

Dexter Friis-Hecht
Dokyun Kim
Dominic Salmieri
Maurice Ampane
12/13/2023

VORA Ethics Statement

VORA (Voice-Operated Robot Assistant) is designed to, as the name suggests, assist its users. By integrating a machine-learning model and computer vision algorithms into the Neato, it is able to recognize the speaker and act accordingly when given a voice command. The intended use for VORA was to serve as an assistant that retrieves the user's desired object from a "loading bay."

The two main technologies used in this project are:

1. Custom-trained voice classifier using **machine learning**
2. Object detection/tracking using **computer vision**

In the section below, we will cover the possible ethical concerns that these technologies might raise.

Regarding the voice-classifier:

1. Not only does VORA listen to what the user is saying, it determines the identity of the user. There are many helpful applications of this, like live transcription of conversations for people who are hard of hearing. Unfortunately, it could be also used for surveillance or other monitoring applications that can negatively impact people without them even knowing. However, our classifier needs a significant amount of training data, and that would limit the ability of classifying the voices of people without their knowledge.
2. VORA uses a wake word and is therefore constantly listening. This is a definite privacy concern, and could be used to gather data on individuals that could be used for a multitude of purposes, for example personalized advertising. Many people already have suspicions about this with voice-activated assistants like Apple's Siri and Amazon's Alexa. We do not store or share any of the collected data from the transcription or voice classifier, and the only thing that information is used for is controlling the Neato.
3. Currently, VORA classifies one of our group members much better than the others. We are not sure what the root cause of this is, but in collecting our data and training our model we need to ensure that we are using unbiased data that will not cause our model to have biases in classification.

Regarding the object detector/tracker:

1. Currently, VORA utilizes Ultralytic's pre-trained YOLOv8 model, which has been trained on the COCO dataset. By default, the model is configured to detect and track specific objects such as cups, bottles, cell phones, computer mice, books, and scissors. It is technically possible to modify the program to include human detection since the COCO dataset already contains images of people. However, the likelihood of the system being used to facilitate harm is extremely low. This is due to the Neato having a maximum speed of 0.2 m/s, making it easily avoidable.
2. As mentioned earlier, VORA has the technical capability to detect and track humans if modified. This may raise concerns related to human rights. While it is true that certain vision technologies can be modified in ways that violate individual privacy, VORA, in its current form, simply follows one of the detected objects. The likelihood of VORA being used as a method of mass surveillance is extremely low. We believe that its surveillance capability is lower than that of a security camera, given that it only tracks a single object at a time.