CS 6953: Deep Learning Capstone
Assignment 1: Project preference

Ankith Boggaram

January 8, 2025

Here is a list of my preferences

1. **Deep Learning for Estimating Lane Line Quality Using Retroreflectometer Ground Truth**

   Why I want to work on this and my qualifications:

   This problem statement seems the most interesting to me. I think I have the required expertise for this too.

   I have experience working with computer vision models in my previous internship at Magna, where we used a model derived from U-Net to detect defects in industrial equipment. I also have other experience working with computer vision models for detection and classification in many other projects (I don't want to spam this section with my experience, but I'm willing to explain more in person if necessary)

   Apart from the projects and experience I have working with deep learning / computer vision, I've also worked on Tandem algorithms before (RANSAC + Regression), which sounds very similar to this Deep Learning + Regression problem.

   Additionally, I have experience creating robust datasets for deep learning tasks. My team and I created the largest publicly available dataset for Indian Sign Language with over 26,000 images back in my undergrad (https://github.com/RealSign62/RealSign-Indian-Sign-Language-Dataset). We used all sorts of techniques from adjusting lighting to data augmentation to create a robust dataset. I'm also open to experimenting with Blender to create a new dataset for this project, although I have no experience with 3D modeling (and would very much prefer a teammate who knows how Blender works).

   I think all of this makes me a not-so-bad candidate to work on this project. I'd like to see what the professors think about this too.

   Here is how I would approach to solve this problem:

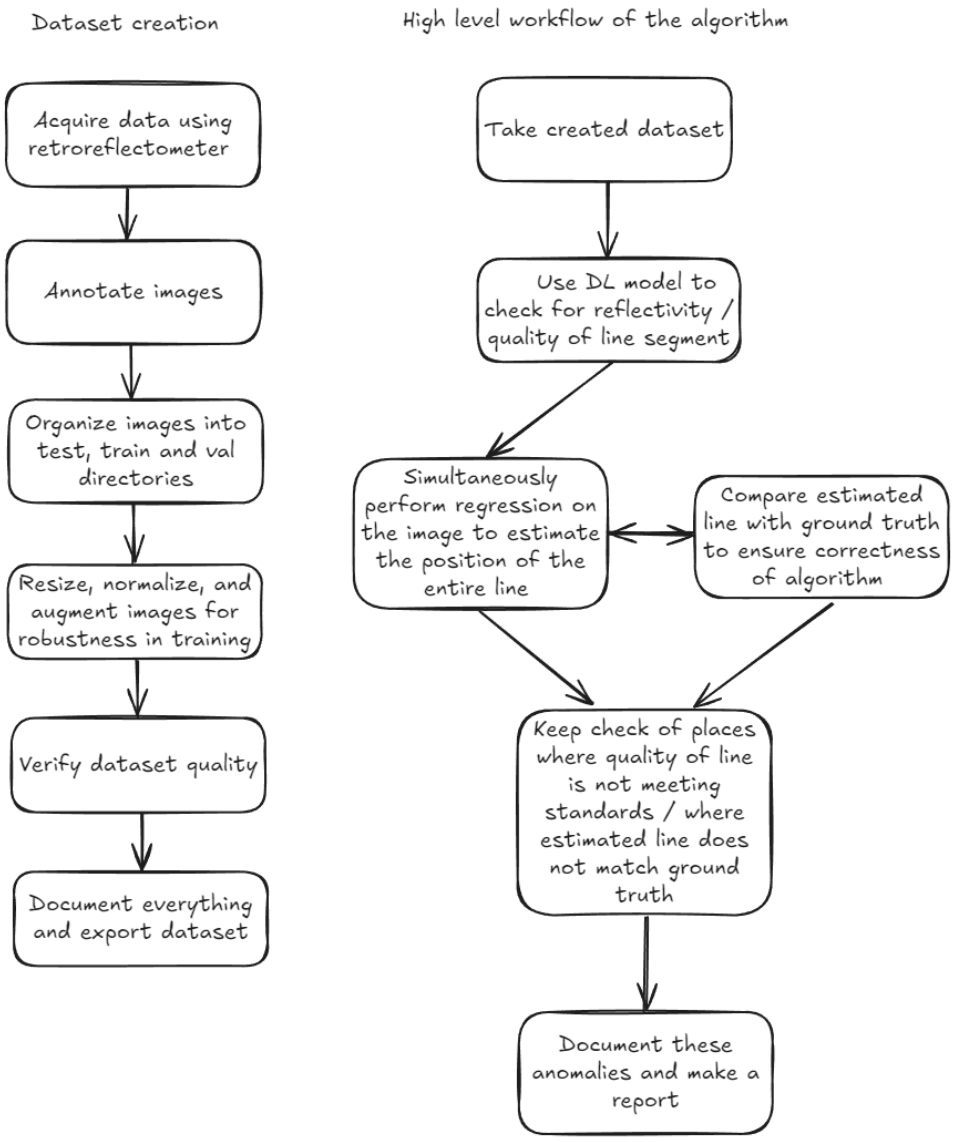

Figure 1: Triangular iteration space

2. **Reinforcement Learning for Road Sign Detection and False Positive Elimination**

I've also have a good background on Reinforcement Learning and this would be my second choice of project. Although I haven't personally tried to finetune deep learning models, I have necessary background with how Reinforcement Learning with Human Feedback (RLHF) works and have also worked on a project related to the same (https://github.com/Sharath04Satish/the-rlhf-repository) where we try to figure out the best form of human feedback for a reinforcement learning agent to better learn it's policy (and better suited to the human to provide).

There are also multiple research papers going over how to fine tune a deep learning model using human feedback in a reinforcement setting. There was a course called Human-AI alignemnt offered at Utah in Fall 2023 (which I took) that went over in great detail about how to use reinforcement learning to "align" deep learning models.

I can make another flowchart for this if necessary. (Let me know if you'd like to see this too)

3. **Deep Learning-Driven 3D Reconstruction, Data Synthesis, and Size Estimation for Road Signs**

   This is my final choice of project.

   I do not know how Blender works, and would appreciate if I was teamed up with a person who does if I get assigned this project.

   I can help out with the CNN / transformer aspect of this project as that is where I'm knowledgeable at.