

Managing intensive rotational grazing using reinforcement learning based livestock simulations

Project Summary

Recently developed machine learning (ML) techniques have found applications in many facets of agriculture. These technologies have been used to guide crop yield maximization, weed and disease detection, animal welfare and livestock production, as well as water and soil quality [1]. Despite these advancements, not much work has been done concerning data-driven approaches to livestock management that take long-term environmental constraints into consideration, such as potential defoliation from improper grazing.

Overgrazing is one of the foremost problems faced by the agricultural sector. Improper grazing, characterized by heavy grazing which exceeds the recovery capacity of a pasture, can lead to lower nitrogen content, lower moisture content, and other soil temperature/chemical changes after continued overuse. This leads to a net loss of fertile land, and frequently a population explosion of unpalatable grasses and an overall loss of biodiversity in affected grasslands [2][3]. In all cases, these issues can lead to systemic economic inefficiency and underutilization of arable grasslands for food production. In extreme cases, this loss of arable land can lead to persistent food insecurity, particularly in developing countries where political and economic issues impede effective agricultural management practices, resulting in “tragedy of the commons” outcomes.

One proposed solution to the issue of overgrazing within the agricultural literature is managed intensive rotational grazing (MIRG), in which herds are rotated between pasture areas in such a way that the long-term biomass of foraged vegetation is maximized and areas of bare dirt are minimized [4]. Moreover, intensive grazing procedures which create healthier soil can help sequester carbon, providing a potential source of climate change mitigation [5]. Automated systems for planning and managing rotational grazing are likely to make this practice more cost-effective and feasible, particularly in developing countries with weaker agricultural institutions.

The goal of the project in this proposal is to design and implement a data-driven system for planning and managing rotational grazing, using a reinforcement learning based simulation. Herds are to be modelled as agents in a dynamic grassland environment, and at each period the agents will learn where to be located next in order to maximize vegetation and soil health. Specifically, data on soil quality (nitrogen, carbon, moisture, etc.) will be estimated using spectroscopy, and vegetation biomass/bare soil area estimated using aerial imaging. These features will be used to train agents in a simulation using a deep reinforcement learning algorithm, where the learned policy is a direction of movement or placement of herds, and the reward function is some function of the aforementioned features collected at each period and of livestock health. The efficacy of this system will be evaluated first in robust simulated examples, and then in a study using a variety of willing small-scale farms over a period of time.

Intellectual Merit

The primary intellectual contribution of this project would be the novel extension of recent reinforcement learning techniques to an agricultural problem which has heretofore not been addressed much by the machine learning community – namely, learning ideal grazing patterns with the goal of optimizing long-term vegetation and soil health. Traditional statistical approaches to this problem have run into numerous issues: due to the high variability of the impact of grazing on vegetation depending on soil type, foliage types, climate, etc. an objective function for traditional supervised learning is hard to define. Indeed, the problem of designing a general way to manage grazing is difficult with a model-rich algorithm. Reinforcement learning addresses this problem by alleviating the requirements of structured models, instead formulating the problem in terms of actions and observable consequences resulting from them.

The use of reinforcement learning in solving this task would push the boundaries of applied machine learning in agricultural problems, by showing a viable machine learning approach to problems where supervised objective functions are difficult to define, or where extensive input/output examples are difficult to come by. The completion of this project would also entail an end-to-end system, which itself would lead to incorporation of previous research regarding estimation of soil content and vegetation biomass in a novel way.

Broader Impact

The broader impact of the project in this proposal generally falls under two categories: 1) reducing global food insecurity, and 2) mitigating climate change through decreasing the impact of agriculture on carbon emissions. The health of grasslands has a large impact on both of these categories – as the world population increases, larger food supplies are required, and even small increases in efficiency of livestock grazing can greatly increase the global food supply [6]. Furthermore, grasslands are very important in storing excess carbon from the atmosphere, and have the potential to sequester more carbon with effective grazing techniques [6][5].

A reinforcement learning driven approach to planning and managing effective rotational grazing techniques is expected to simultaneously increase the accuracy and cost efficiency of grazing management, as well as make grazing management individually feasible in developing nations which lack stable central agricultural institutions, preventing prevalent “tragedy of the commons” outcomes. The resulting increases in food security, and mitigation of the negative effects of climate change on developing nations, would go towards strengthening the economic and geopolitical positions of countries previously underrepresented on the global stage.

Proposed Research

The purpose of this project is to design an end-to-end system for simulating ideal rotation of livestock grazing which can be deployed and effectively used in real-world scenarios. As such, the project decomposes into several objectives.

First, cost-effective methods for estimating the data required for MIRG – chiefly vegetation biomass, soil quality, and area of exposed soil – are required. Previous research has successfully employed ML techniques towards the first two tasks. One such study used a

artificial neural network (ANNs) and an adaptive neuro-fuzzy inference system (ANFIS) to estimate vegetation biomass in managed grasslands using satellite imagery [7]. Another study used support vector machines (SVM) and Cubist regression to estimate total soil nitrogen, organic carbon, and moisture content using infrared spectroscopy [8]. Using ANNs to map weeds in fields using unmanned aircraft systems (UAS) has also been explored [9], an approach which can potentially be extended to detect bare soil in grasslands. As a first step, we will make use of these techniques to train models necessary for our purposes. These trained models will serve as a cost-effective way of estimating the objectives of MIRG using more readily available measures, such as spectroscopy and aerial imaging.

Secondly, a simulation needs to be designed. Research will be done into how to represent grasslands in a simulation which is both effective and sufficiently realistic, taking into account parameters such as water location, shade from foliage, etc. Before designing the reinforcement learning system, design choices need to be made regarding how to model the agent. There are two general ways to do this, which we will explore: 1) individual herd animals will be agents in the simulation, and will have available actions consisting of a direction and magnitude of movement in each period, subject to constraints (e.g. fencing), or 2) the agent in the simulation is the manager itself, which chooses where to place a herd in each period.

In the third stage, a reinforcement learning algorithm will be researched, designed, and tested. Formally, a model for general reinforcement learning consists of a discrete set of environmental states, a set of possible agent actions, and a reward function which maps the environmental state to a scalar reinforcement signal. The goal of the agent is to find a policy mapping states to actions which maximizes the agent's long-term reward (typically geometrically discounted over an infinite horizon) [10]. Frequently, reinforcement learning algorithms assume that the world resembles a Markov Decision Process, where the current state is conditionally independent to states and actions other than the previous one. This assumption appears to work well with this particular application insofar as vegetation and soil health depends only on vegetation and soil health in the previous period, as well as the action of the agent. However, possible extensions which will be considered involve adding environmental noise to the simulation, such as drought or rain, during the training process in order to make the learned policies more robust.

In our case, the reward function (regardless of the choice of agent made above) will be a mapping of the environment to a function of the aforementioned features (vegetation biomass, soil quality, etc.) as well as features related to livestock health/production in order to constrain outcomes to be financially viable. In refining the formalization of the reinforcement learning task, we will also choose an algorithm learning. One architecture for reinforcement learning which has seen high popularity and effectiveness in practical applications is Q-learning, which is known to converge independently of exploration during learning.

In order to carry out this research, our primary resource requirements consist of adequate computing power to train the requisite models and simulations, a small team consisting of qualified experts in agriculture and machine learning, and access to agricultural datasets containing the data required in the first stage.

Evaluation Plan

In order to evaluate the reinforcement learning algorithm and any hyperparameters during development, simulations will be created using simplified plant growth models under various conditions of soil type, foliage type, climate, etc. The efficacy of the rotational grazing policy learned by the agent will be compared to agents in the simulation which instead use hard-coded or manually-operated policies, such as continuous grazing or non-automated rotational grazing.

Following this, a variety of small farms willing to participate in a study will be split into a control group and an experimental group, and our system will be used by the latter group to manage rotational grazing for an extended time. Vegetation biomass and soil quality will be carefully monitored and compared between the experimental group and the control group during each period to evaluate the success of the reinforcement learning based MIRG system.

References

1. K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, "Machine Learning in Agriculture: A Review". *Sensors*, 2018.
2. Q. Wang and O. Batkhisig, "Impact of Overgrazing on Semiarid Ecosystem Soil Properties: A Case Study of the Eastern Hovsgol Lake Area, Mongolia". *Journal of Ecosystem & Ecography*, 2014.
3. N. Hassani, H. R. Asghari, A. S. Frid, and M. Nurberdief, "Impacts of Overgrazing in a Long Term Traditional Grazing Ecosystem on Vegetation Around Watering Points in a Semi-Arid Rangeland of North-Eastern Iran". *Pakistan Journal of Biological Sciences*, 2008.
4. R. Smith, G. Lacefield, R. Burris, D. Ditsch, B. Coleman, J. Lehmkuhler, and J. Henning, "Rotational Grazing". *College of Agriculture, University of Kentucky*, 2011.
5. M. Bogaerts, L. Cirhigiri, I. Robinson, M. Rodkin, R. Hajjar, C. Costa Junior, P. Newton, "Climate Change Mitigation Through Intensified Pasture Management: Estimating Greenhouse Gas Emissions on Cattle Farms in the Brazilian Amazon". *Journal of Cleaner Production*, pp. 1539-1550, 2017.
6. F. P. O'Mara, "The Role of Grasslands in Food Security and Climate Change". *Annals of Botany*, 2012.
7. I. Ali, F. Cawkwell, E. Dwyer, S. Green, "Modeling Managed Grassland Biomass Estimation by Using Multitemporal Remote Sensing Data – A Machine Learning Approach". *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2016.
8. A. Morellos, X. Pantazi, D. Moshou, T. Alexandridis, R. Whetton, G. Tziotzios, J. Wiebensohn, R. Bill, A. M. Mouazen, "Machine Learning Based Prediction of Soil Total Nitrogen, Organic Carbon and Moisture Content by Using VIS-NIR Spectroscopy". *Biosystems Engineering*, 2016.
9. X. E. Pantazi, A. A. Tamouridou, T. K. Alexandridis, A. L. Lagopodi, J. Kashefi, D. Moshou, "Evaluation of Hierarchical Self-Organising Maps for Weed Mapping Using UAS Multispectral Imagery". *Computers and Electronics in Agriculture*, 2017.
10. L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement Learning: A Survey". *Journal of Artificial Intelligence Research*, 1996.