Agbots: Weeding a Field with a Team of Autonomous Robots

Wyatt McAllister, Joshua Whitman, Allan Axelrod, Joshua Varghese, Adam S. Davis, Girish Chowdhary

CPS PI meeting 2019

CPS: Robust Deep Learning for Mechanical Weeding Agbots; NSF#1739874

Abstract

We present techniques for predictive modeling of weed growth and an improved planning index to be used in conjunction for the purpose of improving the performance of coordinated weeding algorithms for robotic agriculture.

Background

Agbot System We work with a team of agbots used for the mechanical control of herbicide-resistant weeds [1].

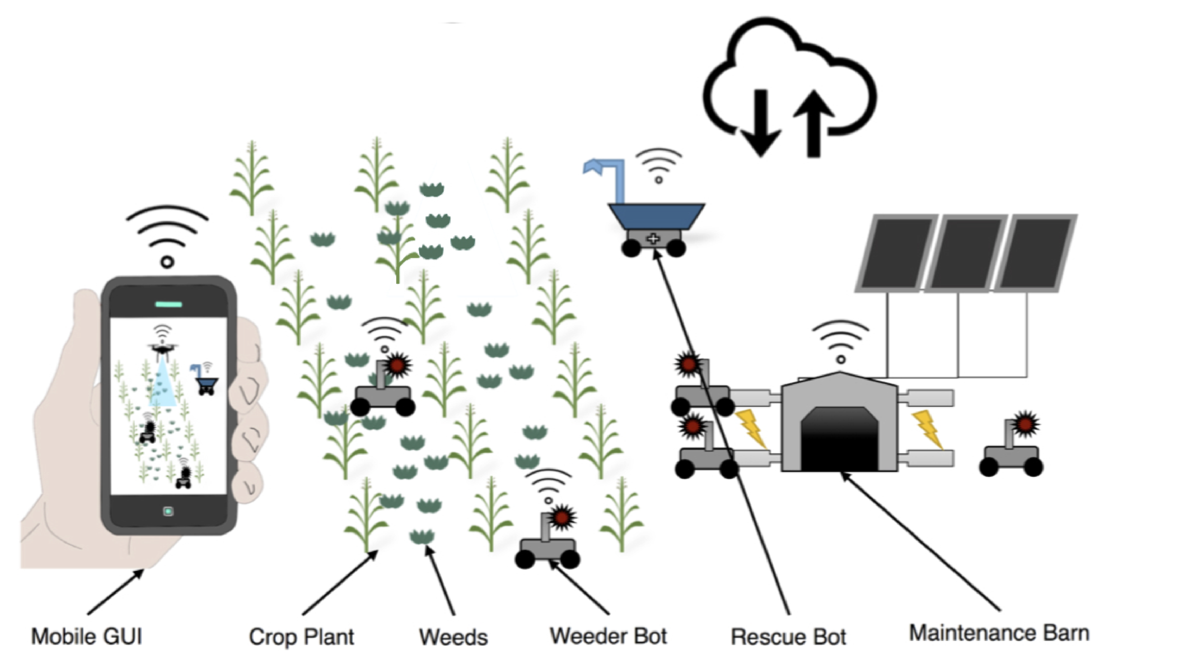

Figure 1: Our solution for persistent autonomous weed control is a collaborative team of mechanical weeding robots.

Simulation Framework For this work, we use an ecologically realistic simulation called Weed World [1].

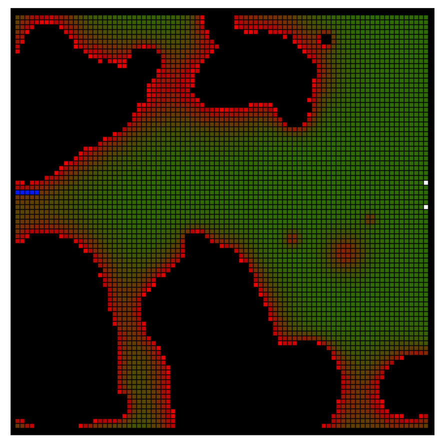

Figure 2: Weed World, designed in Python. Each cell represents a small 0.8 m square portion of the field. The colors of the squares represent weed seed bank density, from light to dark. The agents are shown in solid blue.

This simulation was developed with respect to the physical characteristics and capabilities of the TerraSentia robot [2].

Figure 3: Terra Sentia Robot

Figure 4: Weeder

Methods

Evolving Gaussian Processes (E-GP): This work makes use of a prediction scheme, Evolving Gaussian Processes (E-GPs)[3, 4], where the weights, w_T, in a gaussian process over the spatial domain are evolved forward using a linear operator, A. We can make predictions of weed density, y, anywhere in the spatial domain using the feature map $\hat{\phi}(x, y)$.

$$w_T = A w_{T-1} \tag{1}$$

$$y(x,y) = (w_T)^T \hat{\phi}(x,y) = (A w_{T-1})^T \hat{\phi}(x,y) \tag{2}$$

Predictive Baseline: We wish to compare E-GP to a more simplified model-based approach, which will take advantage of known properties of the weed growth model, such as the fact that it approaches the seed bank density, $S_0(x,y)$, at an exponentially decaying rate. We propose the following candidate emergence model.

$$\zeta(x,y,t) = S_0(x,y)\left(1 - 0.1^{\frac{t}{2\,\text{mo.}}}\right) \tag{3}$$

In this model, to predict the total emerged weed density, $\zeta(x,y,t)$, we use a single time-invariant Gaussian process. This process takes in the observations, $\bar{\zeta}(x,y,t)$, and predicts the initial seed bank density $\hat{S}_0(x,y)$.

Entropic Value at Risk (EVaR): Each robot must make a decision about which row to weed next after it has finished a row. This is essentially a bandit problem. Gittin's Index, $G(X_i)$, is known to be optimal metric for planning on tasks with an uncertain termination time and known statistics [5]. We compute the value for weeding a particular row with Gittins index using EVaR [6], which is a principled way to optimize in terms of the reward and information gain. Here, $T_i(x_i(t), a_i(t))$ is the proposed time to weed row $a_i(t)$ given agent i is at position $x_i(t)$, γ is a learning rate, and $R_i(a_i(t))$ is the estimated reward for weeding row $a_i(t)$.

$$G(X_i) = \sup_{a_i}\left\{\frac{\gamma^{T_i(x_i(t),a_i(t))}\text{EVaR}\left[R_i(a_i(t)); 1 - e^{-D_{KL}(Q\|P)}\right]}{\sum_{t=0}^{T_i(x_i(t),a_i(t))}\gamma^t}\right\}$$

Results

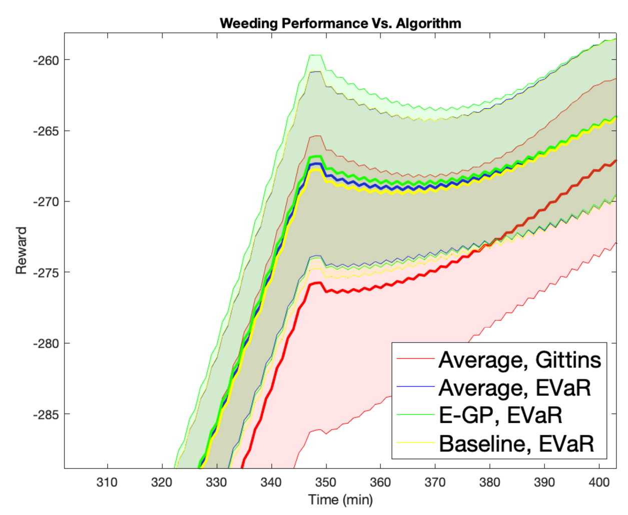

Figure 5: Comparison of Algorithms

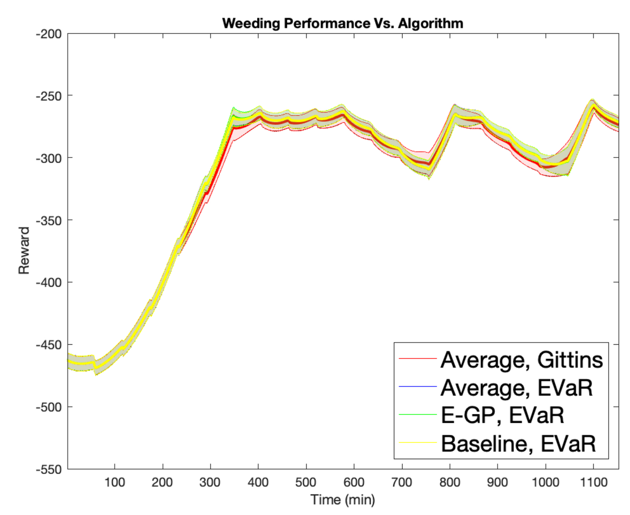

Figure 6: Comparison of Algorithms

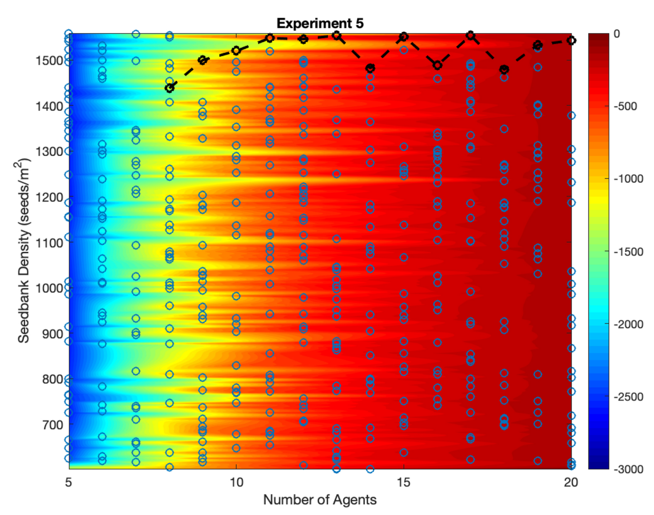

Figure 7: Old Planner and Predictor

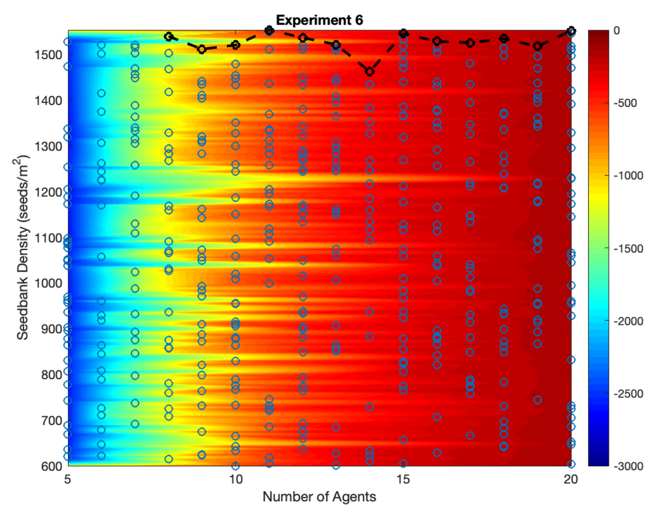

Figure 8: New Planner, Old Predictor

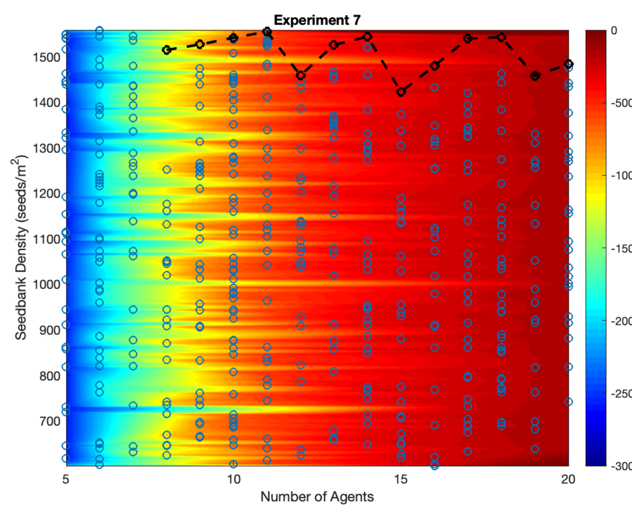

Figure 9: New Planner, E-GP

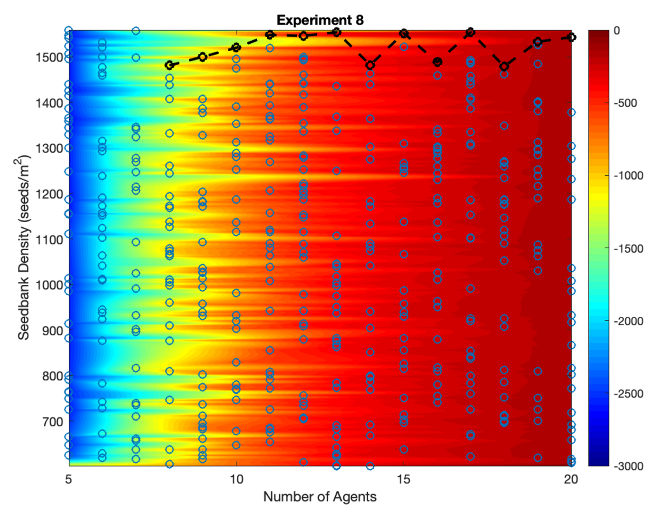

Figure 10: New Planner, Baseline

Results

The results show that when EVaR is used for the planning index, decision making is improved, even without a refined prediction scheme such as E-GP, or the baseline predictor. From Figures 7, and 8, we see that the use of EVaR improves robustness, allowing the algorithms using EVaR to succeed in cases when the algorithm using only Gittin's index does not.

With the addition of refined prediction strategies, we see only slight performance changes as shown in Figure 5. In terms of how many agents can be used to weed fields with a given average seed bank density, the system using E-GP has comparable robustness to the systems using other prediction strategies.

The notable contribution of E-GP is that it provides higher predictive accuracy to the previous course prediction scheme, while using less information than the baseline predictor.

Conclusion

+ The addition of EVaR improves both planning performance, and the robustness of the system to weeding fields with high seed bank densities using a limited number of agents.
+ EVaR used in conjunction with E-GP or the Baseline Predictor provides a further performance improvement, though the robustness remains the same.
+ The similar robustness between algorithms with various prediction schemes is theorized to be due to the information gain term in EVaR dominating decision making during early stages of weeding.

References

[1] Wyatt McAllister, Denis Osipychev, Girish Chowdhary, and Adam Davis. Multi-agent planning for coordinated robotic weed killing. In *Intelligent Robots and Systems (IROS), 2018 IEEE/RSJ International Conference on*. IEEE, 2018.

[2] TerraSentia. Terrasentia robot - earthsense, inc. *TerraSentia*, 2017.

[3] Joshua Whitman and Girish Chowdhary. Learning dynamics across similar spatiotemporally-evolving physical systems. In *Conference on Robot Learning*, pages 472–481, 2017.

[4] Hassan Kingravi, Harshal Maske, and Girish Chowdhary. Kernel observers: Systems theoretic modeling and inference of spatiotemporally varying processes. In *Advances in Neural Information Processing Systems (NIPS)*, Barcelona, Spain, 2016. accepted.

[5] Richard Weber. On the gittins index for multiarmed bandits. *The Annals of Applied Probability*, pages 1024–1033, 1992.

[6] Amir Ahmadi-Javid. Entropic value-at-risk: A new coherent risk measure. *Journal of Optimization Theory and Applications*, 155(3):1105–1123, 2012.

Acknowledgement

Supported by NSF CPS 1739874 under USDA NIFA 2018-67007-28379