Model-based scenario analysis for effective site-specific weed control on grassland sites

Lukas Petrich[a,*], Georg Lohrmann[b], Fabio Martin[b], Albert Stoll[b], Volker Schmidt[a]

[a] *Institute of Stochastics, Ulm University, D-89069 Ulm, Germany*
[b] *Hochschule für Wirtschaft und Umwelt Nürtingen-Geislingen, Neckarsteige 6-10, D-72622 Nürtingen, Germany*

Abstract

The site-specific management of weeds in grassland is often challenging because different weed control strategies have different trade-offs regarding the required resources and treatment efficiency. So, the question arises whether a wide tractor-based system with section control or a small agricultural robot has a higher weed control performance for a given infestation scenario. For example, a small autonomous robot moving from one weed to the next might have much shorter travel distances (and thus lower energy and time costs) than a tractor-mounted system if the locations of the weeds are relatively isolated across the field. However, if the plants are highly concentrated in small areas so-called clusters, the increased width of the tractor-mounted implement could be beneficial because of shorter travel distances and greater working width.

An additional challenge is the fact that there is no complete knowledge of the weed locations. Weeds may not have been detected, for example, due to their growth stage, occlusion by other objects, or misclassification. Weed control strategies must therefore also be evaluated with regard to this issue. Thus, in addition to the driving distance, other metrics are also of interest, such as the number of plants that were actually controlled or the size of the total treatment area.

We performed this investigation for the treatment of the toxic *Colchicum autumnale*, which had been detected in drone images of extensive grassland sites. In addition to real data, we generated and analyzed simulated weed locations

*Corresponding author
Email address: `lukas.petrich@uni-ulm.de` (Lukas Petrich)

using mathematical models of stochastic geometry. These offer the possibility to simulate very different spatial distributions of toxic plant locations. Different treatment strategies were then virtually tested on this data using Monte Carlo simulations and their performance was statistically evaluated.

Keywords: weed control strategy, tractor-mounted implement, autonomous robot, stochastic model, partial information

1. Introduction

In the last years the interest in site-specific weed management tools has grown substantially as they allow for a minimum impact on the environment and, at the same time, reduce operational costs (Schellberg et al., 2008; Christensen et al., 2009; Wegener, 2020). This interest led to a diverse set of treatment tools, all of which with different advantages and disadvantages (Christensen et al., 2009; Machleb et al., 2020; Stoll, 2020; Martin et al., 2022). So for example the question arises how a treatment tool fares with different spatial distributions and severities of weeds on a field. In addition to this, site-specific methods rely on precise localization of the targeted weeds. Regardless of how these locations are acquired, a complete and perfect survey of the weed population is practically infeasible. It is therefore also important to investigate how weed management based on partial information performs with respect to the whole (unobserved) weed population. The easiest and most flexible way to answer these questions is to rely on computer simulations. Various models for a broad variety of weeds have been proposed in the past (Holst et al., 2007; Freckleton and Stephens, 2009; Somerville et al., 2020). With these, answers to questions have been found predominantly about cultural weed management such as the benefits of landsharing versus landsparing (Colbach et al., 2018), the management of parasitic plants (Pointurier et al., 2021), the effect of sowing strategies with regard to weeds (Andrew and Storkey, 2017), or even complex crop-weed interactions (Colbach et al., 2021). Nevertheless, it is still an open question how the treatment performance of site-specific treatment tools can be compared for different infestation scenarios. This holds even more so when focusing on grassland sites instead of arable land (although there are studies for similar problems such as route planning (Maini et al., 2022)) and accounting for partial observations of the weeds.

For these reasons we performed a virtual scenario analysis to compare different tools for physical weed management in grassland. Simulated weed locations allowed to create several prototypical infestation scenarios with varying severities. While in the real world these scenarios may not occur in isolation, they form recurring patterns that make up realistic infestations and understanding their performance characteristics allows to transfer the results of the present paper to some extent to an unobserved real field. Additionally, experimental observations of *Colchicum autumnale* plants on an extensive grassland site were also employed to provide a realistic baseline for the simulation study. Here, *C. autumnale* are toxic plants, which pose a thread to farm animals especially in hay or silage. Note, however, that the present paper does not present a thorough (stochastic) model for the experimental *C. autumnale* locations, which would require much more empirical data over a longer timespan. The focus is to simulate the whole process of (partial) observation of the weed locations on the field, routing of the tools, and the treatment of the targeted weeds. With the proposed simulation framework it is possible to illuminate the aforementioned questions in silico without the need for extensive experimental setup (e.g., building of the treatment tool, finding suitable fields, etc.). This provides an inexpensive way to test machine specifications before constructing a prototype of the treatment tool, or helps potential buyers to decide which tool on the market best fits their usage scenarios. For the present paper, the considered treatment tools are inspired by the non-chemical weed control tools developed in Stoll (2020) and Martin et al. (2022)—that is a small autonomous robot with circular mower and a water-hydraulic tractor-based system with section control—, but the results apply also to some extent to other tools such as site-specific herbicide sprayers, or spot-spraying through unmanned aerial vehicles (UAV) if overlapping treatment areas can be excluded, for example, with suitable application maps. Here, the two considered treatment tools differ primarily in terms of flexibility—the autonomous robot is able to move from one targeted location to the next, while the tractor traverses the field in meandering lines—and working width. One of the main goals of the present paper is thus to evaluate in which infestation scenarios one treatment tool outperforms the other especially in the light of unobserved weed locations.

2. Materials and methods

The present paper deals with two kinds of weed locations—experimentally observed and simulated ones. In this section, we first describe the acquisition and preprocessing of the experimental dataset. After that, the methods for simulating weed locations and the virtual weed control tools are laid out.

2.1. Data acquisition

The experimental dataset consists of locations of *C. autumnale* flowers that have been extracted from drone images. The images were acquired on September 3rd and 10th, 2019, on an extensive grassland field near Nürtingen, Germany. The field had an area of about $8256\,\mathrm{m}^2$, see Figure 1. At the two observation dates different *C. autumnale* plants were blooming leading to varying numbers of visible flowers. The camera, a Sony alpha 7 RII with a CMOS full-frame 42.4 MP image sensor and its lens with 24 mm focal length, was mounted on a HiSystems MK ARF-OktoXL 4S12 octocopter. The route of the drone was about 10 m above ground and was chosen such that the resulting images had an overlap of about 55%.

The individual drone images where stitched to two orthomosaics (one for each observation date) using the Agisoft Metashape software. In the same step, the orthomosaics were georeferenced by matching markers (ground control points) that were placed on the field and whose GPS positions were obtained by real-time kinematic positioning with their corresponding pixel coordinates.

2.2. Data preprocessing

2.2.1. Image registration

In addition to the georeferencing, the overlap between the two orthomosaics has been further improved by pixel-based image registration. For this, matching pairs of control points in the two images have been created by visual inspection at objects such has fence posts or trees that remained unchanged between the two observations. Then, MathWorks MATLAB was used to find a coordinate transformation that transforms one orthomosaic to match the second one. Here, the 'projective' transformation type (Jähne, 2005; MathWorks, 2021) was chosen out of 'nonreflective', 'similarity', 'affine', 'projective', and 'polynomial' with degrees up to 4 (see MathWorks (2021)) as a trade-off between minimizing the Euclidean distance between the first set of control points and the transformed control points of the second orthomosaic and visual goodness-of-fit. For both aligned orthomosaics a region of interest $\mathcal{W} \subset \mathbb{R}^2$ was defined by removing all parts of the images that did not show the considered field or where information from one observation was missing. In the following, only this cutout of the orthomosaics is considered.

2.2.2. Observed weed locations

For the two aligned orthomosaics, *C. autumnale* flowers were automatically detected using the predictor developed in Petrich et al. (2020). Manual checking of each predicted weed location and visual inspection of the remaining images ensured an accurate survey of weed locations at the two observation dates. Here, we considered locations that are closer than 5 cm to each other to be the same weed and replaced these locations with their centroid.

The resulting weed locations are shown in Figures 1a and 1b. In total there were $n_1^{(\textsc{Exp})} = 550$ detected flowers in the observation from September 3rd, 2019 and $n_2^{(\textsc{Exp})} = 1792$ flowers from September 10th. In the following we refer to the first as $\textsc{Exp}^1$ and to the latter as $\textsc{Exp}^2$. It is interesting to note that only about 20% of the weed locations in the smaller dataset $\textsc{Exp}^1$ had weed locations closer than 2 m in the other dataset $\textsc{Exp}^2$. This indicates that a large amount of

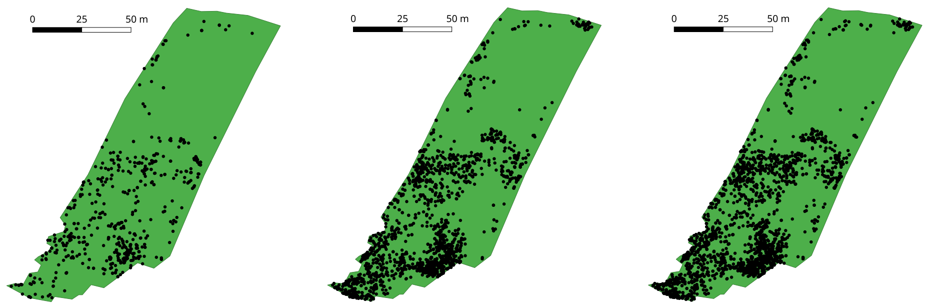

(a) Experimentally observed weed locations from the 3rd September, 2019 (Exp^1).

(b) Experimentally observed weed locations from the 10th September, 2019 (Exp^2).

(c) Combination of the observed weed locations Exp^1 and Exp^2 acting as the ground truth locations Exp.

Figure 1: Visualization of experimental data. The green area corresponds to the whole field.

information is gained by considering both datasets instead of simply the larger one.

2.2.3. Experimental ground truth weed locations

Obviously, the two datasets Exp^1 and Exp^2 were partial observations of the complete weed population on the field. This incompleteness of the observations could have multiple reasons such as the plants were not yet in a growth state in which they can be detected, they were occluded by other objects, etc. For the present simulation study, we made the simplifying assumption that the locations of the whole population can be obtained by combining the observed datasets Exp^1 and Exp^2, see Figure 1c. This is possible since we do not aim to provide an accurate modeling of the true weed population, but rather only use the experimental data as a rough baseline for our scenario analysis. Note, however, that for the combination, we again considered locations closer than 5 cm to each other to be the same weed and replaced these locations with their centroid. The resulting experimental ground truth weed location dataset, denoted Exp in the following, comprised a total of $n^{(\textsc{Exp})} = 2313$ locations (29 locations were thus found in both observed datasets). So in summary, we have a ground truth dataset (Exp) comprised of all weed locations and two subsets thereof

(Exp^1 and Exp^2) corresponding to the weeds that were observed at the two observation dates (September 3rd and 10th, respectively). This general setup will be maintained for the simulated data, see Section 2.3.

2.3. Simulated data

In addition to real experimental data, we also considered simulated data, i.e., we generated and analyzed simulated weed locations using mathematical models of stochastic geometry, see Chiu et al. (2013). This allowed us to investigate various scenarios not readily available with experimental data. By investigating these prototypical scenarios, clues can be obtained for the treatment performances on a given infestation. The general setup was to first simulate ground truth weed locations. However, in practice only a subset of the actual weed locations is observed. We imitate this partial information problem with a second step where we remove some locations from the ground truth datasets, which 'were not observed'.

2.3.1. Ground truth weed location model

The simulated ground truth weed locations $s_1^{(\mathrm{gt})}, \ldots, s_{n^{(\mathrm{gt})}}^{(\mathrm{gt})}$ were drawn from a stochastic point-process model, where $n^{(\mathrm{gt})}$ denotes the total number of weed locations generated in the area under consideration. More specifically, the sequence of simulated locations $s_1^{(\mathrm{gt})}, \ldots, s_{n^{(\mathrm{gt})}}^{(\mathrm{gt})}$ was obtained as a realization of an inhomogeneous Poisson point process $S_1^{(\mathrm{gt})}, S_2^{(\mathrm{gt})}, \ldots$. In the following, we give a short introduction of the mathematical background and refer to Chiu et al. (2013) for a more in-depth discussion.

Consider a bounded sampling window $\mathcal{W} \subset \mathbb{R}^2$, which in our case coincides with the considered field visualized in Figure 1, and the expected number of weed locations $\Lambda(B)$ for a subset $B \subset \mathcal{W}$ ('parts of the considered field') given by the integral $\Lambda(B) = \int_B \lambda(x)\,\mathrm{d}x$ of a (non-negative) intensity function $\lambda : \mathcal{W} \to [0, \infty)$. Then, one says that a sequence of random locations $S_1, S_2, \ldots \subset \mathcal{W}$ follows an inhomogeneous *Poisson point process* with intensities $\Lambda(B), B \subset \mathcal{W}$, if the following conditions are fulfilled: Consider the random number of points

$\Phi(B) = \#\{S_i : S_i \in B \text{ for } i = 1, 2, \dots\}$ in any test set $B \subset \mathcal{W}$, where $\#(\cdot)$ denotes set cardinality, and assume that

(a) the random variable $\Phi(B)$ is Poisson distributed, i.e.

$$\mathbb{P}(\Phi(B) = n) = \frac{\Lambda(B)^n}{n!} \exp(-\Lambda(B)) \qquad \text{for each } n = 0, 1, \dots,$$

(b) the random numbers of points $\Phi(B_1), \dots, \Phi(B_k)$ in k pairwise disjoint (i.e. non-overlapping) test sets $B_1, \dots, B_k \subset \mathcal{W}$ are independent of each other, for each $k = 2, 3, \dots$.

Note that condition (a) implies that the expectation of the random number of points $\Phi(B)$ in the set $B \subset \mathcal{W}$ is given by $\mathbb{E}\,\Phi(B) = \Lambda(B)$. Thus, indeed, $\Lambda(B)$ measures the expected number of weeds in a given part $B \subset \mathcal{W}$ of the field, and the intensity function $\lambda : \mathcal{W} \to [0, \infty)$ governs the spatial distribution of the simulated weeds.

In Section 3, we consider various (virtual) scenarios where the weed locations have different spatial distributions. These scenarios were modeled by corresponding choices of the intensity function λ. For this, the function λ was chosen such that the expected number of points $\mathbb{E}\,\Phi(\mathcal{W}) = \int_{\mathcal{W}} \lambda(x)\,\mathrm{d}x$ in the sampling window $\mathcal{W}$ was set equal to the number of weed points $n^{(\textsc{Exp})}$ in the experimental ground truth dataset Exp multiplied by some factor $\lambda^* > 0$, which we call intensity factor hereinafter, i.e., $\int_{\mathcal{W}} \lambda(x)\,\mathrm{d}x = \lambda^* n^{(\textsc{Exp})}$. More precisely, we chose λ by considering a certain basis function $\lambda_0 : \mathcal{W} \to [0, \infty)$ and a normalizing factor $\alpha > 0$ such that $\int_{\mathcal{W}} \alpha\lambda_0(x)\,\mathrm{d}x = n^{(\textsc{Exp})}$. Then, in a second step, we multiply $\alpha\lambda_0(x)$ by the intensity factor λ^*, i.e., $\lambda(x) = \lambda^*\alpha\lambda_0(x)$ for each $x \in \mathcal{W}$.

Note that by considering different kinds of basis functions λ_0 we were able to generate different types of weed distribution patterns, see Figure 2. Moreover, the intensity factor λ^* allowed us to investigate different degrees of severity of the weed infestation without changing its spatial distribution.

The stochastic ground truth dataset models—given primarily by their corresponding basis function λ_0—are described in the following.

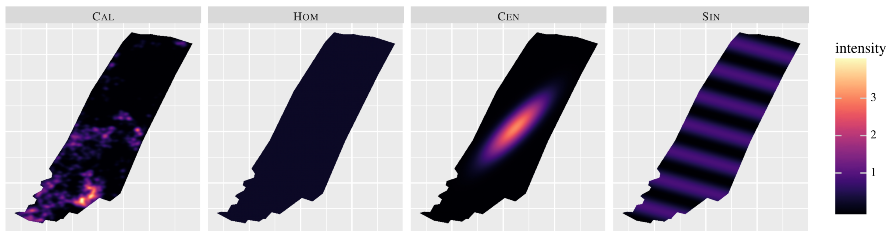

Figure 2: Contour plots of the intensity functions λ with intensity factor $\lambda^* = 1$ of the inhomogeneous Poisson point processes used for the simulated ground truth datasets.

Calibrated ground truth model Cal. The model Cal was calibrated to the experimental ground truth dataset Exp by employing a non-parametric kernel smoothing function (see, e.g., Hastie et al. (2009)) as basis functions λ_0. More specifically,

$$\lambda(x) = \lambda^* \alpha \sum_{i=1}^{n^{(\mathrm{Exp})}} k^{(\mathrm{Cal})}\left(\|x - s_i^{(\mathrm{Exp})}\| \right) \qquad \text{with } k^{(\mathrm{Cal})}(z) = \exp\left(-\frac{z^2}{2 h^{(\mathrm{Cal})^2}} \right)$$

for $x \in \mathcal{W}, z \geq 0$, the normalization factor α, and some bandwidth parameter $h^{(\mathrm{Cal})} > 0$, where $\|x - s\|$ denotes the Euclidean distance of $x, s \in \mathbb{R}^2$. Furthermore, $(s_1^{(\mathrm{Exp})}, \ldots, s_{n^{(\mathrm{Exp})}}^{(\mathrm{Exp})})$ are the weed locations of the experimental ground truth dataset Exp. The bandwidth $h^{(\mathrm{Cal})}$ was chosen by drawing values at random from a gamma distribution with mean 1.5 and standard deviation 1.1. The best value was selected by considering the resulting intensity function as kernel density estimator (through normalization with $1/n^{(\mathrm{Exp})}$) and maximizing the likelihood using cross-validation given the locations in Exp (Loader, 1999; Hastie et al., 2009). In summary, Cal thus closely models the spatial distribution of the experimental data and provides the most realistic simulated dataset, see Figure 2. Compared to Exp, however, in Cal the exact locations of the weeds are drawn at random and can be influenced by choosing different values for the intensity factor λ^*.

Homogeneous ground truth model Hom. The model Hom is based on a homogeneous Poisson process. Then, the intensity function λ is constant, i.e., $\lambda(x) = \lambda^* \alpha$ for each $x \in \mathcal{W}$ with $\lambda_0(x) = 1$ for each $x \in \mathcal{W}$ and $\alpha = n^{(\mathrm{Exp})}/|\mathcal{W}|$, where $|\mathcal{W}| > 0$ denotes the area of the sampling window $\mathcal{W}$. This model represents the case where the weeds are located completely at random across the field $\mathcal{W}$, see Figure 2.

Centered ground truth model CEN. The weed locations of the model CEN are drawn using a bivariate normal distribution with expectation vector $\mu^{(\textsc{Cen})} \in \mathcal{W}$ being the centroid of the sampling window $\mathcal{W}$. More precisely,

$$\lambda(x) = \frac{\lambda^* \alpha}{\sqrt{(2\pi)^2 \det \Sigma^{(\textsc{Cen})}}} \exp\left(-\frac{1}{2}(x - \mu^{(\textsc{Cen})})^{\mathrm{T}} (\Sigma^{(\textsc{Cen})})^{-1} (x - \mu^{(\textsc{Cen})})\right)$$

for each $x \in \mathcal{W}$, where the positive definite dispersion matrix $\Sigma^{(\textsc{Cen})} \in \mathbb{R}^{2\times 2}$ was set equal to

$$\Sigma^{(\textsc{Cen})} = \begin{pmatrix} 218.8 & 324.1 \\ 324.1 & 549.4 \end{pmatrix}.$$

The entries of $\Sigma^{(\textsc{Cen})}$ were determined by visual examination and rescaled to meters. In summary, in the model CEN, the weed locations are concentrated in a central cluster around the expectation vector $\mu^{(\textsc{Cen})} \in \mathcal{W}$, while only few weeds are generated near the boundary of the field $\mathcal{W}$, see Figure 2.

Sinusoidal ground truth model SIN. The intensity function λ of the model SIN arranges the weed locations in sinusoidal waves through the sampling window, where λ is given by

$$\lambda(x) = \lambda^* \alpha \sin\left(\frac{2\pi \left\langle u^{(\textsc{Sin})}, x\right\rangle}{\lambda^{(\textsc{Sin})}}\right) + 2 \qquad \text{for each } x \in \mathcal{W}.$$

Here $< \cdot, \cdot >: \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$ is the dot product with the (unit) normal vector of the wave front $u^{(\textsc{Sin})} \in \mathbb{R}^2$ and the wave length $\lambda^{(\textsc{Sin})} > 0$. The wave direction $u^{(\textsc{Sin})}$ has been set equal to the (normalized) direction of the longest side of the rectangular bounding box of the sampling window $\mathcal{W}$, and $\lambda^{(\textsc{Sin})} = 28.3\,\mathrm{m}$ was chosen by visual examination and rescaled to meters. Therefore, the dataset model SIN produces weed location in a wave-like pattern, see Figure 2.

2.3.2. Observation model

In practice, only some of the weeds in the population can be observed. This can have several reasons such as the plants were not yet in a growth state in which they can be detected, they were occluded by other objects, etc. For our simulation study, we imitated this by eliminating some of the weed locations from the ground truth datasets (so-called *independent thinning*). The remaining locations acted as the observed datasets. Here, we distinguished between two different cases analogous to the experimental dataset: the thinning probability was chosen such that the expected number of locations in the resulting dataset was (i) equal to that in $\textsc{Exp}^1$, and (ii) equal to that in $\textsc{Exp}^2$.

More formally, each of the weed locations $s_1^{(\mathrm{gt})}, \ldots, s_{n^{(\mathrm{gt})}}^{(\mathrm{gt})}$ of a simulated ground truth dataset model (i.e., Cal, Hom, Cen, or Sin) was eliminated with a certain probability independent of its position and the deletion of any other weed location (so-called independent thinning, see Chiu et al. (2013)). As mentioned above, two different thinning probabilities were considered. With $n^{(\textsc{Exp})}, n_1^{(\textsc{Exp})}$, and $n_2^{(\textsc{Exp})}$ the number of weed locations in Exp, $\textsc{Exp}^1$, and $\textsc{Exp}^2$, respectively, the first thinning probability is given by $p_1^{(\mathrm{obs})} = n_1^{(\textsc{Exp})}/n^{(\textsc{Exp})}$. The resulting models correspond to $\textsc{Exp}^1$ and are denoted with superscript '1', e.g., $\textsc{Hom}^1$. The second thinning probability, on the other hand, is given by $p_2^{(\mathrm{obs})} = n_2^{(\textsc{Exp})}/n^{(\textsc{Exp})}$. The resulting models correspond to $\textsc{Exp}^2$ and are denoted with superscript '2', e.g., $\textsc{Hom}^2$. As seen in Section 2.2.2, $\textsc{Exp}^1$ contained fewer observed weeds compared to $\textsc{Exp}^2$. Naturally, this fact translates to the corresponding simulated datasets. For the formal treatment of the present simulation study both thinnings are treated analogously. For simplicity, we therefore denote the $n^{(\mathrm{obs})} \leq n^{(\mathrm{gt})}$ "observed" weed locations selected from the originally simulated weed locations $s_1^{(\mathrm{gt})}, \ldots, s_{n^{(\mathrm{gt})}}^{(\mathrm{gt})}$ as $s_1^{(\mathrm{obs})}, \ldots, s_{n^{(\mathrm{obs})}}^{(\mathrm{obs})}$.

2.4. Treatment strategies

Now that the generation of the simulated datasets is established, we describe two different treatment strategies, which are applied in Section 3 to the experimental and simulated datasets. The treatment strategies are a combination of three parts: (i) the action threshold (see Section 2.4.1) to select the targeted weed locations, (ii) the treatment tool, which is either an autonomous robot (see Section 2.4.2), or a tractor with attachment (see Section 2.4.3), and (iii) the (parameter) configuration of the virtual treatment tool. The considered tools are based on the real-world analogues presented in Stoll (2020) and Martin et al. (2022).

2.4.1. Action threshold

For both treatment strategies considered in this paper, we apply an *action threshold*, which is used as a simple preprocessing method to improve the treatment instead of simply targeting each observed weed locations. More specifically, we ignore isolated weeds, which might drastically increase the distance a treatment tool has to drive, while having only a very limited effect on the treatment quality (see Section 3.1 for a description of suitable performance measures). Furthermore, in Section 3 varying settings for the action threshold are investigated.

Formally, the action threshold $\ell_a > 0\,\mathrm{m}$ caps the nearest neighbor distance (in meter) of the treated weed locations, i.e., given the observed weed locations $s_1^{(\mathrm{obs})}, \dots, s_{n^{(\mathrm{obs})}}^{(\mathrm{obs})}$, the $n^{(\mathrm{tgt})} \leq n^{(\mathrm{obs})}$ targeted weed locations $s_1^{(\mathrm{tgt})}, \dots, s_{n^{(\mathrm{tgt})}}^{(\mathrm{tgt})}$ that are fed into the treatment tools are those points $s_i^{(\mathrm{obs})}$ with $i = 1, \dots, n^{(\mathrm{obs})}$ such that there is some $j \neq i$ with $\|s_i^{(\mathrm{obs})} - s_j^{(\mathrm{obs})}\| \leq \ell_a$. Note that by setting $\ell_a = \infty\,\mathrm{m}$, all observed weed locations are being targeted.

2.4.2. Autonomous robot

The first treatment tool, denoted as Robot, imitates an autonomous robot that is able to drive directly from one target location to another (see, e.g., Stoll (2020)). The treatment is performed, e.g., by activating a circular mower. In reality, this would result in an oblong treatment area along the driving direction as no cutting occurs in the center of the mower. For simplicity sake, however, we assume a disk-shaped area. After the treatment, the robot continues straight to the next target location, see Figure 3a.

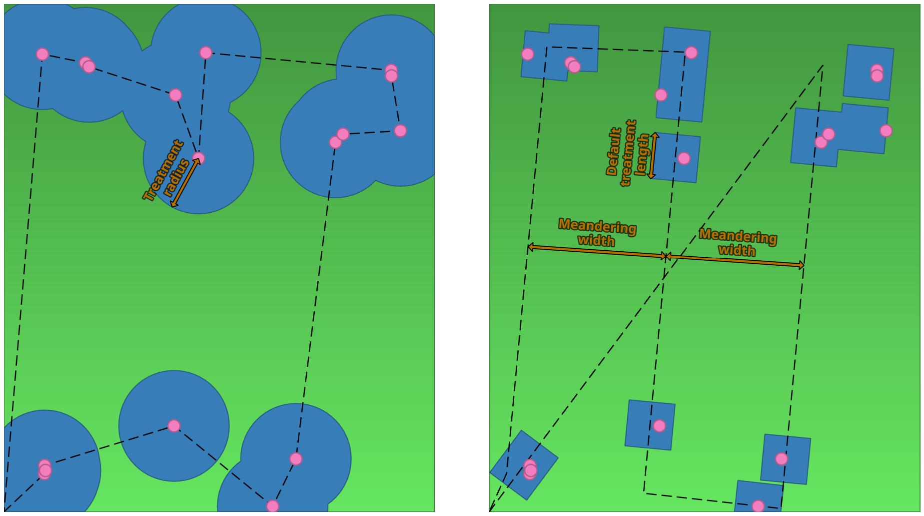

(a) The tool Robot goes directly from one targeted weed location to the next, where a disk-shaped area is treated.

(b) The tool Tractor meanders through the weed infested area such that each targeted weed location is covered by one of its separately controllable treatment sections. The meandering width is thus equal to the width of the attachment. Once the tool crosses a targeted weed location a rectangular area with a given side length along the driving direction is treated. In the illustrated example the treatment tool consists of three sections.

Figure 3: Illustration of the treatment tools Robot (left) and Tractor (right) and their parameters. The treatment tools start at a starting point x_0 (bottom left corner) and drive over the field $\mathcal{W}$ (green area) along a specified route (dashed line). Every time it crosses a targeted weed location (pink dots), a tool-specific area is treated (blue area).

The parameter configuration consists of a treatment radius $r_{\mathrm{t}} > 0\,\mathrm{m}$ (in meter), which specifies the radius of the treated circular area centered at each targeted weed location $s_1^{(\mathrm{tgt})}, \dots, s_{n^{(\mathrm{tgt})}}^{(\mathrm{tgt})}$. For our simulation study we assume that both treatment tools start in a fixed point $x_0 \in \mathcal{W}$, which will also be the point to where they return after they have finished. The route of the treatment tool Robot through the field is determined as the shortest tour through all target locations $s_1^{(\mathrm{tgt})}, \dots, s_{n^{(\mathrm{tgt})}}^{(\mathrm{tgt})}$ and the starting point x_0. This is known as the (Euclidean) traveling salesman problem (Jungnickel, 2008), which we solved using the *OR-Tools* library (Perron and Furnon, 2019). The library produced a not necessarily optimal, but reasonably good route as finding the optimum can be very time consuming. The initial tour was chosen by iteratively adding the weed location with the shortest distance to the previous location beginning with the starting location. It is worth mentioning that as a consequence the route is independent from the treatment radius r_t. More sophisticated strategies would also have been possible where instead of targeting observed weed locations directly, artificial target locations could have been computed depending on r_t such that multiple (observed) weeds were treated simultaneously. However, these lead to extra complexity through additional constraints in practice such as through the imperfect positioning accuracy on the field.

2.4.3. Tractor with attached treatment tool

The second treatment tool, denoted as Tractor, behaves like a tractor with an attached treatment tool (such as the water-hydraulic tool proposed in Stoll (2020) and Martin et al. (2022)) that covers the weed infested area in a winding path. The attachment consists of several separately engageable sections for a site-specific treatment. Each of these sections treat a rectangular area surrounding a targeted weed location, see Figure 3b.

For the precise definition of Tractor, we consider the targeted weed locations $s_1^{(\mathrm{tgt})}, \ldots, s_{n^{(\mathrm{tgt})}}^{(\mathrm{tgt})}$. The weed infested area $M_\mathrm{inf} \subset \mathcal{W}$ is then given by the convex hull (de Berg et al., 2008) of these targeted weed locations, and the primary driving direction $u_\mathrm{d} \in \mathbb{R}^3$ with $\|u_\mathrm{d}\| = 1$ is the direction of the longest side of the (arbitrarily oriented) minimum rectangular bounding box of M_inf (or equivalently of the targeted weed locations). As described in Section 2.4.2 for Robot, the tractor starts in a fixed point x_0. From there it traverses M_inf primarily in parallel line segments determined by the vector u_d. The distance between these line segments is given by the meandering width $w_\mathrm{m} > 0\,\mathrm{m}$ (in meter), which is equivalent to the width of the attached treatment tool. Relative to x_0, the farthest parallel line segment is located such that its distance to the farthest targeted weed location (which is on the boundary of M_inf) is equal to

$w_{\mathrm{m}}/2$. The tractor crosses orthogonally from one line segment to the other such that the infested area M_{inf} is fully covered (i.e., the distance to the boundary of M_{inf} is at most $w_{\mathrm{m}}/2$). The positive turning radius of a real tractor is neglected in this model. After M_{inf} has been traversed, the tractor returns to the starting point x_0. Note that there are two possibilities in which direction to start the traversal, namely u_{d} or $-u_{\mathrm{d}}$. We chose the one that leads to the shortest route. An example route for Tractor can be seen in Figure 3b.

While the tractor moves along the route described above, the $n_{\mathrm{s}} > 0$ individual sections of the attached treatment tool perform a site-specific weed control. For this, the line segment of length w_{m} perpendicular to the driving route is divided into n_{s} parts. These parts correspond to the sections of the treatment tool, which engage (at least) $l_{\mathrm{d}}/2$ in front of and disengage (at least) $l_{\mathrm{d}}/2$ behind a targeted weed location with the default treatment length $l_{\mathrm{d}} > 0\,\mathrm{m}$. This leads to rectangular treatment areas for isolated weeds with side lengths $l_{\mathrm{d}} \times w_{\mathrm{m}}/n_{\mathrm{s}}$. Note, however, that for targeted weed locations close to each other these rectangles can merge. For the simulation study considered in Section 3, we set the meandering width equal to $w_{\mathrm{m}} = 2.5\,\mathrm{m}$ and the default treatment length to $l_{\mathrm{d}} = w_{\mathrm{m}}/n_{\mathrm{s}}$. The number of sections n_{s} will be varied.

3. Results

Now that the entire simulation framework, i.e. the generation of the simulated weed locations and the treatment strategies, has been laid out, we describe the results of our case study.

3.1. General setting

The general setup for the case study is as follows. First, a set of ground truth weed locations $\{s_1^{(\mathrm{gt})}, \ldots, s_{n^{(\mathrm{gt})}}^{(\mathrm{gt})}\}$ and the observed subset $\{s_1^{(\mathrm{obs})}, \ldots, s_{n^{(\mathrm{obs})}}^{(\mathrm{obs})}\}$ with $n^{(\mathrm{gt})} \geq n^{(\mathrm{obs})} \geq 0$ were obtained, either from the experimental datasets (see Sections 2.2.3 and 2.2.2, respectively) or as realizations from stochastic models (see Section 2.3). Recall from Section 2.3.1 that for the simulated datasets, it is possible to vary the mean number of ground truth weed locations relative to the experimentally observed ones in Exp through the intensity factor λ^*. In the following, we consider three different cases, i.e. $\lambda^* \in \{0.5, 1, 2\}$, which are denoted with a subscript '<' (such as $\text{Cal}^1_<$), '0' (such as Cal^1_0), or '>' (such as $\text{Cal}^1_>$), respectively. Moreover, two observations were simulated, where either fewer or more weed locations were observed, just like with the experimental datasets Exp^1 and Exp^2, respectively. Based on the observed weed locations, the $n^{(\text{tgt})} \geq 0$ targeted weed locations $s_1^{(\text{tgt})}, \dots, s_{n^{(\text{tgt})}}^{(\text{tgt})}$ were selected by applying an action threshold $\ell_a \in \{2.5\,\text{m}, 5\,\text{m}, \infty\,\text{m}\}$ eliminating locations with a nearest neighbor distance larger than ℓ_a. Note that for $\ell_a = \infty\,\text{m}$ no locations were thus removed from the observed dataset. In total six different configurations of treatment tools were considered, three for Robot with treatment radius $r_\text{t} \in \{0.2\,\text{m}, 0.4\,\text{m}, 1.25\,\text{m}\}$ and three for Tractor with $n_\text{s} \in \{1, 5, 10\}$ separately controllable sections. Only the targeted weed locations were given to each of these individual strategies, which then produced a set of treated weed locations $\{s_1^{(\text{t})}, \dots, s_{n^{(\text{t})}}^{(\text{t})}\}$ with $n^{(\text{gt})} \geq n^{(\text{t})} \geq n^{(\text{tgt})}$ and a treated subset M_t of the field $\mathcal{W}$. From the definition of the treatment tools, it is clear that every targeted weed location was treated, but there might have been some weed locations near a targeted location that were also treated as 'collateral damage'. The point where the treatment tools started and finished their tour x_0 was set to the lowest point on the left of the field $\mathcal{W}$ and remained fixed for all simulations.

3.2. Performance measures

In order to quantify the performance of a treatment, we computed the following performance measures, which are minimized by an opimal treatment:

(a) the distance d_d driven by the treatment tool (in meter, including the distance from and to the starting point x_0),

(b) the number f_r of remaining weed locations relative to the total number of ground truth plants, where $f_\text{r} = \frac{n^{(\text{gt})} - n^{(\text{t})}}{n^{(\text{gt})}}$,

(c) the maximum density ρ_{r} of remaining weeds in a disk of radius 2 m, where

$$\rho_{\mathrm{r}} = \frac{1}{(2\,\mathrm{m})^2\,\pi}\ \max_{x\in G} \#\left\{ i : \|s_i^{(\mathrm{gt})} - x\| \leq 2\,\mathrm{m} \text{ and } s_i^{(\mathrm{gt})} \notin \{s_1^{(\mathrm{t})}, \ldots, s_{n^{(\mathrm{t})}}^{(\mathrm{t})}\} \right\}$$

and $G \subset \mathcal{W}$ is a square 5 cm-grid of the field $\mathcal{W}$,

(d) the treated area A_{t} relative to the area of the whole field, where $A_{\mathrm{t}} = \frac{|M_{\mathrm{t}}|}{|\mathcal{W}|}$ and $|\cdot|$ denotes the area of a given set, and

(e) the treated area A_{eff} per treated weed (treatment efficiency), where $A_{\mathrm{eff}} = \frac{|M_{\mathrm{t}}|}{n^{(\mathrm{t})}}$ in m².

Note that the radius of the disk for ρ_{r} was chosen as a trade-off between capturing varying local weed concentrations, while still being large enough to cover more than one weed in a cluster.

A great advantage of simulated data drawn from stochastic weed location models, compared to experimental data, is that the virtual treatments can be repeated without changing the considered scenario (in terms of the expected number of weed locations, their spatial distribution, etc.). Through these replications independent samples of the treatment results can be obtained, which leads to great statistical reliability. For this reason, we drew 10 samples from each of the stochastic simulation models. In the following only the mean values obtained for the quality measures (a) – (e) are presented.

3.3. Comparison of treatment strategies

The first question, we want to answer is which strategy is the best one. Unfortunately, since no strategy outperforms its alternatives with respect to all performance measures (a) – (e) in all scenarios, the definition of optimality has to be relaxed. The importance of a performance measure might vary from context to context. We therefore aim to investigate common trade-offs, which might support decision making by identifying strictly inferior strategies.

This problem can be tackled using the so-called *Pareto-optimality* known from multiobjective optimization, see, e.g., Miettinen (2012). In our case a treatment strategy is Pareto-optimal if every other strategy with a better outcome regarding one performance measure would have a worse result regarding another measure. If for a strategy, on the other hand, there is no performance measure where it surpasses the others, this strategy is clearly inferior as the outcome can be improved regardless of what trade-off a decision maker is willing to make. This means more formally that when minimizing the objective functions (in our case performance measures) $f_1, \dots, f_k : \mathbb{R}^d \supset D \to \mathbb{R}$ for some integers d, k, the decision vector (in our case strategy) $x^* \in D$ is Pareto-optimal if there is no other vector $x \in D$ with $f_i(x) \leq f_i(x^*)$ for all $i = 1, \dots, k$ and $f_j(x) < f_j(x^*)$ for at least one index j (Miettinen, 2012).

In the following, we consider only two performance measures at a time as more would result in practically all strategies being Pareto-optimal. Moreover, visualizing three or more measures simultaneously is much harder. The results of each treatment strategy with respect to the considered performance measures are shown in a scatter plot. The criterion for the Pareto-optimality is illustrated as a line—the so-called Pareto frontier—that separates the Pareto-optimal strategies on the line from the inferior ones on the top right of the line.

In order to reduce complexity, we only consider the datasets corresponding to the early observation date (of September 3rd, 2019) and set the intensity factor $\lambda^* = 1$ for the simulated data (i.e. $\textsc{Exp}^1$, $\textsc{Cal}_0^1$, $\textsc{Hom}_0^1$, ...) when investigating the Pareto-optimality of the strategies. In Figure 4 the driving distance d_d versus the treated area per treated weed A_eff are shown. Here it turned out that the Pareto-optimal strategies were those with the smallest individual treatment area (i.e., the smallest considered treatment radius $r_\mathrm{t} = 0.2\,\mathrm{m}$ for Robot, or the maximum number of sections $n_\mathrm{s} = 10$ for Tractor). Furthermore, compared to Robot, Tractor drove a much larger distance, but treated a slightly smaller area per weed. The scenario $\textsc{Cen}_0^1$ was the only one, however, where there was almost no difference between the two tools. Note that the driving distance d_d does not change when varying the parameters r_t or n_s as the route of a treatment tool depends only on its type (Robot or Tractor) and the targeted weed

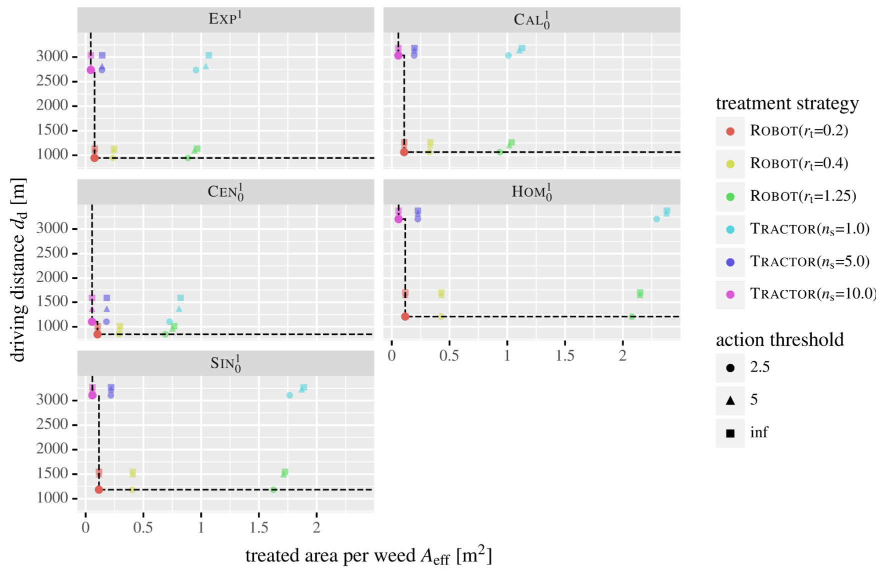

Figure 4: Scatter plots together with the corresponding Pareto frontiers of the driving distance d_d and the treated area per treated weed A_eff for various treatment strategies and infestation scenarios.

locations, which are a function of the action threshold ℓ_a. Regarding ℓ_a, lower values, which reduced the number of targeted weed locations, led to lower driving distances as might have been expected. Consequently, $\ell_a = 2.5\,\mathrm{m}$ produced the lowest values of d_d.

In Figure 5, the trade-off between the fraction of remaining weeds f_r and the treated area per treated weed A_eff are illustrated. Here, apparently almost all strategies were Pareto-optimal which points to the strong antagonistic dependency between the treated area and the number of remaining weeds. Two groups of treatment strategies could be made out: the first one comprised the strategies with the largest individual treatment area (i.e. ROBOT with $r_\mathrm{t} = 1.25\,\mathrm{m}$ and TRACTOR with $n_\mathrm{s} = 1$) and treated significantly more area per weed at the benefit of hitting a larger percentage of the whole weed population compared to the remaining strategies in the second group. Interestingly enough, no big difference between the two tools was visible, only between their configuration.

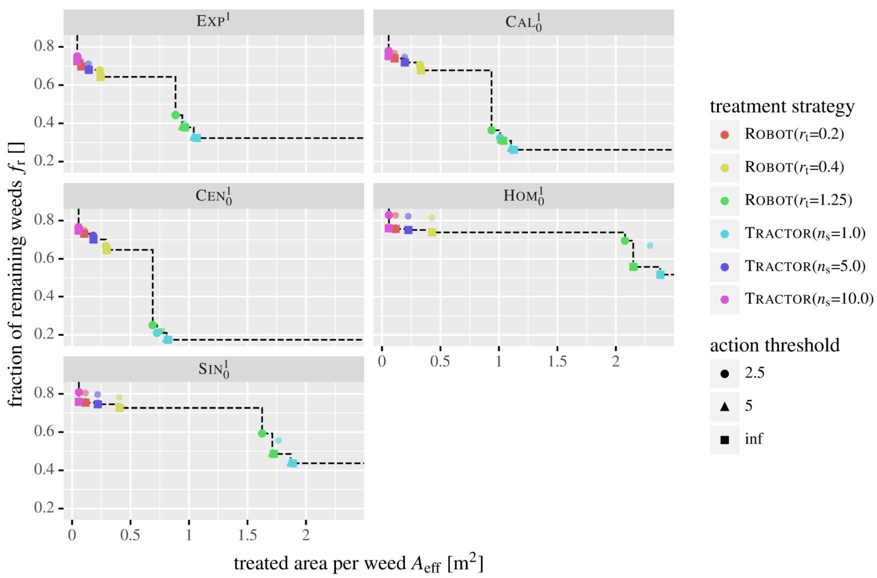

Figure 5: Scatter plots together with the corresponding Pareto frontiers of the fraction of remaining weeds f_r and the treated area per treated weed A_eff for various treatment strategies and infestation scenarios.

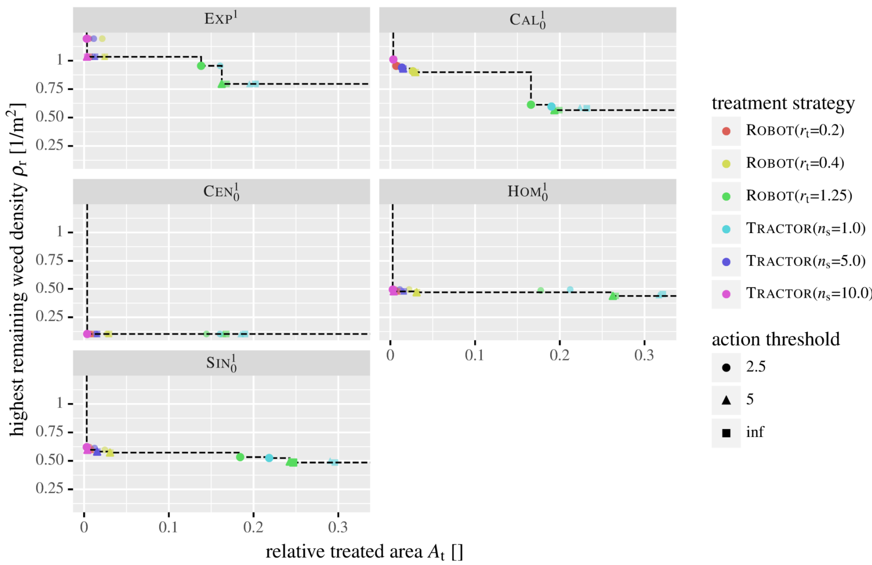

Figure 6: Scatter plots together with the corresponding Pareto frontiers of the relative treated area A_t and the highest remaining weed density ρ_r for various treatment strategies and infestation scenarios.

The differences between $\ell_a = \infty\,\mathrm{m}$ and $\ell_a = 5\,\mathrm{m}$ were negligible but there was visible contrast compared to $\ell_a = 2.5\,\mathrm{m}$. As opposed to Figure 4, the latter action threshold usually performed the worst.

Another important pair of competing performance measures is the relative treated area A_t and the highest remaining weed density ρ_r. The corresponding results are shown in Figure 6. Here, TRACTOR with $n_s = 10$ was Pareto-optimal in all scenarios. The other strategies, ROBOT with $r_t = 1.25\,\mathrm{m}$ most of all, were in some cases able to decrease the highest remaining weed density ρ_r. The biggest difference was achieved for EXP^1 and CAL$_0^1$, whereas for the remaining scenarios little or no improvement was obtained with respect to ρ_r. Generally speaking a smaller value for the action threshold ℓ_a yielded better results.

3.4. Influence of action threshold on treatment performance

Another question that we want to investigate is how the action threshold ℓ_a affected the treatment performance. Like above, we considered only the early

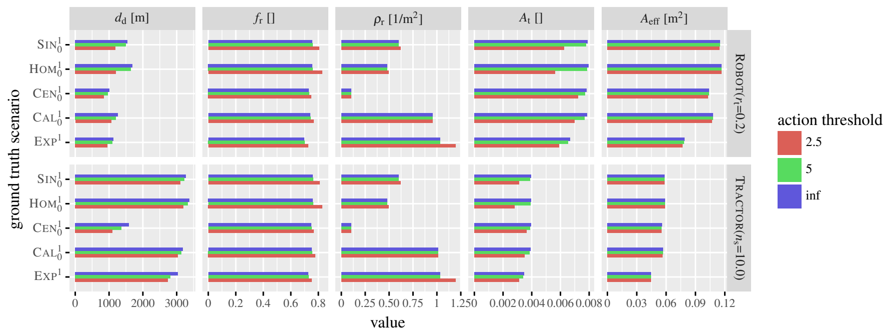

Figure 7: Influence of the action threshold ℓ_a on the considered performance measures for various infestation scenarios.

observation date and set the intensity factor $\lambda^* = 1$, but restricted the tool configurations to the ones described in Stoll (2020) and Martin et al. (2022), namely ROBOT with treatment radius $r_\mathrm{t} = 0.2\,\mathrm{m}$ and TRACTOR with section count $n_\mathrm{s} = 10$. The results are shown in Figure 7. By setting the action threshold $\ell_a = 5\,\mathrm{m}$, small reductions in the driving distance d_d compared to the baseline $\ell_a = \infty\,\mathrm{m}$ could be observed, but practically no change for the other metrics. So a small net win could be achieved. For the smallest action threshold, $\ell_a = 2.5\,\mathrm{m}$, more weeds remained untreated. It is also noteworthy, that especially for the TRACTOR the treated area per weed A_eff was mostly unaffected by the action threshold.

3.5. Effect of observation date

In order to study the effect of the observation date, we set the intensity factor $\lambda^* = 1$ and considered only the action threshold $\ell_a = \infty\,\mathrm{m}$ for the two treatment configurations, ROBOT with treatment radius $r_\mathrm{t} = 0.2\,\mathrm{m}$ and TRACTOR with section count $n_\mathrm{s} = 10$. The resulting metrics are visualized in Figure 8. Between the two observation dates, large differences in almost all considered metrics could be observed. The datasets corresponding to September 10th, 2019, where more weeds were observed ($n_2^{(\mathrm{EXP})} = 1792$ versus $n_1^{(\mathrm{EXP})} = 550$ for September 3rd, 2019) had a much better treatment result. Notable outliers

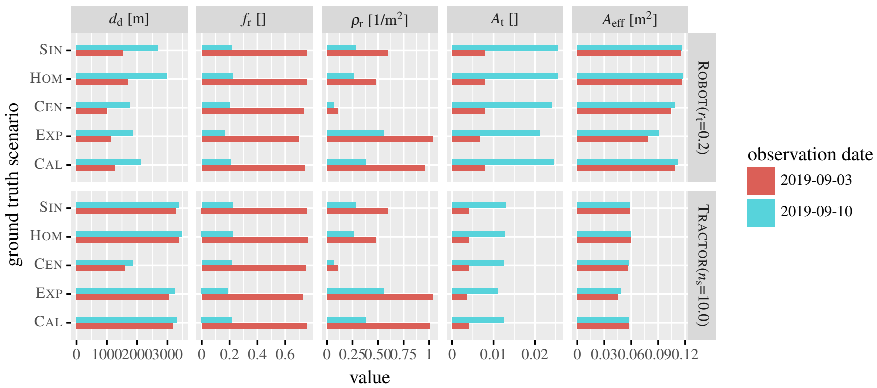

Figure 8: Influence of the observation date on the considered performance measures for various infestation scenarios.

were, however, the driving distance d_d for TRACTOR, and the treated area per treated weed for both tool, which changed only marginally. It should be kept in mind that the individual treatment areas differ and except for the driving distance d_d a comparison between the TRACTOR and the ROBOT would not be reasonable. All ground truth scenarios produced the same qualitative behavior of the considered performance measures.

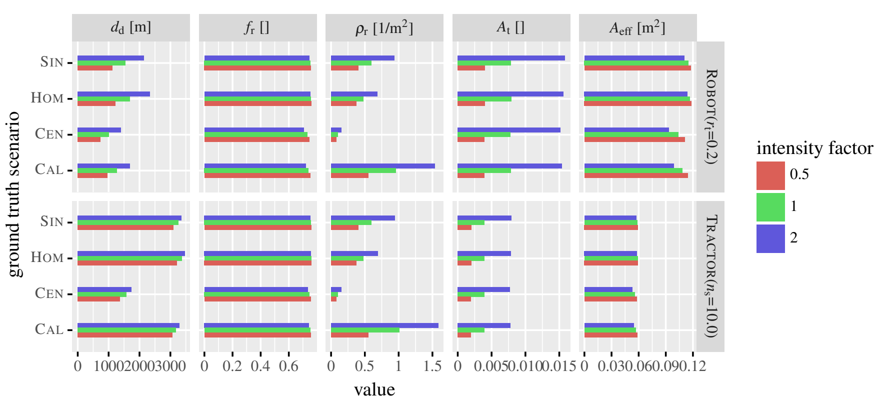

Figure 9: Influence of the intensity factor λ^* on the considered performance measures for various infestation scenarios.

3.6. Influence of the intensity factor λ^*

In Figure 9 the dependency of the intensity factor λ^* on the considered performance measures is shown, which illuminates the question how the treatment strategies performed when the number of (ground truth) weed locations varies. Of course, only simulated datasets could be considered. As before, we focused on Robot with treatment radius $r_\mathrm{t} = 0.2\,\mathrm{m}$ and Tractor with section count $n_\mathrm{s} = 10$, and we ignored the action threshold, i.e. $\ell_a = \infty\,\mathrm{m}$. Apparently, the driving distance d_d did not scale linearly with the mean number of weeds on the field. So, a low number of weeds lead to much larger distances (and therefore costs) per weed compared to cases where this number was already quite high. The Tractor tool was less affected by the increase in d_d than Robot. For the treated area per treated weed A_eff, it was also observable that A_eff decreased with increased intensity factor λ^*. So the efficiency rose as more and more unobserved weeds stood near targeted weed locations and were treated.

4. Discussion

When it comes to evaluating the different treatment strategies, the Figures 4, 5, and 6 showed how difficult it is to give general suggestions. However, certain strategies can be eliminated when considering only a few performance measures that are the most important ones in the current context. Especially when keeping the results shown in Figure 9 in mind, the Robot strategy generally performed quite well. Only with an increasing number of weeds Tractor obtained better results in particular if the weeds are clustered. For real world applications, additional constraints come into play, such as a limited fuel tank of an automated robot resulting in a shorter reach, or the varying costs of personal. The presented approach could also be extended to serve as a framework to investigate such more intricate questions that arise from planning a treatment tool to managing weeds on a large number of fields with a limited number of available treatment tools.

It turns out that by employing a (finite) action threshold, the driving distance d_{d} could be reduced. For the smallest value $\ell_a = 2.5\,\mathrm{m}$, however, a noticeable degradation of the treatment performance could be observed, see Figure 7. The value $\ell_a = 5\,\mathrm{m}$ might be used as a compromise between a smaller driving distance and practically the same treatment performance compared to omitting the action threshold. For future research it might be worthwhile to implement more sophisticated methods to determine the set of targeted weed locations. For example, using point processes would allow for a model based prediction of unobserved weed locations. With this, additional virtual locations could be created in areas where unobserved weeds are assumed. Another approach would be to predict unobserved weed locations using machine learning techniques. However, both approaches would require additional experimental data to build accurate models.

In Figure 8, a strong dependency of the treatment success on the number of observed weeds can be seen. The importance of this relationship is only increased, when keeping in mind that only for simplicity's sake we assumed the experimental locations Exp comprised the whole weed population. In reality, however, there might have been much more weeds on the field that were not captured during the two drone mappings. This partial information problem affects both online and offline treatment methods that either identify the targeted weed locations while the treatment tool is traversing the field (e.g., through attached cameras and real-time image processing), or decouple the acquisition of the image data/detection of weed locations with their treatment, respectively. For offline methods the weed population might change between the observation and the treatment (e.g., drone imaging in fall where the weeds are easiest to identify versus treatment in spring where the plants are most vulnerable). This might lead to an even larger amount of discrepancy between the observed and the ground truth weed locations. However, an in-depth analysis of this phenomenon is beyond the scope of the present paper. For online methods, on the other hand, it is not possible to combine information from several observations, e.g., from different days. These additional data sources could be necessary to obtain an accurate survey of the whole weed population. Hybrid approaches that combine the benefits of online and offline techniques could be a solution.

As said in the introduction, the findings might also be applicable for other treatment tools such as site-specific herbicide sprayers, or spot-spraying through UAVs. For this, it is necessary to account for tool-specific requirements. So while overlapping treatment areas are of no importance for a circular mower, they have to be eliminated when considering sprayers. However, this can be achieved by preprocessing the application maps prior to the treatment, or by tracking the treated areas during operation.

Various extensions to the proposed framework are possible. First off, more sophisticated point-process models like (Poissonian) cluster point processes for the ground truth weed location and dependent thinning for the observed subset could be investigated to obtain more variability and more complex interdependencies in the resulting point patterns. Ideally, a thorough stochastic model of real weed locations could be built. However, this would require more experimental data from different fields at several time points. Another point for enhancement would be the addition of further treatment tools and configurations.

5. Conclusions

In the present paper, a simulation framework for comparing weed control tools under varying infestation scenarios was developed. This allowed a novel view on the pros and cons of very different kinds of treatment tools such as an autonomous robot, or a tractor-attached implement, where previously only one of these was studied in isolation. An additional unique feature was the focus on grassland sites and the accounting of partial observation of the weed population, which we found to have a strong influence on the overall treatment success. Furthermore, it turned out that the considered autonomous robot performed quite well in most scenarios, but the results for the tractor-mounted implement improved as the number of weeds increased. For both treatment tools slight improvements could be achieved when isolated weeds were not treated.

Acknowledgements

The project is supported by funds of the Federal Ministry of Food and Agriculture (BMEL) based on a decision of the Parliament of the Federal Republic of Germany via the Federal Office for Agriculture and Food (BLE) under the innovation support program.

CRediT authorship contribution statement

Lukas Petrich: Methodology, Software, Writing - Original Draft. **Georg Lohrmann**: Data Curation. **Fabio Martin**: Data Curation. **Albert Stoll**: Funding acquisition, Project administration, Writing - Review & Editing. **Volker Schmidt**: Supervision, Conceptualization, Writing - Review & Editing.

References

Andrew, I.K., Storkey, J., 2017. Using simulation models to investigate the cumulative effects of sowing rate, sowing date and cultivar choice on weed competition. Crop Protection 95, 109–115.

de Berg, M., Cheong, O., van Kreveld, M., Overmars, M., 2008. Computational Geometry. 3rd ed., Springer.

Chiu, S.N., Stoyan, D., Kendall, W.S., Mecke, J., 2013. Stochastic Geometry and Its Applications. 3rd ed., J. Wiley & Sons.

Christensen, S., Søgaard, H.T., Kudsk, P., Nørremark, M., Lund, I., Nadimi, E.S., Jørgensen, R., 2009. Site-specific weed control technologies. Weed Research 49, 233–241.

Colbach, N., Colas, F., Cordeau, S., Maillot, T., Queyrel, W., Villerd, J., Moreau, D., 2021. The FLORSYS crop-weed canopy model, a tool to investigate and promote agroecological weed management. Field Crops Research 261, 108006.

Colbach, N., Cordeau, S., Garrido, A., Granger, S., Laughlin, D., Ricci, B., Thomson, F., Messéan, A., 2018. Landsharing vs landsparing: How to reconcile crop production and biodiversity? A simulation study focusing on weed impacts. Agriculture, Ecosystems & Environment 251, 203–217.

Freckleton, R.P., Stephens, P.A., 2009. Predictive models of weed population dynamics. Weed Research 49, 225–232.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning. 2nd ed., Springer.

Holst, N., Rasmussen, I.A., Bastiaans, L., 2007. Field weed population dynamics: a review of model approaches and applications. Weed Research 47, 1–14.

Jähne, B., 2005. Digital Image Processing. Springer.

Jungnickel, D., 2008. Graphs, Networks and Algorithms. Springer.

Loader, C., 1999. Local Regression and Likelihood. Springer.

Machleb, J., Peteinatos, G.G., Kollenda, B.L., Andújar, D., Gerhards, R., 2020. Sensor-based mechanical weed control: Present state and prospects. Computers and Electronics in Agriculture 176, 105638.

Maini, P., Gonultas, B.M., Isler, V., 2022. Online coverage planning for an autonomous weed mowing robot with curvature constraints. IEEE Robotics and Automation Letters 7, 5445–5452.

Martin, F., Lohrmann, G., Stoll, A., 2022. Selective weed control in grassland using high-pressure water jets, in: VDI Wissensforum GmbH (Ed.), LAND.TECHNIK 2022: The Forum for Agricultural Engineering Innovations, Düsseldorf. pp. 105–110.

MathWorks, 2021. Image processing toolbox: Reference (Matlab R2021b). URL: `https://mathworks.com/help/pdf_doc/images/images_ref.pdf`. (accessed 2021-11-08).

Miettinen, K., 2012. Nonlinear Multiobjective Optimization. Springer.

Perron, L., Furnon, V., 2019. OR-Tools. URL: https://developers.google.com/optimization/.

Petrich, L., Lohrmann, G., Neumann, M., Martin, F., Frey, A., Stoll, A., Schmidt, V., 2020. Detection of colchicum autumnale in drone images, using a machine-learning approach. Precision Agriculture 21, 1291–1303.

Pointurier, O., Gibot-Leclerc, S., Moreau, D., Colbach, N., 2021. How to pit weeds against parasitic plants. A simulation study with Phelipanche ramosa in arable cropping systems. European Journal of Agronomy 130, 126368.

Schellberg, J., Hill, M.J., Gerhards, R., Rothmund, M., Braun, M., 2008. Precision agriculture on grassland: applications, perspectives and constraints. European Journal of Agronomy 29, 59–71.

Somerville, G.J., Sønderskov, M., Mathiassen, S.K., Metcalfe, H., 2020. Spatial modelling of within-field weed populations; a review. Agronomy 10, 1044.

Stoll, A., 2020. Selektive, nicht-chemische Bekämpfung von Giftpflanzen in extensiven Grünlandbeständen (SELBEX), in: Bundesanstalt für Landwirtschaft und Ernährung (Ed.), Tagungsband Innovationstage 20./21.Oktober 2020, Bonn. pp. 191–193.

Wegener, J.K., 2020. Gezielter und flexibler – Trends in der Pflanzenschutztechnik, in: Frerichs, L. (Ed.), Jahrbuch Agrartechnik 2019. Institut für mobile Maschinen und Nutzfahrzeug, Braunschweig, pp. 1–7.