

Tech United Eindhoven @Home

2024 Team Description Paper

A. Aggarwal, M.F.B. van der Burgh, P. van Dooren, L. Messing,
R. Martin Núñez, S. Narla, G. Cordova Cardenas, M.L.M. van Vroonhoven,
and M.J.G. van de Molengraft

Eindhoven University of Technology,
Den Dolech 2, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
<http://www.techunited.nl>, techunited@tue.nl,
<https://github.com/tue-robotics>

Abstract. This paper provides an overview of the main developments of the Tech United Eindhoven RoboCup@Home team. Tech United uses an advanced world modeling system called the Environment Descriptor. It allows straightforward implementation of localization, navigation, exploration, object detection & recognition, object manipulation and robot-robot cooperation skills based on the most recent state of the world. Other important features include object and people recognition via deep learning methods, a GUI & use of HSR's display, natural language interpretation & context-free grammar based speech recognition and a chat interface combined with a custom keyboard and conversation engine. Ongoing developments are focused on improving the perception pipelines.

1 Introduction

Tech United Eindhoven¹ (established 2005) is the RoboCup student team of Eindhoven University of Technology² (TU/e), which joined the ambitious @Home League in 2011. The RoboCup@Home competition aims to develop service robots that can perform everyday tasks in dynamic and cluttered 'home' environments. The team has been awarded multiple world vice-champion titles in the Open Platform League (OPL) of the RoboCup@Home competition during previous years, and world champion titles in 2019 and 2022 in the Domestic Standard Platform League (DSPL).

Tech United Eindhoven consists of (former) PhD and MSc. students and staff members from different departments within the TU/e. The software base is developed to be robot independent, which means that the years of development on AMIGO and SERGIO are currently being used by HERO. Thus, a large part of the developments discussed in this paper have been optimized for years, whilst the DSPL competition has only existed since 2017³. All the software discussed

¹ <http://www.techunited.nl>

² <http://www.tue.nl>

³ <https://athome.robocup.org/robocuphome-spl>

in this paper is available open-source at GitHub⁴, as well as various tutorials to assist with implementation.

2 Environment Descriptor (ED)

The TU/e Environment Descriptor (ED) is a Robot Operating System (ROS) based 3D geometric, object-based world representation system for robots. ED is a database system that structures multi-modal sensor information and represents this such that it can be utilized for robot localisation, navigation, manipulation and interaction. Figure 1 shows a schematic overview of ED.

ED has been used on our robots in the OPL since 2012 and was also used this year in the DSPL. Previous developments have focused on making ED platform independent, as a result ED has been used on the PR2, Turtlebot, Dr. Robot systems (X80), as well as on multiple other @Home robots.

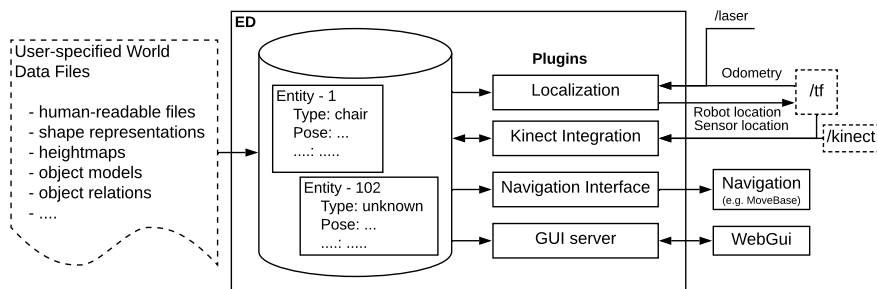


Fig. 1. Schematic overview of TU/e Environment Descriptor. Double sided arrows indicate that the information is shared both ways, one sided arrows indicate that the information is only shared in one direction.

ED is a single re-usable environment description that can be used for a multitude of desired functionalities such as object detection, navigation and human machine interaction. Improvements in ED reflect in the performances of the separate robot skills, as these skills are closely integrated in ED. This single world model allows for all data to be current and accurate without requiring updating and synchronization of multiple world models. Currently, different ED plug-ins exist that enable robots to localize themselves, update positions of known objects based on recent sensor data, segment and store newly encountered objects and visualize all this in RViz and through a web-based GUI, as illustrated in Figure 6. ED allows for all the different subsystems that are required to perform challenges to work together robustly. These various subsystems are shown in Figure 2, and are individually elaborated upon in this paper.

⁴ <https://github.com/tue-robotics>

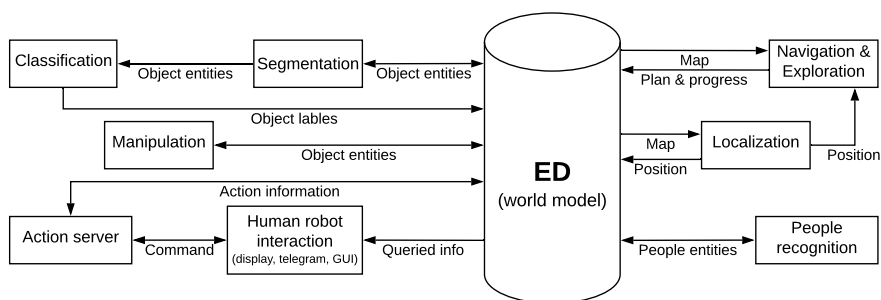


Fig. 2. A view of the data interaction with robot skills that ED is responsible for.

2.1 Localization, Navigation and Exploration

The *ed_localization*⁵ plugin implements AMCL based on a 2D render of the central world model. With use of the *ed_navigation* plugin⁶, an occupancy grid is derived from the world model and published. With the use of the *cb_base_navigation* package⁷ the robots are able to deal with end goal constraints. The *ed_navigation* plugin allows to construct such a constraint w.r.t. a world model entity in ED. This enables the robot to navigate not only to areas or entities in the scene, but to waypoints as well. Figure 3 shows the navigation to an area. Modified versions of the local and global ROS planners available within *move_base* are used.




Fig. 3. A view of the world model created with ED. The figure shows the occupancy grid as well as classified objects recognized on top of the cabinet.

2.2 Detection & Segmentation

ED enables integrating sensors through the use of the plugins present in the *ed_sensor_integration* package. Two different plugins exist:

⁵ https://github.com/tue-robotics/ed_localization

⁶ https://github.com/tue-robotics/ed_navigation

⁷ https://github.com/tue-robotics/cb_base_navigation

1. *laser_plugin*: Enables tracking of 2D laser clusters. This plugin can be used to track dynamic obstacles such as humans.
2. *kinect_plugin*: Enables world model updates with use of data from a RGBD camera. This plugin exposes several ROS services that realize different functionalities:
 - (a) *Segment*: A service that segments sensor data that is not associated with other world model entities. Segmentation areas can be specified per entity in the scene. This allows to segment object ‘on-top-of’ or ‘in’ a cabinet. All points outside the segmented area are ignore for segmentation.
 - (b) *FitModel*: A service that fits the specified model in the sensor data of a RGBD camera. This allows updating semi-static obstacles such as tables and chairs.

The *ed_sensor_integration* plugins enable updating and creating entities. However, new entities are classified as unknown entities. Classification is done in *ed_perception* plugin⁸ package.

2.3 Object grasping, moving and placing

The system architecture developed for object manipulation is focused on grasping. In the implementation, its input is a specific target entity in ED, selected by a Python executive and the output is the grasp motion joint trajectory. Figure 4 shows the grasping pipeline.




Fig. 4. Custom grasping pipeline base on ED, MoveIt and a separate grasp point determination and approach vector node.

MoveIt! is used to produce joint trajectories over time, given the current configuration, robot model, ED world model (for collision avoidance) and the final configuration. The grasp pose determination uses the information about the position and shape of the object in ED to determine the best grasping pose. The grasping pose is a vector relative to the robot. Placing an object is approached in a similar manner to grasping, except for that when placing an object, ED is queried to find an empty placement pose.

⁸ https://github.com/tue-robotics/ed_perception

3 Image Recognition

The *image_recognition* packages apply state of the art image classification techniques based on Convolutional Neural Networks (CNN).

1. **Object recognition:** Tensorflow™ with retrained top-layer of a Inception V3 neural network.
2. **Face recognition:** OpenFace⁹, based on Torch.
3. **Pose detection:** OpenPose¹⁰.

Our image recognition ROS packages are available on GitHub¹¹ and as Debian packages: *ros-noetic-image-recognition*. Our ongoing efforts are to upgrade these to a YOLOv8 based pipeline for object recognition and pose detection, and FaceNet for face recognition, all implement in PyTorch 2.1 with TensorRT.

4 People Recognition

As our robots need to operate and interact with people in a dynamic environment, our robots' need people detection skills. For this a generalized system capable of recognizing people in 3D is developed. An RGB-D camera is used to capture the scene information after which a recognition sequence is completed in four steps. First, people are detected in the scene using OpenPose and if their faces are recognized as one of the learned faces in the robot's database, they are labeled using their known name using OpenFace. The detections from OpenPose are associated with the recognitions from OpenFace by maximizing the IoUs of the face ROIs. Then, for each of the recognized people, additional properties such as age, gender and the shirt color are identified. Furthermore, the pose keypoints of these recognitions are coupled with the depth information of the scene to re-project the recognized people to 3D as skeletons. Finally, information about the posture of each 3D skeleton is calculated using geometrical heuristics. This allows for the addition of properties such as "pointing pose" and additional flags such as 'is_waving', 'is_sitting', etc.

4.1 Pointing detection

In the previous section, our approach to people recognition is explained. This recognition includes information about the posture of each 3D skeleton. Once the people information is inserted into the world model, additional properties can be added to the persons that take also other entities in the world model into account, e.g. 'is_pointing_at_entity'. This information is used by the toplevel state machines to implement challenges. However an additional check based on spatial queries is inserted to ensure that the correct operator is found. By using such

⁹ <https://cmusatyalab.github.io/openface/>

¹⁰ <https://github.com/CMU-Perceptual-Computing-Lab/openpose>

¹¹ https://github.com/tue-robotics/image_recognition

a query it is possible to filter out people based on their location. Finally, to determine at which entity the operator is pointing, ray-tracing is implemented. Figure 5 shows an example of the ray-tracing.




Fig. 5. Ray-tracing based on pose detection with AMIGO.

5 Human-Robot Interface

We provide multiple ways of interacting with the robot in an intuitive manner: WebGUI (Section 5.1), Speech Recognition (Section 5.2), Telegram™ interface (Section 5.3), and the Head Display (Section 5.4).

5.1 Web GUI

In order to interact with the robot, apart from speech, we have designed a web-based Graphical User Interface (GUI). This interface uses HTML5¹² with the Robot API written in Javascript and we host it on the robot itself.



Fig. 6. Illustration of the 3D scene of the WebGUI with AMIGO. User can long-press objects to open a menu from which actions on the object can be triggered

Figure 6 represents an instance of the various interactions that are possible with the Robot API.

¹² https://github.com/tue-robotics/tue_mobile_ui

5.2 Speech Recognition

For natural interaction using spoken commands, we use Picovoice¹³ as our Automatic Speech Recognition (ASR) engine. We prepare various (Feature) Context-free Grammar (FCFG) based language models for the ASR model, some specific to the challenges and the rest general and have developed an interface that allows lazy loading of these models in the challenges.

5.3 Telegram™

The Telegram interface¹⁴ to our robots is a ROS wrapper around the *python-telegram-bot* library. The software exposes four topics, for images and text resp. from and to the robot. The interface allows only one master of the robot at a time. All the reasoning behind this interface is handled by the *conversation_engine*.

Conversation Engine

The *conversation_engine*¹⁵ bridges the gap between text input and an action planner (called *action_server*). The input text is parsed according to a FCFG, resulting in an action description in the form of a nested mapping. In the action description, (sub)actions and their parameters are filled in. This may include references such as “it”. Based on the action description, the *action_server* tries to devise a sequence of actions and parameterize those with concrete object IDs. Furthermore, the *conversation_engine* engages with the user for additional input to fill in missing information, and keeps the user “informed” whilst actions are being performed by reporting on the current subtask.

Custom Keyboard, Telegram HMI ¹⁶

To reduce the operator error, we only present the user with a limited number of buttons in the Telegram app using its *custom_keyboards* feature. This feature has been employed to compose commands word-for-word. After the user has entered a partial command, for example “Bring me the ...” they are presented with only those words that might follow that text according to the grammar, eg. “apple”, “orange” etc. This process iterates until a full command has been composed.

5.4 Head Display

For most people, especially people who do not deal with robots in their day-to-day life, interaction with robots is not as easy as one would like it to be. To remedy this, the head display (pre-integrated in the Toyota HSRs) of HERO is used¹⁷. By default, our Tech United @Home logo is shown on the screen, as

¹³ https://github.com/tue-robotics/picovoice_ros

¹⁴ https://github.com/tue-robotics/telegram_ros

¹⁵ https://github.com/tue-robotics/conversation_engine

¹⁶ https://github.com/tue-robotics/hmi_telegram

¹⁷ <https://github.com/tue-robotics/hero-display>

depicted in Figure 7. When the robot is speaking the spoken text is displayed, when the robot is listening a spinner along with an image of a microphone is shown. It is also possible to display (live) images for various tasks.



Fig. 7. The default status of HERO's head display.

6 Re-usability of the system for other research groups

Tech United takes great pride in creating and maintaining open-source software and hardware to accelerate innovation. Tech United initiated the Robotic Open Platform website¹⁸, to share hardware designs. All our software is available on GitHub¹⁹. All packages include documentation and tutorials. Tech United and its scientific staff have the capacity to co-develop (15+ people), maintain and assist in resolving questions.

References

1. Morgan Quigley, Ken Conley, Brian P. Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. ROS: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.
2. D. Fox. Adapting the sample size in particle filters through kld-sampling. *The International Journal of Robotics Research*, 22(12):985–1003, 2003.
3. D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Magazine on Robotics & Automation*, 4(1):23–33, 1997.
4. Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015.
5. Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. Technical report, CMU-CS-16-118, CMU School of Computer Science, 2016.

¹⁸ <http://www.roboticopenplatform.org>

¹⁹ <https://github.com/tue-robotics>

A HSR's Software and External Devices

We use a standard Toyota™ HSR robot. To differentiate our unit, we named it HERO. We wanted to link it's name to our AMIGO and SERGIO domestic service robots.

An overview of the software used by the Tech United Eindhoven @Home robots can be found in Table 1. All our software is developed open-source at GitHub²⁰.

Table 1. Software overview

Operating system	Ubuntu 20.04 LTS Server
Middleware	ROS Noetic [1]
Simulation	Gazebo
World model*	Environment Descriptor (ED), custom https://github.com/tue-robotics/ed
Localization*	Monte Carlo [2] using Environment Descriptor (ED), custom https://github.com/tue-robotics/ed_localization
SLAM	Gmapping
Navigation*	CB Base navigation, custom https://github.com/tue-robotics/cb_base_navigation Global: custom A* planner Local: modified ROS DWA [3]
Arm navigation	MoveIt!
Object recognition*	Inception based custom DNN [4], custom https://github.com/tue-robotics/image_recognition
People recognition*	Custom implementation using contour matching https://github.com/tue-robotics/people_recognition
Face recognition*	OpenFace [5], custom https://github.com/tue-robotics/image_recognition
Speech recognition*	Picovoice, custom https://github.com/tue-robotics/picovoice_ros.git
Speech synthesis	Toyota™ Text-to-Speech
Task executors*	SMACH, custom https://github.com/tue-robotics/tue_robocup

External Devices *HERO* relies on the following external hardware:

- Official Standard Laptop
- USB power speaker
- Gigabit Ethernet Switch
- Wi-Fi adapter

Cloud Services *HERO* does not use any cloud services

²⁰ <https://github.com/tue-robotics>