# JAMES RENNIE BEQUEST

# REPORT ON EXPEDITION / PROJECT / CONFERENCE

| | |
|---|---|
| **Expedition/Project/ Conference Title:** | Computer Vision 4 Ecology Summer School at the California Institute of Technology |
| **Travel Dates:** | July 25th to August 21st, 2022 |
| **Location:** | Pasadena, California, USA |
| **Group member(s):** | Alixandra N. Prybyla, PhD Student |
| **Aims:** | The aim of this expedition was to learn the skills necessary to build machine learning (ML) models that could be used in service of ecological questions. Specifically, I aimed to build an ML model to answer a question for my PhD thesis. |

**Photography consent form attached:** *(please refer to your award letter)*

☒ Yes
☐ No

**OUTCOME (a minimum of 500 words):-**

Computer vision is a field of artificial intelligence focused on training computer systems to extract information from digital inputs (such as images, sounds, and videos) and offer insights based on what has been "seen." While the field has existed since the 1960s, only recently has it been applied to ecological questions, thus allowing researchers to approach large, unwieldy datasets in novel ways. Its novelty means that there is only a small number of world experts—and it was because of the James Rennie Bequest that I was able to access many of those experts for the benefit of my PhD thesis. The CV4Ecology Summer School was a transformative experience for me and has changed the trajectory of my PhD work—and my career—for the better.

To give you a better idea about my learning outcomes while at the CV4Ecology Summer School, I will break down what we did week-by-week. The first week of the course was dedicated to machine learning basics: we were given lectures on dataset prototyping and visualisation, working on virtual machines, data splitting and data poisoning, training our models (and all the intricacies that come with that, such as configurations, launching, monitoring, checkpointing, and organisation), and working with open-source computer vision codebases. These lectures were broken up by working time where we could apply what we had learned to craft our own models with the datasets we brought with us. We were also given the option to join a reading group of our choice. I opted to join a journal club on the "long-tail problem," or the statistical phenomenon that occurs when you have an uneven distribution of data (something that I experience in my own dataset) and how to approach this in a machine learning context. In between lectures, journal club, and working time with teaching assistants that helped us with our programming, we were invited to listen to guest lectures. These lecturers were academics and professionals that had changed the scope of the field.

The second week of the course was dedicated to evaluating the models we had begun building in the first week, offline evaluation and analysis, ways that we could improve the results we were getting from the first iterations of our models, understanding the benefits and drawbacks of data augmentation, and ways in which we could improve training datasets

(with things like weak supervision, self-supervision, targeted relabelling, and anomaly detection). At the end of this week we visited Conservify, an open-source hardware development group that focused on conservation issues. There I was able to learn about potential field tools I would be able to use for my own research. I appreciated this because, for most of the academic year, I am gathering data in the field.

The final week was all about polishing the models we had been running. Many of us had goals about what we wanted to leave the Summer School with, and so there were many long hours of programming, debating, sharing ideas, sharing issues, and working together as students, teachers, and TAs. The lectures this week focused on fair comparisons and ablation studies, efficient models and speed vs. accuracy, and thoughts on serving, hosting, and deploying models. At the end of the course we were invited to present our work to the group and share how much we had learned, how far we'd come, and communally brainstorm about how we could improve on what we had built together there. I left the course with three iterations of the model I had wanted to build for my PhD work. It was more than I could have hoped for.

In addition to the tangible skills I gained during the Summer School, I also gained access to a community of individuals who are also pushing the boundaries on what is possible in the field of ecology through artificial intelligence. These people will be future collaborators and, I am certain, long-time friends.

The field that I am now navigating with competence is often seen as highly experimental. It is therefore often difficult to access funding for my experiments. In light of this, I would just like to add a note here to say thank you to the James Rennie Bequest team: you took a chance on me and the results are better than I could have hoped for. Thank you for placing your trust in me.