 7–9 May 2015.