

The CORSMAL Challenge

Audio-visual object classification
for human-robot collaboration

Performance scores

Performance scores

The CORSMAL challenge evaluates and ranks the teams by assigning a 100 point-based score that accounts for a set of objective performance scores and an assessment of the submitted source code for reproducibility. The integrity of the submitted results is ensured by having a public test set with no annotations available to the teams and a private test set with both the data and the annotations not available to the teams. For the public test set, teams run their models on their own and submit their results in a constrained amount of time. For the private test set, organisers will install, run and evaluate the teams' models.

To provide a sufficient granularity into the behaviour of the various components of the pipeline, we use 10 performance scores for the challenge tasks across public and private sets. The first 7 scores quantify the accuracy of the estimations for the 5 main tasks. The last 3 scores are an indirect evaluation of the impact of the estimations on the quality of human-to-robot handover and delivery of the container by the robot. Performance scores will be also computed individually for the public CCM test set, the private CCM test set, and their combination. The scores cover filling level, filling type, container capacity, container width at the top, width at the bottom, and height, container mass, filling mass, object mass (container + filling), and the delivery of the container upright and at a pre-defined target location (see Tab. 1).

T1 and T2

For the tasks of filling level and filling type classification, we compute precision, recall, and F1-score for each class k across all the configurations belonging to class k , J_k . *Precision* is the number of true positives over the total number of true positives and false positives for each class k (P_k). *Recall* is the number of true positives over the total number of true positives and false negatives for each class k (R_k). The *F1-score* is the harmonic mean of precision and recall, defined as

$$F_k = 2 \frac{P_k R_k}{P_k + R_k}. \quad (1)$$

We then compute the weighted average F1-score, \bar{F}_1 , across the K classes,

$$\bar{F}_1 = \sum_{k=1}^K \frac{J_k F_k}{J}, \quad (2)$$

where J is the total number of configurations (e.g., for either the public test set, the private test set, or their combination) and J_k is the number of configurations associated to class k . Note that $K = 3$ for filling level classification and $K = 4$ for filling type classification.

T3, T4 and T5

For the tasks of container capacity estimation and container mass estimation, we compute the relative absolute error between the estimated measure $a \in \{\tilde{\gamma}^j, \tilde{m}_c^j\}$, and the true measure, $b \in \{\gamma^j, m_c^j\}$, for each configuration j ,

$$\varepsilon(a, b) = \frac{|a - b|}{b}. \quad (3)$$

For the task of container dimensions estimation, where $a \in \{\tilde{w}_t^j, \tilde{w}_b^j, \tilde{h}^j\}$ and b is the corresponding annotation, we use the normalisation function $\sigma_1(\cdot, \cdot)$ [2]:

$$\sigma_1(a, b) = \begin{cases} 1 & \text{if } b = a = 0 \\ 1 - \frac{|a-b|}{b} & \text{if } |a-b| < b \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Note that the first condition is never used in our case, as $b > 0$.

For the indirect high-level filling mass estimation¹, we compute the relative absolute error between the estimated,

¹Note that the score for filling mass is not a linear combination of the scores outputted for filling level classification, filling type classification, and capacity estimation, but it takes into consideration the formula for computing the filling mass based on the estimations of each task for each configuration. This means that a method with lower T1, T2 and T3 scores can obtain a higher filling mass score compared to other methods because the performance on each run is more accurate in general.

\tilde{m}_f^j , and the true filling mass, m_f^j , for each configuration j , unless the annotated mass is zero (empty filling level),

$$\epsilon(\tilde{m}_f^j, m_f^j) = \begin{cases} 0, & \text{if } m_f^j = 0 \wedge \tilde{m}_f^j = 0, \\ \tilde{m}_f^j, & \text{if } m_f^j = 0 \wedge \tilde{m}_f^j \neq 0, \\ \frac{|\tilde{m}_f^j - m_f^j|}{m_f^j}, & \text{otherwise.} \end{cases} \quad (5)$$

With reference to Tab. 1, we compute the score, $s_i \in [0, 1]$ (where 1 is best), across all the configurations J for each measure as follows:

$$s_i = \begin{cases} \frac{1}{J} \sum_{i=1}^J \mathbb{1}e^{-\epsilon(a,b)} & \text{if } a \in \{\tilde{\gamma}^j, \tilde{m}_c^j\} \\ \frac{1}{J} \sum_{j=1}^J \mathbb{1}\sigma_1(a,b) & \text{if } a \in \{\tilde{w}^j, \tilde{w}_b^j, \tilde{h}^j\} \\ \frac{1}{J} \sum_{i=1}^J \mathbb{1}e^{-\epsilon(a,b)} & \text{if } a = \tilde{m}_f^j, \end{cases} \quad (6)$$

with $i = \{3, \dots, 8\}$. The indicator function $\mathbb{1} \in \{0, 1\}$ is 0 only when $a \in \{\tilde{\gamma}^j, \tilde{m}_c^j, \tilde{w}_t^j, \tilde{w}_b^j, \tilde{h}^j, \tilde{m}_f^j\}$ in recording j is not estimated. Note that estimated and annotated measures are strictly positive, $a > 0$ and $b > 0$, except for filling mass that can be 0 when the container is empty.

Object safety and accuracy of delivery

An innovative feature of the CORSMAL challenge is the use of the estimation results by the teams in the context of a specific use-case. For this reason, we evaluate the object safety and the delivery accuracy using a human-to-robot handover simulator [1].

Object safety is the probability that the predicted force applied by the robot (based on the object mass given by the sum of container and filling masses) will enable a gripper to hold the container without dropping it or breaking it. We compute object safety as an exponential function that accounts for the difference between the predicted normal force \hat{F}_j and the required normal force at the real object mass (as available to the simulator), \tilde{F}_j , for each configuration j :

$$\psi_j = e^{\frac{|\tilde{F}_j - \hat{F}_j|}{\tilde{F}_j} \ln(1-c)}, \quad (7)$$

where c controls the sensitivity of ψ_j . A negative difference represents an increase in the probability of dropping the container and a positive difference represents an increase in the probability of breaking the container. We

T1	T2	T3	T4	T5	Description	Unit	Measure	Score	Weight	Type	R2S
●	○	○	○	○	Filling level		λ^j	$s_1 = \bar{F}_1(\lambda^1, \dots, \lambda^J, \hat{\lambda}^1, \dots, \hat{\lambda}^J)$	$\pi_1 = 1/8$	D	○
○	●	○	○	○	Filling type		τ^j	$s_2 = \bar{F}_1(\tau^1, \dots, \tau^J, \hat{\tau}^1, \dots, \hat{\tau}^J)$	$\pi_2 = 1/8$	D	○
○	○	●	○	○	Capacity	mL	γ^j	$s_3 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j e^{-\epsilon^j(\gamma^j, \hat{\gamma}^j)}$	$\pi_3 = 1/8$	D	○
○	○	○	●	○	Container mass	g	m_c^j	$s_4 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j e^{-\epsilon^j(m_c^j, \hat{m}_c^j)}$	$\pi_4 = 1/8$	D	○
○	○	○	○	●	Width at top	mm	w_t^j	$s_5 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j \sigma_1(w_t^j, \hat{w}_t^j)$	$\pi_5 = 1/24$	D	○
○	○	○	○	●	Width at bottom	mm	w_b^j	$s_6 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j \sigma_1(w_b^j, \hat{w}_b^j)$	$\pi_6 = 1/24$	D	○
○	○	○	○	●	Height	mm	h^j	$s_7 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j \sigma_1(h^j, \hat{h}^j)$	$\pi_7 = 1/24$	D	○
●	●	●	○	○	Filling mass	g	m_f^j	$s_8 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j e^{-\epsilon^j(m_f^j, \hat{m}_f^j)}$	$\pi_8 = 1/8^*$	I	○
●	●	●	●	●	Object mass	g	m^j	$s_9 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}_j \psi^j(m^j, \hat{F}^j)$	$\pi_9 = 1/8^*$	I	●
●	●	●	●	●	Pose at delivery	(mm, °)	(α^j, β^j)	$s_{10} = \frac{1}{J} \sum_{j=1}^J \Delta_j(\alpha^j, \beta^j, \eta, \phi)$	$\pi_{10} = 1/8^*$	I	●
●	●	○	○	○	Joint filling type and level			$s_{11} = \bar{F}_1(\lambda^1, \tau^1, \dots, \hat{\lambda}^1, \hat{\tau}^1, \dots)$	—	D	—
○	○	●	○	●	Container capacity and dimensions			$s_{12} = s_3/2 + (s_4 + s_5 + s_6)/6$	—	D	—
●	●	●	●	●	Overall score			$S = \sum_{l=1}^{10} \pi_l s_l$	—	I	—

KEY – T: task, D: direct score, I: indirect score, R2S: measured in the real-to-simulation framework [1].

* weighted by the number of performed tasks.

Table 1: Performance scores for a given test set (public, private or their combination). For a measure a , its corresponding ground-truth value is \hat{a} . All scores are normalised and presented in percentages (the higher, the better). $\bar{F}_1(\cdot)$ is the weighted average F1-score. Filling amount and type are sets of classes (no unit).



Figure 1: Example of successful and unsuccessful container delivery to the target location using the CORSMAL simulator [1].

approximate the force applied by the robot in simulation for holding the object to

$$\tilde{F} \approx \frac{\tilde{m}(g + a_{max})}{\mu}, \quad (8)$$

where \tilde{m} is the predicted mass of the container (and content, if any); $g = 9.81 \text{ m/s}^{-2}$ is the gravitational acceleration; a_{max} is the acceleration of the robot arm when holding and delivering the object; and μ is the coefficient of friction between the container and the gripper ($\mu = 1.0$ [1]).

We quantify the **accuracy of delivering** the container upright and within the target area as

$$\Delta_j = \begin{cases} 1 - \frac{\alpha}{\eta}, & \text{if } (\alpha < \eta) \text{ and } (\beta < \phi), \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where α is the distance from the centre of the container base to the target location \mathbf{d} ; η is the maximum distance allowed from the pre-defined delivery location; β is the angle between the vertical axis of the container and the vertical axis of the world coordinate system; and ϕ is the value of β at which the container would tip over. Fig. 1 shows examples of container deliveries.

With reference to Tab. 1, we compute the score, $s_9 \in [0, 1]$ (where 1 is best), across all the configurations J for object safety as

$$s_9 = \frac{1}{J} \sum_{j=1}^J \mathbb{1}\psi^j(m^j, \hat{F}^j), \quad (10)$$

and the score $s_{10} \in [0, 1]$ (where 1 is best), across all the configurations J for the accuracy of delivering the container upright and within the target area as

$$s_{10} = \frac{1}{J} \sum_{j=1}^J \Delta_j(\alpha^j, \beta^j, \eta, \phi). \quad (11)$$

The indicator function $\mathbb{1} \in \{0, 1\}$ is 0 only when either the filling mass or the containers mass is not estimated in recording j .

References

- [1] Y. L. Pang, A. Xompero, C. Oh, and A. Cavallaro, “Towards safe human-to-robot handovers of unknown containers,” in *IEEE International Conference on Robot and Human Interactive Communication*, Virtual, 8–12 Aug. 2021.
- [2] 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.