Grassland African Road Images (GARI): A Driving Dataset from Kenyan Highways and National Parks

Khushal Brahmbhatt
AI Researcher
Kenya
khushalbr@gmail.com

Ronald Ojino
Lecturer
The Cooperative University of Kenya
Kenya
ronojinx@gmail.com

I. Introduction

The current state of autonomous driving technology is disproportionately developed around the ecosystems of the Global North and poorly adapts to the needs of the Global South. A contributing factor to this is the lack of driving datasets from developing countries to fuel research in this direction. Nearly all driving datasets available have been collected on well-defined roads and environments from developed countries. The models developed using these datasets do not adapt well to road conditions in developing countries.

It is with this motivation that we set out to collect a driving dataset from Kenya in unstructured environments lacking good road infrastructure. The Grassland African Road Images (GARI) dataset consists of images extracted from driving video recorded on highways and suburbs in western Kenya, as well as driving trails across two national parks in Kenya. The dataset also consists of sensor measurements such as speed and steering angle from the vehicle showing the driving behaviour. From the driving data collected in national parks, we also sought to address a different problem in wildlife conservation. We aimed to investigate the technological feasibility of deploying unmanned ground vehicles for automated wildlife patrol in national parks in order to meet shortages in ranger workforce and increase patrol efficiency[1]. The GARI dataset was collected with the intention of predicting steering angles as well as other parameters such as brake and accelerator pedal positions using behavioural cloning. To the best of our knowledge, this is the first driving dataset from Africa.

II. GARI Dataset

A. Data Collection Platform

The data was collected between two different vehicles: A Toyota Prius 2012, and a Toyota Aqua (Prius C) 2012. Both vehicle models are nearly identical aside from some performance differences. The video was recorded using an Apeman A80 cuboid action camera center-mounted on the dashboard. The camera was connected to a laptop running Ubuntu Linux 18.04 with an 8-core Intel i5-8300H 2.3GHz processor, and 8GB RAM. The video was recorded with a 720p resolution at a frame rate of 30 fps. The individual video frames were captured and timestamped using OpenCV 4.2 and encoded to a video file on the computer using XVID codec and .avi container.

The driving signals were captured from the vehicle's Engine Control Unit (ECU) through the Controller Area Network (CAN) bus that is usually exposed on the On-board Diagnostics-II (OBD-II) port on most vehicles. The driving signals were logged and timestamped on a Raspberry Pi 3B+ that was interfaced to the vehicle's CAN bus through the PiCAN2 CAN-bus board. GPS information was recorded on a smartphone using the BasicAirData GPS Logger app.

B. Data Descriptions

The data was collected between 20th January 2020 and 17th March 2020. It was collected during different times of day to get an even representation of different lighting conditions in the dataset. The GARI dataset is split into 2 subdatasets:

GARI-Cities: This subdataset consists of approximately 9 hours of driving data recorded over 425 km of tarmacked highways in western Kenya and dirtroads in the suburban areas of Kisumu city.

GARI-Savannah: This subdataset consists of approximately 11 hours of driving data recorded over 145 km of driving trails across two national parks in Kenya. 8.5 hours of data was recorded in Nairobi National Park over 115 km of dirtroad, and 2.5 hours of data was recorded in Ruma National Park over 30 km of dirtroad. It is worth noting that the Nairobi National Park section of the dataset was recorded on a Toyota Aqua 2012, while the rest of the dataset was recorded on

[1] https://github.com/KhushalB/Autonomous-Driving-Wildlife-Park

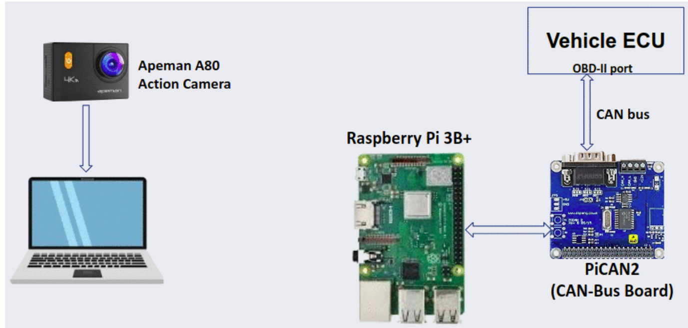

Fig. 1: Data collection setup

a Toyota Prius 2012.

The attributes in the dataset are described below:

i) *Image data*
The images in the dataset are extracted from the recorded videos at a frame rate of 15 fps and have a resolution of 1280×720. Segments of the video containing bad data such as U-turns, intersections, or vehicle stopped at an obstacle were removed before extracting the frames. Segments containing driving in bad road conditions that required extreme steering corrections or going around the obstacles (except for the suburban portion of the GARI-Cities subdataset) were also removed as they presented navigation challenges beyond the scope of our project.

ii) *Steering angle*
The steering angle values given in the dataset are the raw values recorded from the CAN bus. They are not the physical angle of the steering wheel. If you need the measurements as actual angles, a python script[2] is available to make the conversion based on empirical observations from our measurements of the steering value at different angles. Positive values represent anticlockwise angles while negative values represent clockwise angles.

iii) *Steering torque*
Three steering torque values are given as raw recorded values. They could be from three different torque sensors. We were unable to determine the units for the steering torque.

iv) *Speed*
Speed of the vehicle is given in km/h.

v) *Accelerator pedal*
The raw recorded value of the accelerator pedal position. The values did not change linearly, hence we could not a find a way to give them as percentages.

vi) *Brake pedal*
The raw recorded value of the brake pedal position. The values did not change linearly, hence we could not a find a way to give them as percentages.

vii) *Engine rpm*
The raw recorded value of the Internal Combustion Engine (ICE) rpm. We were unable to verify the accuracy of this attribute. Use with caution.

viii) *Individual tire speeds*
The raw recorded value of the rotational speed of each of the four wheels. We were unable to determine which wheel each of the values corresponded to. We could also not determine the units of the attribute. The individual tire speeds are

[2] https://github.com/KhushalB/Autonomous-Driving-Wildlife-Park/blob/master/src/data_collection/steer_angles.py

not available for the section of the dataset recorded in Nairobi National Park.

ix) GPS data

The latitute and longitude are given in decimal degrees. The altitude is the height above sea level given in metres. The GPS coordinates have an accuracy between 3 and 15 metres. As the GPS data is only updated every 1 second while the other parameters have higher update rates, it may be repeating between data samples.

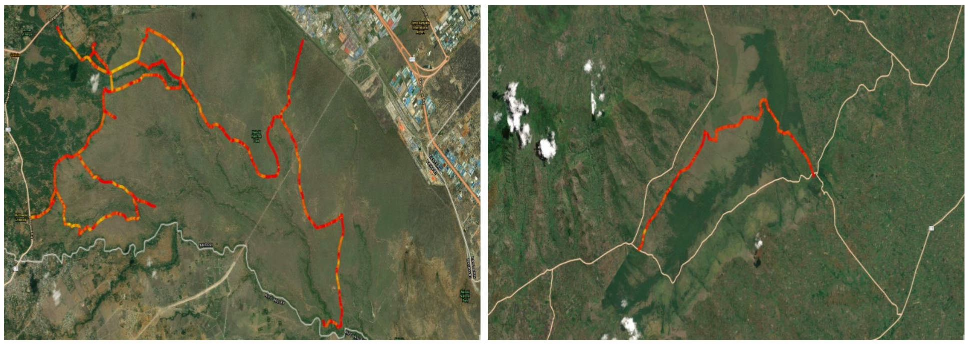

(a) GARI-Savannah (Left: Nairobi National Park, Right: Ruma National Park)

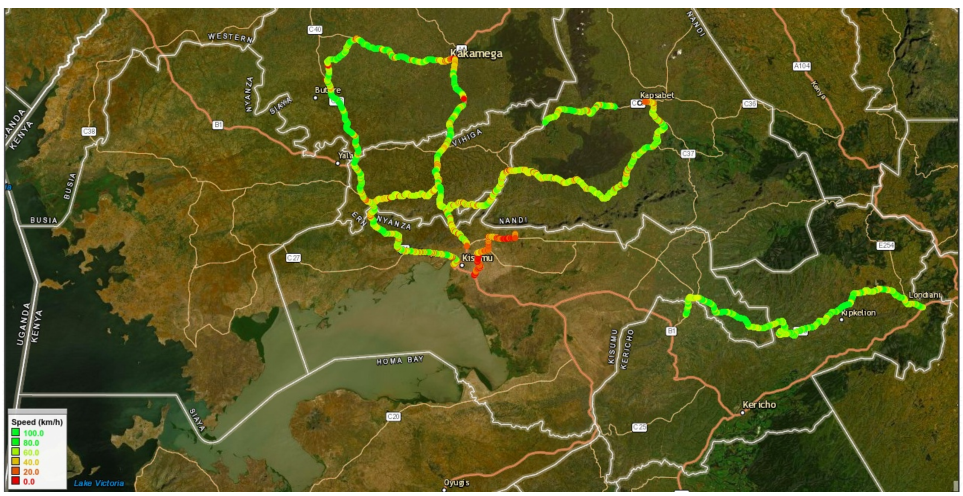

(b) GARI-Cities

Fig. 2: GPS map of areas covered in GARI dataset, colorized by speed.

C. Data Statistics

1) GARI-Cities: The highways section of this subdataset consists of 350,670 images, while the suburban section (*"CITIES-offroad"*) has 68,370 images. The GARI-Cities subdataset is 29 GB in size in total.

Figures 3a and 3b show the count distributions of some of the attributes in the GARI-Cities subdataset. We can see that the steering angles are normally distributed with a mean of around 0. This means that a significantly larger portion of samples have the steering wheel at 0 degrees, and it may be helpful to trim the dataset accordingly when designing your training and validation sets in order to prevent overfitting in your models. The variation of steering angles in the highways section is also smaller (roughly between -50 and 50) than the variation of steering angles in the suburban section (roughly between -200 and 200). This can be explained by the fact that the vehicle was travelling at much higher speeds on the highways with less turns, hence less steering correction was required, while more steering correction was required when travelling over rough dirtroads at low speeds in the suburbs. This relationship between speed and steering angles can also be seen from the contour plots in figures 4a and 4b.

2) GARI-Savannah: The section of this subdataset recorded in Nairobi National Park has 326,115 images while the section recorded in Ruma National Park has 123,911 images. The GARI-Savannah subdataset is 30.1 GB in size in total.

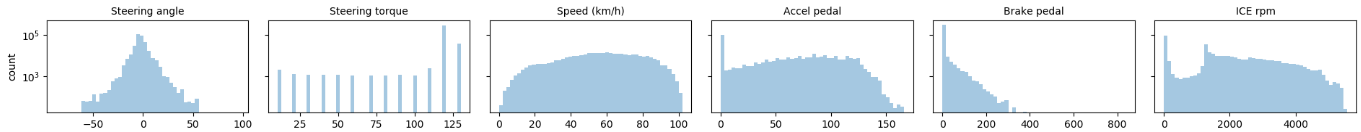

(a) GARI-Cities Highways

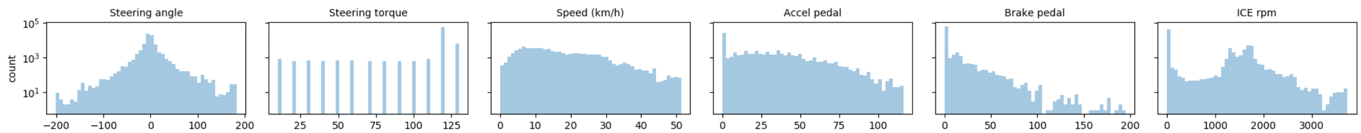

(b) GARI-Cities Suburbs

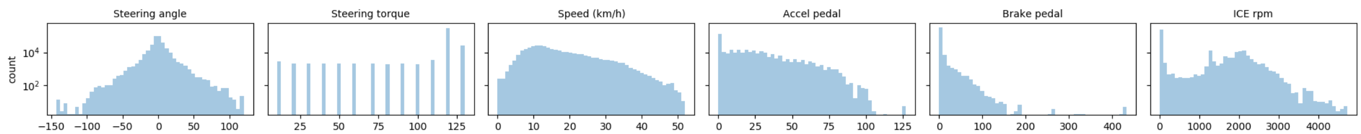

(c) GARI-Savannah Nairobi National Park

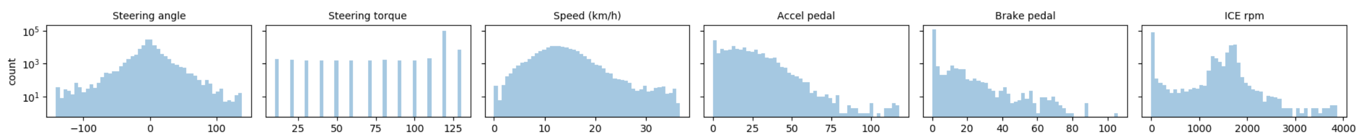

(d) GARI-Savannah Ruma National Park

Fig. 3: Histograms of sensor measurements from the dataset.

Figures 3c and 3d show the count distributions of some of the attributes in the GARI-Savannah subdataset. Both Nairobi National Park and Ruma National Park have approximately similar normal distributions of steering angle values with a mean of around 0. It would be advisable to trim out the large number of samples around 0 degrees and augment the samples with larger steering values to achieve a closer to symmetric distribution when building your training and validation sets in order to prevent overfitting in your models. The speed distributions have rightward skews in both parks, although there is a larger variation in steering values and smaller variation in speed in Ruma National Park as seen from figures 4c and 4d. This can be explained by the fact that the driving trails in Nairobi National Park were much smoother, allowing higher speeds to be reached in some sections, while the rougher roads in Ruma National Park required more cautious driving and more steering corrections.

III. Acknowledgements

This project was made possible through a research grant from the Knowledge 4 All Foundation. It was conducted as part of the Artificial Intelligence for Development (AI4D) initiative in sub-Saharan Africa by the International Development Research Centre (IDRC).

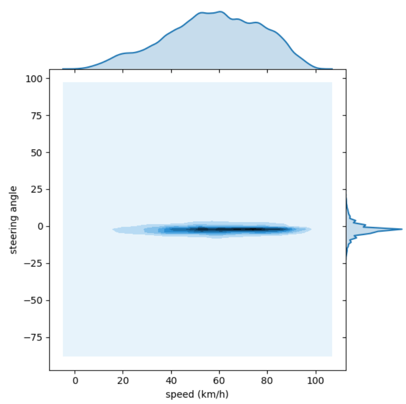

(a) GARI-Cities Highways

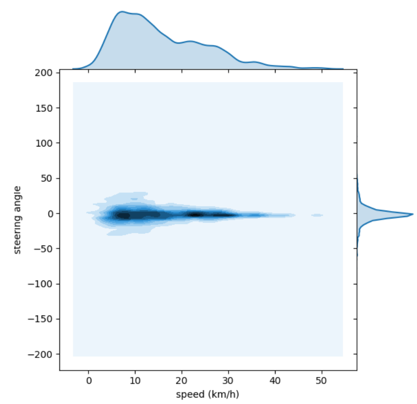

(b) GARI-Cities Suburbs

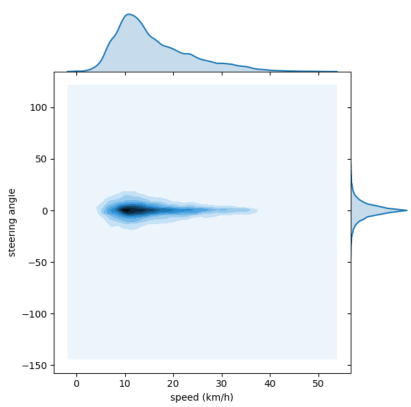

(c) GARI-Savannah Nairobi National Park

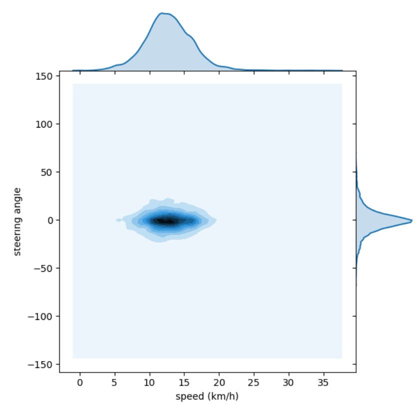

(d) GARI-Savannah Ruma National Park

Fig. 4: KDE contour plots of speed vs steering angle.