

The CORSMAL Challenge: Multi-modal Fusion and Learning for Robotics

Call for participation

A major challenge for human-robot cooperation in household chores is enabling robots to predict the properties of previously unseen containers with different fillings. Examples of containers are cups, glasses, mugs, bottles and food boxes, whose varying physical properties such as material, stiffness, texture, transparency, and shape must be inferred on-the-fly prior to a pick-up or a handover. The challenge focuses on the estimation of the capacity and mass of containers, as well as the type, mass, and percentage of the content, if any.

CORSMAL distributes a multi-modal dataset with 1140 visual-audio-inertial recordings of people interacting with containers, for example while pouring a liquid in a glass or moving a food box. The dataset is collected with four multi-sensor devices (one on a robotic arm, one on the human chest and two third-person views) and a circular microphone array. Each device is equipped with an RGB camera, a stereo infrared camera, and an inertial measurement unit. In addition to RGB and infrared images, each device provides synchronized depth images that are spatially aligned with the RGB images. All signals are synchronized, and the calibration information for all devices, as well as the inertial measurements of the body-worn device, is also provided.

Contributions to the ICPR 2020 CORSMAL Challenge are invited on the following topics:

- Machine learning for multi-modal fusion (inertial, audio and video data)
- Multi-modal signal analysis and processing
- Deep learning architectures for estimating the properties of objects
- Multi-modal algorithms for estimating the amount and type of content of a container
- Object, hand and body detection and tracking
- Multi-modal object detection
- Multi-modal and multi-view tracking
- Multi-view object classification
- Machine listening and machine vision for fullness estimation
- Robustness of multi-class classifiers with transparent objects
- Intention prediction

Participants must submit the results of the public testing set, as well as the source code and executable files that will be run by the organizers on the private test set of the provided dataset. The best-performing entries will be presented at the conference venue. Selected participants will be invited to co-author the writing of a paper to discuss and analyze the research outcomes of the challenge.

Schedule

Submission deadline (estimations on the public test set + source codes): 26 October 2020

Release of the results (both public and private testing sets) via the on-online leaderboard: 2 November 2020

Paper submission deadline: 9 November 2020

Acceptance notification: 27 November 2020

Camera-ready due: 4 December 2020

Challenge presentation at ICPR: 10-15 January 2021

Challenge webpage

<http://corsmal.eecs.qmul.ac.uk/ICPR2020challenge.html>