Field Crops Research 337 (2026) 110236

Contents lists available at ScienceDirect

Field Crops Research

journal homepage: www.elsevier.com/locate/fcr

Can convolutional neural networks support agronomic analysis of cereal–legume canopy cover dynamics?

David Kottelenberg [a,*], Lammert Bastiaans [a], Rick van Essen [b], Gert Kootstra [b], Jacob C. Douma [a]

[a] Centre for Crop Systems Analysis, Wageningen University, Bornsesteeg 48, Wageningen 6708 PE, the Netherlands
[b] Agricultural Biosystems Engineering, Wageningen University, Droevendaalsesteeg 1, Wageningen 6708 PB, the Netherlands

ARTICLE INFO

Keywords:
Intercropping
Cereal
Legume
Canopy cover
Lotka-Volterra
Competition
Convolutional neural network

ABSTRACT

Context: Understanding crop–crop and crop–weed interactions is essential for designing overyielding and weed-suppressive intercropping systems. Measurements of canopy cover over time can provide insights into these interactions, but are labour-intensive to collect. Machine learning methods, specifically convolutional neural networks (CNNs), could automatically analyse cover of individual species from canopy cover photos, yet the quality of the cover assessment that is needed to study species interaction remains unclear.
Objective: This study aimed to quantify competitive dynamics in cereal–faba bean intercrops based on canopy cover and assess CNN performance required for reliable analysis.
Methods: We collected RGB images from cereal–faba bean intercrops varying in cereal species (barley, rye, triticale, wheat), triticale:faba bean mixing ratios (1:1, 1:3, 3:1), and spatial design (row or mixed). Canopy cover was manually annotated for 397 images, identifying cereal, faba bean, and weed classes. Four CNN models of varying complexity were trained, the simplest of which were used off-the-shelf. We compared qualitative patterns and Lotka–Volterra competition parameters between ground-truth and CNN-segmented data.
Results: Ground-truth data revealed that rye was the most competitive cereal, and wheat the least, reflected in Lotka-Volterra intrinsic growth rate parameters. Separating cereals and legumes into rows and reducing the cereal proportion in intercrops decreased cereal competitiveness relative to faba bean, resulting in more even canopy cover and more symmetrical competition parameters between species. All CNN models achieved high accuracy (Intersection over Union (IoU) = 0.900–0.926). While CNN-based segmentations matched ground-truth patterns visually, only our most complex model came close to the ground-truth parameter estimates, whereas the other three produced values too uncertain or biased to support the same conclusions.
Conclusion: We conclude that moderate-complexity CNN models are sufficient to qualitatively interpret cover trends, but for more refined ecological analysis more complex CNNs are needed. Sensitivity analysis could aid in quantifying the performance needed before training such a complex CNN.

1. Introduction

Modern industrial agriculture is characterized by input-intensive monocropping systems. These systems have created substantial challenges, including landscape simplification, soil erosion, and biodiversity loss due to synthetic fertilizer and pesticide runoff (Crews et al., 2018). To sustainably meet the demands of a growing population, modern agriculture must reduce inputs, adapt to increasingly variable climate conditions, and prioritize soil health, biodiversity, and ecosystem services (Gaba et al., 2014). This reduction can roughly be achieved by two means. First, by making use of precision agriculture allowing spatial and temporal adjustment of inputs based on the crop's needs, hence lowering inputs. Second, by making use of ecological intensification strategies (Mortensen and Smith, 2020). Intercropping, the simultaneous cultivation of two species on the same land, is such a promising practice as these systems generally have higher yields (Li et al., 2020; Bedoussac et al., 2014), improved resource-use efficiency (Tang et al., 2020; Jensen et al., 2020; Chen et al., 2018; Franco et al., 2018), reduced pest pressure (Li et al., 2021), and increased weed suppression (Gu et al., 2021). Interestingly, precision agriculture can reinforce the

* Corresponding author.
E-mail address: david.kottelenberg@wur.nl (D. Kottelenberg).

https://doi.org/10.1016/j.fcr.2025.110236
Received 19 September 2025; Received in revised form 10 November 2025; Accepted 14 November 2025
Available online 21 November 2025

implementation of ecological intensification strategies by providing frequent and low cost monitoring of crop health, including the detection of pests and weeds. However, it often relies on recognition and classification algorithms, typically employing machine learning methods like deep learning, to provide efficient and detailed insights. Despite this potential, the adoption of these complex methods in both farming and research remains limited (Coulibaly et al., 2022).

Diversified cropping systems, like intercropping, are inherently complex and demand advanced management strategies. Effective management of these systems requires a deep understanding of the interactions between crop species. Cereal-legume intercrops, in particular, have been widely studied due to their complementary ecological niches and synergistic interactions: cereals are strong competitors enhancing the intercrop's weed suppression (Gu et al., 2022), while faba bean has the ability to fix atmospheric nitrogen, increasing nitrogen availability in the soil (Bybee-Finley and Ryan, 2018). However, as cereals are generally stronger competitors than legumes (Gu et al., 2021), this competitive imbalance often results in the suppression of the legume, potentially leading to significant yield losses (Gu et al., 2025; Cheriere et al., 2020). To optimise the performance of such systems, it is important to understand competition dynamics between those crops in time and space. Precision agriculture tools can support both species-specific monitoring for management decisions and research aimed at quantifying competition dynamics to identify optimal species combinations and spatial arrangements. Detailed agroecological analyses can help gain insights into plant interactions that shape these systems and provide useful clues for improved systems design.

Canopy cover is a commonly used agroecological metric for assessing crop performance at different points in time, as it provides a straightforward way to estimate the proportion of ground covered by plants through overhead images (Patrignani and Ochsner, 2015). However, in intercropping and weed-infested systems, it is essential to distinguish between the contributions of different crop species and weeds to the overall canopy cover, as these components differentially affect crop performance, and ultimately yield. Techniques which only estimate total canopy cover fall short in capturing this complexity. Species-specific cover estimates, while more informative, are labour-intensive to obtain and prone to high uncertainty. While manual labelling of crops and weeds in images can yield detailed insights, it is both time-consuming and impractical for large datasets. Furthermore, competitive dynamics can change throughout the growing season, making it essential to track individual species at short time intervals, especially during the early growth stages (Trinder et al., 2013). This increases the volume of data needing labelling, which makes the process of identifying and distinguishing plant species, known as segmentation, challenging and laborious.

Convolutional neural network (CNN) models have been successfully used to identify and segment crop species in images (Koshelev et al., 2023; Ercolini et al., 2022). CNNs learn to recognise species through iterative training on labelled datasets (Gu et al., 2018). By providing the model with manually annotated images, it learns to associate certain characteristics in the image, so-called features, with their respective classes (e.g., crop species and weeds). With sufficient training data, a CNN can accurately predict the category of each pixel in a new image, effectively segmenting the image into different components. This capability makes CNNs a powerful tool for analysing complex cropping systems, enabling efficient and accurate canopy cover assessment at species (group) level.

Gathering input data for neural network training is laboursome, as it requires large amounts of images that need to be manually annotated (i.e., pixels need to be hand-classified to the species of interest). The performance of a CNN can be improved by presenting the algorithm with modified versions of already annotated images, so-called data augmentation. Such modifications include rotated, flipped, mirrored, and colour-shifted images (Su et al., 2021). Additionally, model performance can be improved with data synthetization techniques, whereby new images are generated by, for example, combining different parts of existing data (Sapkota et al., 2022; Fawakherji et al., 2021; Van Essen et al., 2021).

Many CNN models have been published showing performance improvements on the detection of crops and/or weeds in agricultural fields (Mesías-Ruiz et al., 2024; Purohit et al., 2024; Koshelev et al., 2023; Ercolini et al., 2022; Champ et al., 2020; Revanasiddappa et al., 2020; Ma et al., 2019; Abdalla et al., 2019; Kounalakis et al., 2017). However, to our knowledge, there has been no analysis to see whether existing models perform well enough in agroecological contexts and whether their quality is sufficient to provide meaningful insights into competition dynamics in agroecological systems. It seems there is a disconnect between artificial intelligence (AI) experts, who develop increasingly more sophisticated deep-learning models, and agroecology experts, who primarily rely on traditional models to collect and analyse data.

This study addresses two goals. First, from an agroecological perspective, we aim to better understand competition dynamics in cereal–legume intercrops. We do so by applying the Lotka–Volterra competition model to species-specific canopy cover data derived from manually segmented (ground-truth) RGB images of various cereal-faba bean intercropping systems, differing in species combinations, mixing ratios, and planting designs. This allows us to assess how these intercrop system design aspects affect the balance of cereal and faba bean canopy development. Second, from a methodological perspective, we evaluate whether convolutional neural networks (CNNs) can generate sufficiently accurate canopy segmentation and provide similar insights as the analysis based on ground-truth images. We train four CNN models of varying methodological complexity and compare their segmentation quality against the ground-truth data to assess whether key ecological inferences, such as those based on parameter fits from the Lotka–Volterra model, remain robust. In this way, we aim to clarify both the agroecological insights that species-specific canopy cover data can provide and the level of CNN-model performance needed to obtain them.

2. Materials and methods

2.1. Field experimental setup

Canopy cover data were collected from two field experiments (2022 and 2023) described in Kottelenberg et al. (2025a). Relevant experimental details are summarised below.

In 2022, barley (*Hordeum vulgare* cv. Irina; treatment code 'B'), rye (*Secale cereale* cv. Boyce; 'R'), triticale (× *Triticosecale* cv. Santos; 'T'), and wheat (*Triticum aestivum* cv. Quintus; 'W') were grown as cereals in sole crop or in alternate row intercrops of cereal and faba bean (*Vicia faba* cv. Cartouche; 'F'). The experiment also included a sole crop faba bean. Row distances were 12.5 cm and plot sizes were 3 × 9 m. All plots were replicated four times and arranged according to a randomized complete block design. Sowing was done on 19 April for both crops with the Macon trial field sowing machine. On 12 May the fields were treated with 2 L ha^{-1} of the herbicide Basagran (BASF, 480 g L^{-1} bentazon). At the time, the primary objective was to analyse crop-crop interactions under weed-free conditions.

In 2023, the experimental design was modified to address different research objectives and include weed-infested treatments. Triticale (cv. Mazur; different from 2022 due to availability) and faba bean (cv. Cartouche) were grown in sole crops ('T' and 'F'). Additionally, triticale and faba bean were grown in 1:1 row intercrops ('1 T:1 F', the same intercrop system as in 2022), 3:1 and 1:3 row intercrops ('3 T:1 F' and '1 T:3 F'), and 1:1 proportion within-row mixed intercrops ('1 T:1F-M'). Row distances were 12.5 cm and plot sizes were 3 × 11 m. All plots were replicated five times and arranged according to a randomized complete block design. Sowing was done on 2 March for both crops with the Macon trial field sowing machine. Canopy cover images from non-herbicide-treated plots were used to include weed presence in the analysis, while herbicide-treated images were excluded.

Canopy cover images were taken with a Canon Powershot SX200 IS camera fixed on a vertical metal pole, positioned 100 cm above a 100 × 75 cm rectangular frame. The frame defined the photographed area and was placed horizontally at canopy height, perpendicular to the crop rows. This setup ensured consistent top-down images, which were later cropped to the frame boundaries. One picture per crop system was taken weekly from 4 May 2022–14 June 2022 and from 5 April 2023–31 May 2023. Furthermore, only half of the sole crop barley, rye, triticale, and wheat images were chosen, to not inflate the dataset with too much sole crop cereal data, as the focus was mainly on segmenting intercrop images. This resulted in a total of 397 images used for analysis.

Total canopy cover was determined by detection of green pixels, assuming the proportion of green to total number of pixels is the total canopy cover of the crops and weeds, using a canopy cover analysis tool (Kottelenberg, 2024). The tool separates pixels based on user-defined colour thresholds. Green pixels were determined using user-defined hue, saturation, and value (HSV) ranges: hue between 36 and 86 (colour, from a full range between 0°–179°), saturation between 50 and 255 (colour intensity, from a full range of 0–255), and value between 30 and 255 (colour brightness, from a full range of 0–255). Total canopy cover was then calculated as the proportion of green pixels relative to the total number of pixels in the image. To obtain species-specific canopy cover, images were manually segmented by drawing polygons around visible plant structures (leaves, stems, flowers, fruits) of faba bean and weeds only, classifying the enclosed pixels accordingly (Fig. 1). Cereal plants were not explicitly annotated to simplify model training, as their canopy cover could later be derived indirectly from the remaining classes. Specifically, faba bean and weed cover were each computed as the proportion of their annotated pixels relative to the total number of image pixels, and cereal canopy cover was estimated as the residual, by subtracting the faba bean and weed cover from the total canopy cover.

2.2. Fitting the Lotka-Volterra competition model on segmented canopy cover data

Annotated pictures gave ground-truth data on component canopy cover of cereal, faba bean, and weed. Due to the low amount of weed cover and the focus on cereal and faba bean competition, weed cover was excluded from the quantitative analysis. The Lotka-Volterra model of interspecific competition was fitted to the ground truth cover data of cereal and faba bean per intercrop treatment (1B:1 F, 1 R:1 F, 1 T:1 F, 1 W:1 F in 2022; 1 T:1 F, 1 T:1F-M, 3 T:1 F, 1 T:3 F in 2023):

$$\frac{dC}{dt} = r_c C(1 - C - a_{fc} F),$$

$$\frac{dF}{dt} = r_f F(1 - F - a_{cf} C) \tag{1}$$

where C and F are the proportion of cereal and faba bean cover (dimensionless); $\frac{dC}{dt}$ and $\frac{dF}{dt}$ are the change of cereal and faba bean proportion cover over time, respectively (in day^{-1}); r_c and r_f are the intrinsic growth rates of cereal and faba bean cover (in day^{-1}); and a_{fc} and a_{cf} are the interspecific competition coefficients of the effect of faba bean cover on cereal cover and vice versa, respectively (dimensionless). The intrinsic growth rate determines how quickly a species develops canopy cover, while the competition coefficients reflect the competitive pressure exerted per unit cover of one species relative to the other.

In our analysis, we estimated species-specific parameters. For 2022, we estimated separate r_c parameters for each cereal species (r_b for barley, r_r for rye, r_t for triticale, r_w for wheat) and their corresponding competition coefficients (e.g., a_{bf} and a_{fb} for barley-faba bean interactions), while we estimated a single r_f parameter across all treatments since faba bean is the only legume species. For 2023, with triticale as the only cereal species, we estimated a single r_t across treatments. However, to capture effects of spatial arrangement, we estimated treatment-specific competition coefficients (e.g., $a_{tf,1t:1f}$, $a_{tf,1t:3f}$, $a_{tf,3t:1f}$, $a_{tf,1t:1f-m}$, similar for a_{ft} parameters for the different treatments) for the different intercrop configurations.

Comparing parameter values between species and treatments provides insights into expected canopy development. A species with both a high intrinsic growth rate and strong competitive effect on its neighbour can suppress the other species to near-zero cover (competitive exclusion). Conversely, a species may dominate early if it grows faster, but be overtaken later if its competitor exerts stronger competitive pressure per unit of cover. The parameter combinations illustrated in Fig. 2 represent a range of possible cover trajectories over time.

Parameter values for the Lotka-Volterra model were estimated for mean values across repetitions using R 4.4.2 (R Core Team, 2024) with the `optim` function. The model was solved numerically using the `ode` function from the deSolve package (Soetaert et al., 2010). Initial canopy cover values for each crop species were set equal to the observed or predicted mean cover at the earliest time point in the dataset. To fit the model to the data, parameters were optimized by minimizing the negative log-likelihood of a Beta distribution, which is well-suited for values constrained between 0 and 1, such as canopy cover proportions (Douma and Weedon, 2019):

$$f(x) = \frac{\Gamma\alpha + \beta}{\Gamma\alpha\Gamma\beta} x^{\alpha-1}(1-x)^{\beta-1} \tag{2}$$

where Γ is the Gamma function, and α and β are the shape parameters of the probability distribution, with mean $\frac{\alpha}{\alpha+\beta}$ and variance $\frac{\alpha}{(\alpha+\beta)^2(\alpha+\beta+1)}$. In our implementation, the Beta distribution was fitted to the canopy cover

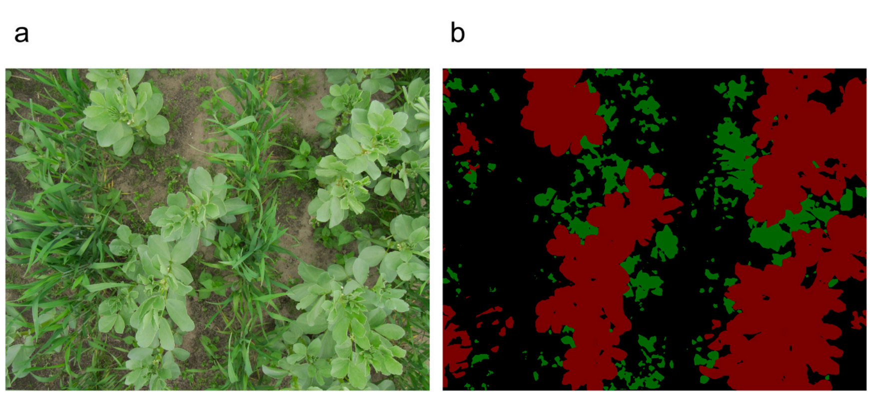

Fig. 1. Cereal-legume intercropping image segmentation. a: Cereal-legume intercropping canopy cover picture. b: Annotation of the picture in 'a' where faba bean and weed are segmented in separate categories. Red is faba bean, green is weed, black is background composed of cereal and soil.

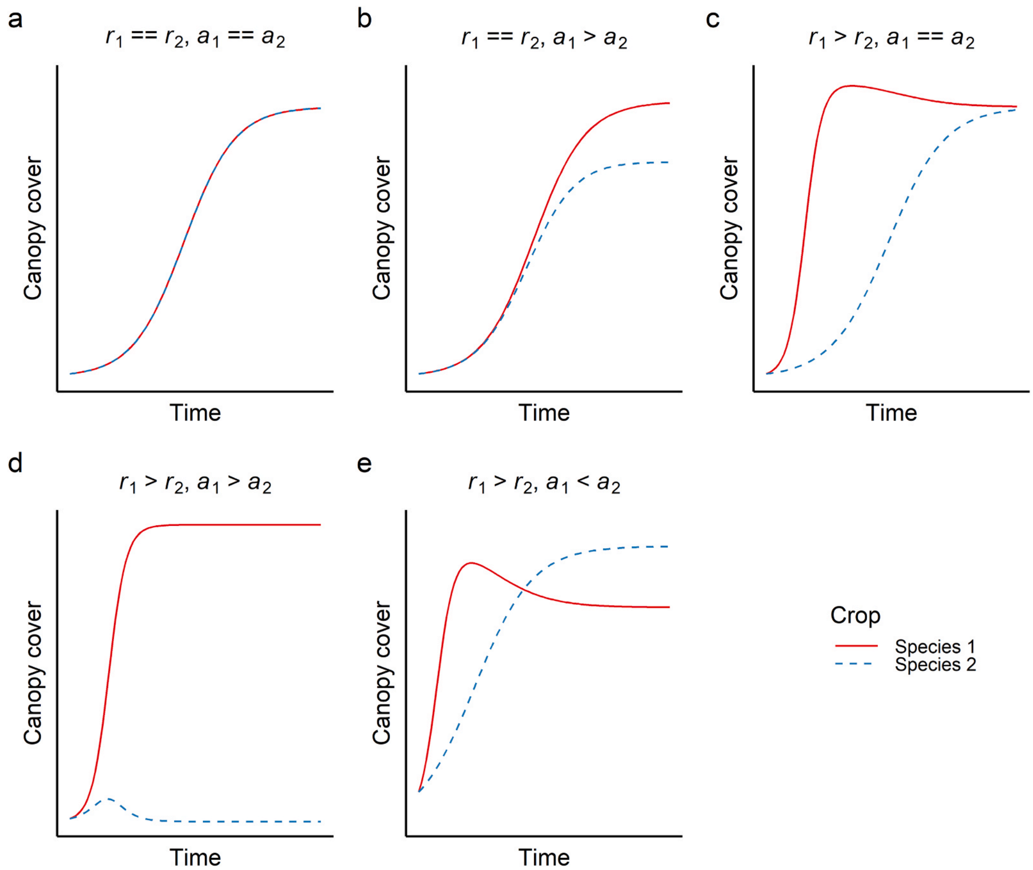

Fig. 2. Simulated canopy cover dynamics in cereal-legume intercrops using Lotka-Volterra competition models. Each panel shows temporal trajectories under different combinations of intrinsic growth rates (r_1, r_2) and competition coefficients (a_1, a_2) for species 1 and 2: (a) equal parameters for both species ($r_1 == r_2$, $a_1 == a_2$); (b) higher competitive effect of species 1 ($r_1 == r_2$, $a_1 > a_2$); (c) higher growth rate for species 1 ($r_1 > r_2$, $a_1 == a_2$); (d) species 1 superior in both growth and competition ($r_1 > r_2$, $a_1 > a_2$); (e) growth-competition trade-off with species 1 having higher growth but lower competitive ability ($r_1 > r_2$, $a_1 < a_2$).

data at each time step, where the data could be either ground-truth data (from manual image annotation) or CNN-model-derived data. The Lotka-Volterra model predicts the proportion of canopy cover for each species at each time step, which we denote as P_c. The beta distribution shape parameters α and β are then defined as $\alpha = P_c\phi$ and $\beta = (1 - P_c)\phi$, where ϕ is the dispersion parameter that is optimized during the optimisation procedure. This approach allowed the model to account for natural variation in cover proportions while anchoring the likelihood to biologically meaningful predictions.

Initial parameter values for the optimisation were manually adjusted to improve convergence and enhance the quality of the final fit. The 95 % confidence intervals (CIs) were estimated by calculating the standard error from the Hessian matrix, which contains the second-order partial derivatives of the log-likelihood function (i.e., the curvature of the negative log-likelihood profile). The CIs were then determined by adding and subtracting 1.96 times the standard error from the parameter estimates (Bolker, 2008).

2.3. Deep learning networks

We trained multiple deep learning models to perform image segmentation, distinguishing among the three annotated classes: faba bean, weed, and background (which is cereal + soil). Cereal cover was estimated indirectly by subtracting faba bean and weed cover from total cover, simplifying model training to three classes instead of four. We employed semantic segmentation, meaning that individual plants within a class were not differentiated; instead, all pixels belonging to the same class were treated as part of a single category.

We employed a U-Net segmentation model (Ronneberger et al., 2015), which has been shown to work well for crop and weed image segmentation (Kim and Park, 2022; Zou et al., 2021; Asad and Bais, 2020). The U-Net follows the standard encoder–decoder design with skip connections: feature maps are extracted at five scales (512 × 512, 256 × 256, 128 × 128, 64 × 64, 32 × 32 pixels) in the encoder, then upsampled and fused with same-scale encoder outputs in the decoder, before a final 1 × 1 convolution produces a three-class softmax map (faba bean, weed, background). To improve performance, we tested two different encoder backbones, VGG-16 and ResNet-50, both pretrained on ImageNet. Although these models have been around for years, they still have recently been shown to perform well in agricultural weed segmentation (Mesías-Ruiz et al., 2024; Purohit et al., 2024; Yang et al., 2023; GC et al., 2022; Olsen et al., 2019).

To reduce overfitting of the segmentation model and artificially increase the size of our dataset, we applied two methods: data augmentation and data synthetization. Data augmentation included a combination of random scaling (zooming in or out), rotating (random between 0°–15°), and mirroring (50 % chance) of the training data. As every image is randomly altered just before being included in training, the model never sees the exact same image more than once. More options for data augmentation were tested, including hue, saturation, and value shifts, spatial shifts, jittering, and shearing, but these methods decreased model performance. The currently used techniques preserved the relative shapes of the plants while altering their absolute positions within the images. This approach helps prevent overfitting by ensuring the model does not rely on specific object locations in the training data.

Data synthetization was performed in two steps. First, images of cereal canopy cover from sole crop plots were split vertically into ten equal-width strips. These strips were then recombined, randomly selecting one strip from each of ten different images taken on the same date, to form new composite cereal canopy backgrounds while preserving temporal consistency. Second, individual faba bean and weed plants that were manually annotated as separate polygon objects were digitally extracted from their original images. These objects were randomly rotated and scaled, their edges softened to blend more naturally, and then randomly placed over the recombined cereal backgrounds to simulate realistic mixed-species canopies. This resulted in synthesized images of recombined strips with additional faba bean and weed plants pasted on over (Suppl. Fig. A.1). For every synthesized image, the number of plants added was randomly chosen from a normal distribution with mean 20 and standard deviation 5. The added plants were chosen based on their size categories of small, medium, and large weed and faba plants (Suppl. Table A.1). As weeds were less common in the images than faba bean plants, a higher amount of weed plants was chosen to add to the images to compensate for this class imbalance. In total, 2000 synthetic images were generated to be included in the training of the model. Data augmentation was also applied on these images before being included in training. The synthetic images were not included in the validation or test sets.

This process yielded two datasets: the original set of 397 images and an expanded set of 2397 images (including 2000 synthetic). Synthetic images permanently increased dataset size, whereas augmentations were applied dynamically during training. Each time an image was presented to the algorithm, random transformations were applied, so the model never encountered the exact same version twice.

As the VGG-16-Unet model performed better than the ResNet-50-Unet model (see Results), data augmentations and synthetization were applied only for the VGG-16-Unet model. The base ResNet-50-UNet model was still compared to the other models, to evaluate how much a change in backbone model matters for model performance and subsequent agroecological analysis. This resulted in the following four CNN models:

- ResNet-50-UNet (ResNet50)
- VGG-16-UNet (VGG16)
- VGG-16-Unet + data augmentation (VGG16 Aug)
- VGG-16-Unet + data augmentation + data synthetization (VGG16 Aug Syn)

The dataset was split into 60 % training data (238 images, or 2238 images with synthesized data), 20 % validation data (79 images, used to validate model performance during training), and 20 % test data (80 images, used only after model training was done, to evaluate the final model). Hyperparameters (model settings that control learning behaviour) were tuned on the validation set. This resulted in each model being trained for 120 epochs using a learning rate of 0.0001. Optimisation was done with the Adam optimiser, using a combination of Dice loss and focal loss as the loss function (Yeung et al., 2022). For CNN model training, we used the TensorFlow2 implementation of U-net from https://github.com/bubbliiiing/unet-tf2. The networks were trained on a computer equipped with an NVIDIA GeForce RTX 4060 graphics card. The adapted version of the implementation is published at https://github.com/David-BK/UNet-CNN-TF2-Canopy-Cover-Segmentation.

2.4. Evaluation of segmentation models

2.4.1. Evaluation of CNN models

CNN model performance was evaluated by calculating the number of true positive (TP), true negative (TN), false positive (FP) and false negative (FN) pixels, and the intersection over union (IoU). IoU measures the overlap between predicted and true pixel-based segmentation masks areas per class (c), providing an indicator of how accurately the model identifies the spatial extent of objects:

$$IoU_c = \frac{TP_c}{TP_c + FP_c + FN_c} \tag{3}$$

Similarly, the Precision, a measure for the proportion of pixels predicted as belonging to a particular category that are correctly classified, and Recall, a measure for the proportion of true pixels of a category that the model successfully identifies, are calculated:

$$Precision_c = \frac{TP_c}{TP_c + FP_c} \tag{4}$$

$$Recall_c = \frac{TP_c}{TP_c + FN_c} \tag{5}$$

Additionally, the weighted average IoU was calculated:

$$average\ IoU = \sum_{c=1}^{c} IoU_c \bullet w_c \tag{6}$$

where:

$$\sum_{c=1}^{c} w_c = 1.0 \tag{7}$$

IoU_c is the IoU of class c and w_c is the weight of that class defined by the proportion of pixels belonging to that class compared to the total number of pixels.

Furthermore, to evaluate how well model predictions perform for the pixel classification proportions, the mean error (ME) and root mean squared error (RMSE) of the predicted and true proportions of each class in the images were calculated:

$$ME_c = \frac{1}{n}\sum_{i=1}^{N}\left(\hat{y}_{i,c} - y_{i,c}\right) \tag{8}$$

$$RMSE_c = \sqrt{\frac{1}{n}\sum_{i=1}^{N}\left(\hat{y}_{i,c} - y_{i,c}\right)^2} \tag{9}$$

where ME_c and $RMSE_c$ are the mean error and root mean squared error of class c (cereal, faba bean, or weed), respectively, N is the number of images, $\hat{y}_{i,c}$ is the predicted proportion of category c in image i, and $y_{i,c}$ is the true proportion.

2.4.2. Evaluation of agroecological analysis using CNN model predicted data

In our original agroecological analysis, we fitted the Lotka–Volterra competition model (Eq. 1) using fully manual segmentations on all 397 images. To compare those "ground-truth" results with CNN-based segmentation outputs without introducing bias, we ensured that each CNN prediction came from a model that has never seen that image during training or validation. To achieve this, we performed a five-fold cross-validation on the entire dataset of 397 images. For each of the five folds we randomly split the images into 60 % training, 20 % validation, and 20 % test. We trained the CNN on the training set while monitoring performance on the validation set to prevent overfitting and select the best model. Once training was complete, we ran the model on the held-out 20 % test set to generate segmentations. Because each fold's test set was a different 20 % of the data, across the five folds every image appeared exactly once in a test set, and its segmentation comes from a model that was not trained on that image. We then concatenated all five test-set predictions to assemble a complete, 397-image predicted dataset. Finally, we fitted the Lotka–Volterra model to these CNN-segmented data and compared the parameter estimates side-by-side with those from the fully manual ground-truth segmentations.

Beyond pixel-level accuracy, we also assessed each CNN's ability to reproduce the overall canopy cover measurements that drive our competition model. For each network (ResNet50, VGG16, VGG16 with augmentation, and VGG16 with augmentation plus synthetic data), we compared the model's estimated cover percentages for cereal, faba bean, weed, and total cover against the manually annotated ground-truth values. We visualized these comparisons with scatterplots against the 1:1 line, and then tested for any consistent offsets or scaling errors by fitting linear mixed-effects models with fixed effects for the predicted canopy cover and random intercepts for plot-level variation. A non-zero intercept would reveal a systematic over- or underestimation of cover, while a slope different from 1.0 would indicate that the model exaggerates or compresses cover variability.

All agroecological data analysis code was done using R 4.4.2 (R Core Team, 2024), and can be found on https://github.com/David-BK/Intercropping—Canopy-Cover-Segmentation-Analyses

3. Results

3.1. Segmented canopy cover of cereal, faba bean, and weed

Manually annotated (ground-truth) canopy cover data revealed that cereals were the dominant crop in all treatments except for the 1:3 ratio triticale-faba bean intercrop (1 T:3 F), where triticale made up only 25 % of the sown proportion and canopy cover appeared more balanced between triticale and faba bean (Fig. 3). This reflects a clear pattern of cereal dominance over faba bean. In most intercrop combinations, faba bean increased its share of the canopy over time as cereal cover increased initially but declined toward the end of the vegetative stage. Exceptions were the rye-faba bean (1 R:1 F) and 3:1 ratio triticale-faba bean (3 T:1 F) intercrops, where rye's strong competitive ability and the high cereal proportion in 3 T:1 F (75 % triticale) likely reinforced cereal dominance throughout the measurement period. In 2022, canopy cover data were collected only from herbicide-treated plots, resulting in consistently low weed cover. Even so, the wheat-faba bean intercrop (1 W:1 F), which showed slightly lower early wheat cover, allowed some weed growth later in the season. In contrast, the 2023 experiment (where no herbicide was applied) revealed a different dynamic, with weed cover more closely mirroring cereal cover dynamics, initially increasing and gradually declining over time later, although at much lower levels.

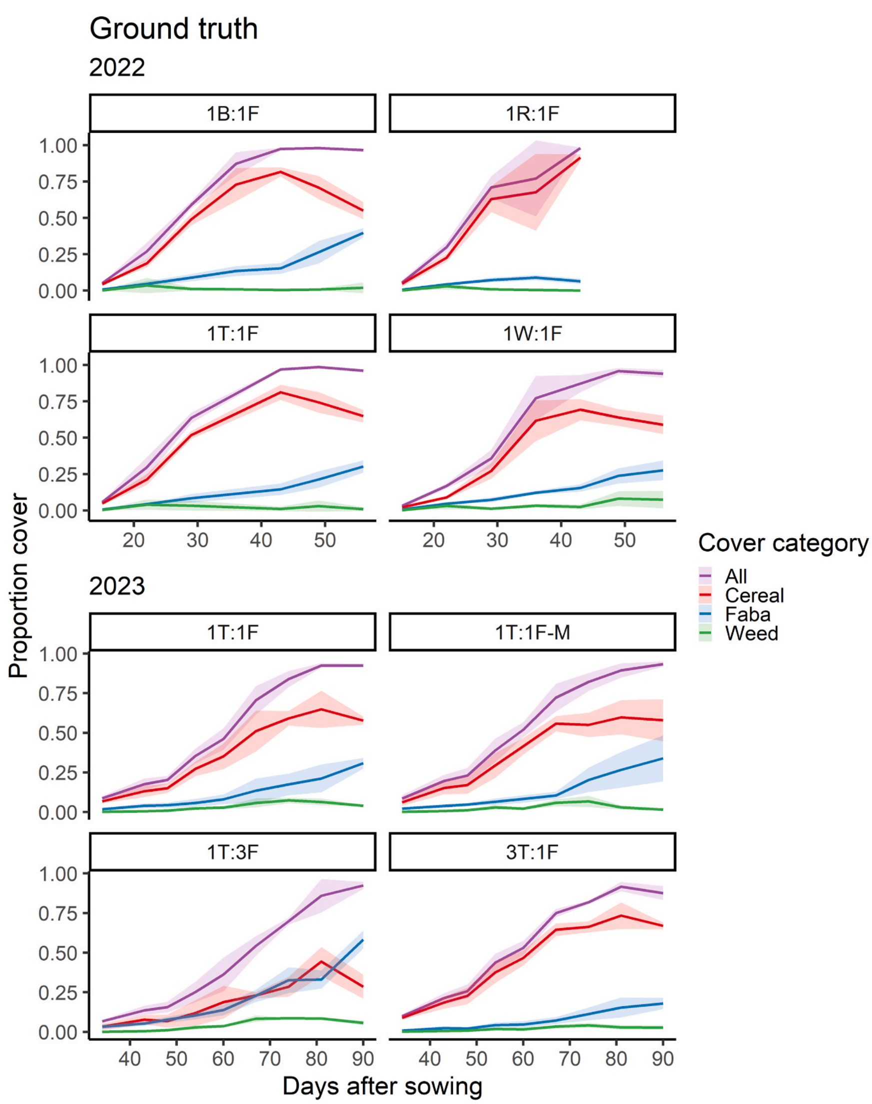

Fig. 3. Ground-truth segmented canopy cover. Proportion of canopy cover over time (days after sowing) for cereal (red), faba bean (blue), and weed (green) categories, as well as total canopy cover (purple), for the treatments in 2022 and 2023, as determined by manual annotation (ground-truth). The treatments included several 1:1 ratio row intercrops of faba bean with different cereals: barley (1B:1 F), triticale (1 T:1 F), rye (1 R:1 F), and wheat (1 W:1 F); and a triticale–faba bean intercrop with a 1:1 ratio mixed within rows (1 T:1F-M), as well as triticale–faba bean row intercrops with unbalanced ratios of 1:3 (1 T:3 F) and 3:1 (3T:1 F). Shaded areas indicate standard deviations.

3.2. Lotka-Volterra competition model parameter fits

To better understand the competitive dynamics between cereal and faba bean, we fitted the Lotka–Volterra competition model to the segmented ground-truth canopy cover data (Fig. 4). The model captured the general canopy cover trends of both crops well, with the estimated parameters providing insights into their interactions.

The intrinsic growth rates revealed distinct patterns among cereals: wheat (r_w) showed a significantly lower growth rate compared to other cereals, similar to that of faba bean (r_f), while rye (r_r), barley (r_b), and triticale (r_t) exhibited higher rates (Fig. 5). In terms of competitive ability per unit cover (a_{cf} and a_{fc}), faba bean was more competitive than barley, triticale, and especially wheat in their respective intercrops, whereas rye outcompeted faba bean. Despite wheat's lower intrinsic growth rate and weaker per-unit competitiveness compared to faba bean, wheat achieved higher overall canopy cover. This resulted from all cereals establishing higher initial canopy cover than faba bean due to earlier emergence, demonstrating how timing advantages can override competitive disadvantages.

The 2023 parameters indicate that triticale is more competitive when it interacts more closely with faba bean, either through within-row mixing (1 T:1F-M) or a higher triticale proportion in the intercrop (3 T:1 F). Conversely, faba bean gains a competitive advantage when present in higher proportion (1 T:3 F), where triticale's competitiveness is markedly reduced. This substantial asymmetry in competition coefficients allowed faba bean cover to overtake triticale cover by the end

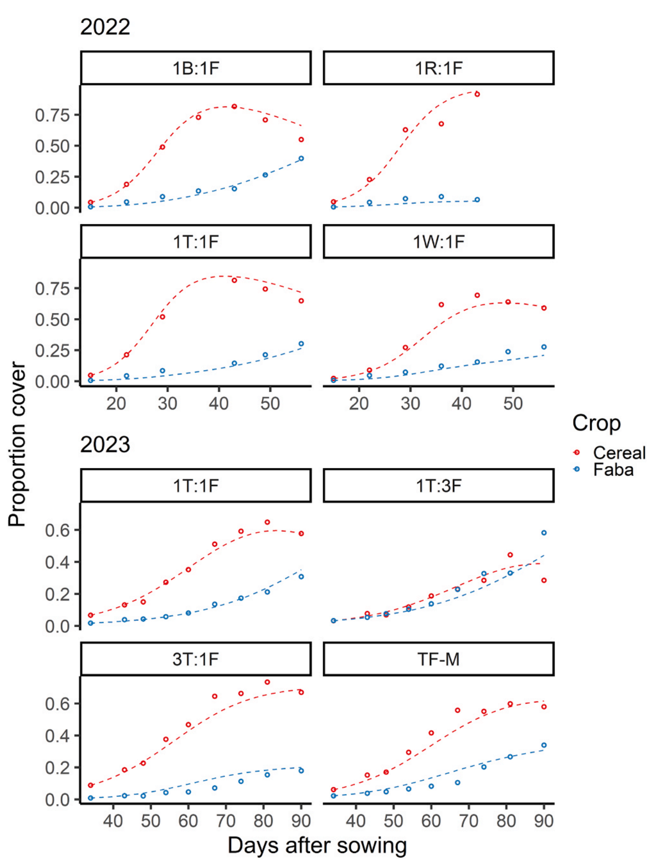

Fig. 4. Lotka-Volterra model fits. Ground-truth mean proportion canopy cover (points) and the Lotka-Volterra model fits (dashed lines) of cereal (red) and faba bean (blue) for the 2022 and 2023 experiments. The treatments include several 1:1 ratio row intercrops of faba bean with different cereals: barley (1B:1 F), triticale (1 T:1 F), rye (1 R:1 F), and wheat (1 W:1 F); and a triticale–faba bean intercrop with a 1:1 ratio mixed within rows (1 T:1F-M), as well as triticale–faba bean row intercrops with unbalanced ratios of 1:3 (1 T:3 F) and 3:1 (3 T:1 F).

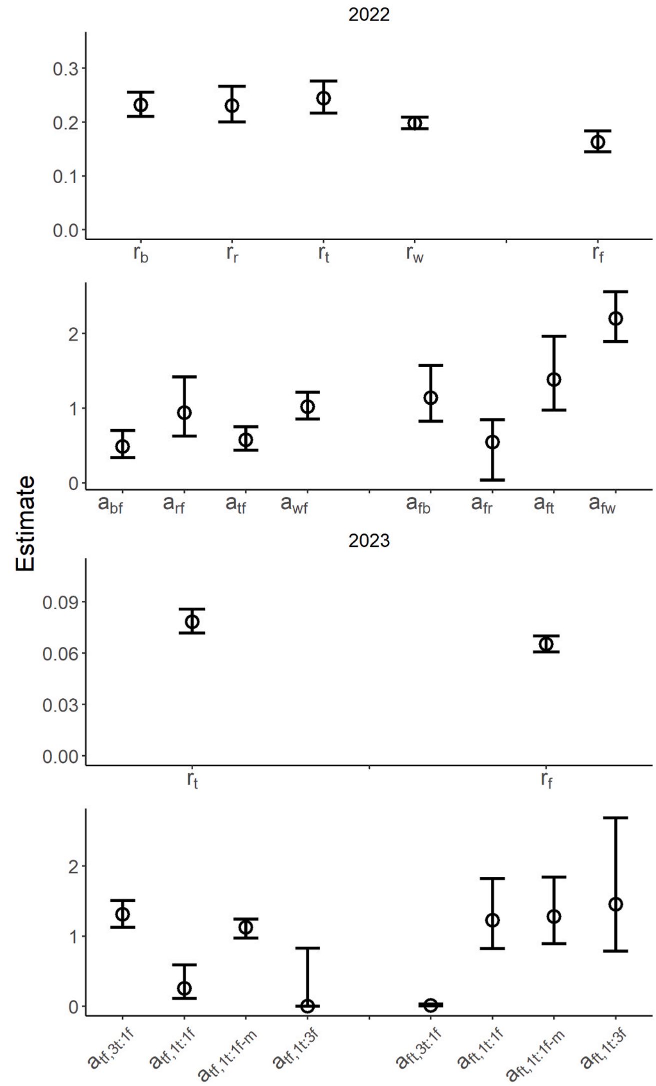

Fig. 5. Lotka-Volterra parameter estimates for different cereal-faba bean intercrops. Species-specific parameter estimates from the Lotka-Volterra model (Eq. 1) fitted to ground-truth data from 2022 and 2023. Parameters shown are intrinsic growth rates (r) and interspecific competition coefficients (a). In 2022, separate parameters were estimated for each cereal species in 1:1 row intercrops: barley (b), rye (r), triticale (t), and wheat (w), yielding species-specific intrinsic growth rates (r_b, r_r, r_t, r_w) and competition coefficients (e.g., a_{bf}, a_{fb} for the effect of barley on faba bean and faba bean on barley, respectively). A single r_f was estimated across all treatments as faba bean (f) was the only legume species. In 2023, with triticale as the only cereal, a single r_t and r_f were estimated, while competition coefficients were treatment-specific to capture effects of different spatial arrangements: 1:1 row intercrop (1t:1 f), 1:1 within-row mixture (1t:1f-m), and unbalanced row intercrops at 1:3 (1t:3 f) and 3:1 (3t:1 f) ratios. Error bars represent 95 % confidence intervals.

of the measurement period (Fig. 4), despite faba bean's lower intrinsic growth rate, a pattern consistent with the theoretical framework presented earlier (Fig. 2e). Despite the wheat faba-bean intercrop having a similar parameter structure (Fig. 5), the difference is smaller, preventing faba bean canopy cover from surpassing that of wheat by the end of the measurement period.

3.3. CNN model performance

3.3.1. Segmentation performance

All CNN models achieved high segmentation performance, with mean IoUs between 0.900–0.926 (Fig. 6, Table 1). VGG16 consistently

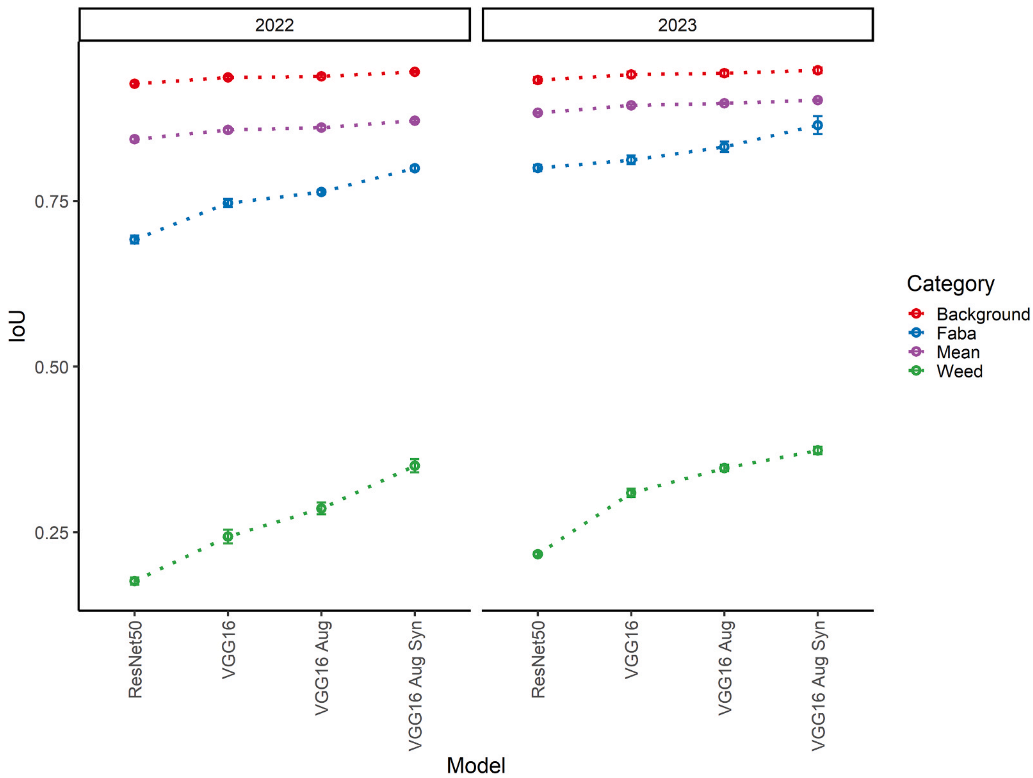

Fig. 6. Intersection over union (IoU) of canopy cover segmentations by four CNN models, based on the ResNet-50 (ResNet50) or VGG-16 backbone (VGG16), with augmentation (VGG16 Aug), and synthetization (VGG16 Aug Syn), for 2022 and 2023 data. Error bars indicate the standard error across the five models that were trained to comprise the full predicted dataset. 'Background' reflects the performance in identifying the combination of cereal and soil, while 'Mean' reflects the weighted average of the three discerned categories, with weighing factors based on the pixel proportion of each category.

Table 1
CNN model intersection over union (IoU), mean errors (ME), and root mean squared errors (RMSE) of canopy cover images for predictions of 2022 and 2023 image data, with standard deviation of IoU values in brackets. Values are given for the Background, Faba, and Weed categories, as well as the Mean, which is weighted based on the number of pixels of the categories in the image.

Year	Model	Background	Faba	Weed	Mean
IoU					
2022	ResNet−50	0.938 (0.009)	0.694 (0.028)	0.183 (0.022)	0.900 (0.014)
2022	VGG−16	0.944 (0.008)	0.752 (0.019)	0.266 (0.034)	0.911 (0.013)
2022	VGG−16 Aug	0.947 (0.007)	0.767 (0.011)	0.309 (0.027)	0.916 (0.012)
2022	VGG−16 Aug Syn	0.953 (0.006)	0.802 (0.017)	0.373 (0.039)	0.926 (0.009)
2023	ResNet−50	0.928 (0.021)	0.786 (0.026)	0.214 (0.013)	0.907 (0.006)
2023	VGG−16	0.936 (0.020)	0.838 (0.024)	0.306 (0.022)	0.918 (0.008)
2023	VGG−16 Aug	0.938 (0.020)	0.848 (0.031)	0.339 (0.022)	0.922 (0.006)
2023	VGG−16 Aug Syn	0.942 (0.020)	0.861 (0.038)	0.368 (0.026)	0.926 (0.006)
ME					
2022	ResNet−50	−0.0023	0.0048	−0.0025	−0.0010
2022	VGG−16	−0.0074	0.0101	−0.0028	−0.0041
2022	VGG−16 Aug	−0.0087	0.0081	0.0006	−0.0054
2022	VGG−16 Aug Syn	−0.0092	0.0071	0.0021	−0.0056
2023	ResNet−50	0.0089	−0.0004	−0.0084	0.0063
2023	VGG−16	0.0029	0.0014	−0.0043	−0.0030
2023	VGG−16 Aug	0.0008	0.0012	−0.0020	0.0013
2023	VGG−16 Aug Syn	0.0023	0.0015	−0.0038	0.0022
RMSE					
2022	ResNet−50	0.0192	0.0129	0.0116	0.0162
2022	VGG−16	0.0168	0.0143	0.0093	0.0145
2022	VGG−16 Aug	0.0157	0.0116	0.0086	0.0129
2022	VGG−16 Aug Syn	0.0152	0.0102	0.0086	0.0122
2023	ResNet−50	0.0106	0.0046	0.0100	0.0081
2023	VGG−16	0.0071	0.0040	0.0075	0.0049
2023	VGG−16 Aug	0.0052	0.0030	0.0052	0.0029
2023	VGG−16 Aug Syn	0.0052	0.0031	0.0058	0.0031

outperformed ResNet50 in terms of IoU, with performance further improving through data augmentation (VGG16 Aug) and combining data augmentation and synthetization (VGG16 Aug Syn). The CNN segments three classes: faba bean, weed, and background (cereal + soil). Cereal is not predicted as a separate class; its cover is later derived by subtracting the predicted faba bean and weed cover from the total canopy. Component-level IoU was highest for the background and lowest for weeds. The largest gains in performance were observed in the weed and faba bean categories, where initial segmentation accuracy left more room for improvement. Similar improvements can be seen in the precision and recall metrics (Suppl. Table B.1).

While the VGG16 model and its enhanced variants (with augmentation and synthetization) substantially improved IoU for some categories, the corresponding reductions in mean error or RMSE for pixel-wise category assignment were modest (Table 1). This indicates that improvements in segmentation accuracy do not necessarily lead to similar improvements in pixel-level category assignments. Mean errors reveal that faba bean cover is often overestimated, while weed cover is often underestimated (Table 1).

3.3.2. Predicted proportions canopy cover

The canopy cover patterns predicted by the deep learning models closely resembled those of the manually annotated ground-truth data, with visual differences being minimal and difficult to detect (Suppl. Figs. B.1-B.4). Statistical comparisons confirmed that these differences were not significant (Suppl. Table B.2 and Suppl. Fig. B.5). This suggests that all trained deep learning models were effective in capturing the temporal dynamics of canopy cover in intercrops and are well suited for visualizing these patterns.

3.3.3. Lotka-Volterra parameter fits

The CNN-derived canopy cover data yielded Lotka–Volterra model fits that closely resembled those based on the ground-truth segmentation data (Suppl. Figs. B.6-B.9). Parameter estimation accuracy varied by intercrop system and CNN methodology (Figs. 7 and 8). For the barley, rye, and triticale intercrops in 2022 (1B:1 F, 1 R:1 F, and 1 T:1 F), all CNN-derived parameter estimates fell within the 95 % confidence intervals of the ground-truth estimates, demonstrating robust performance across these systems. However, the wheat intercrop (1 W:1 F) showed systematic biases: all CNNs overestimated the wheat growth rate (r_w),

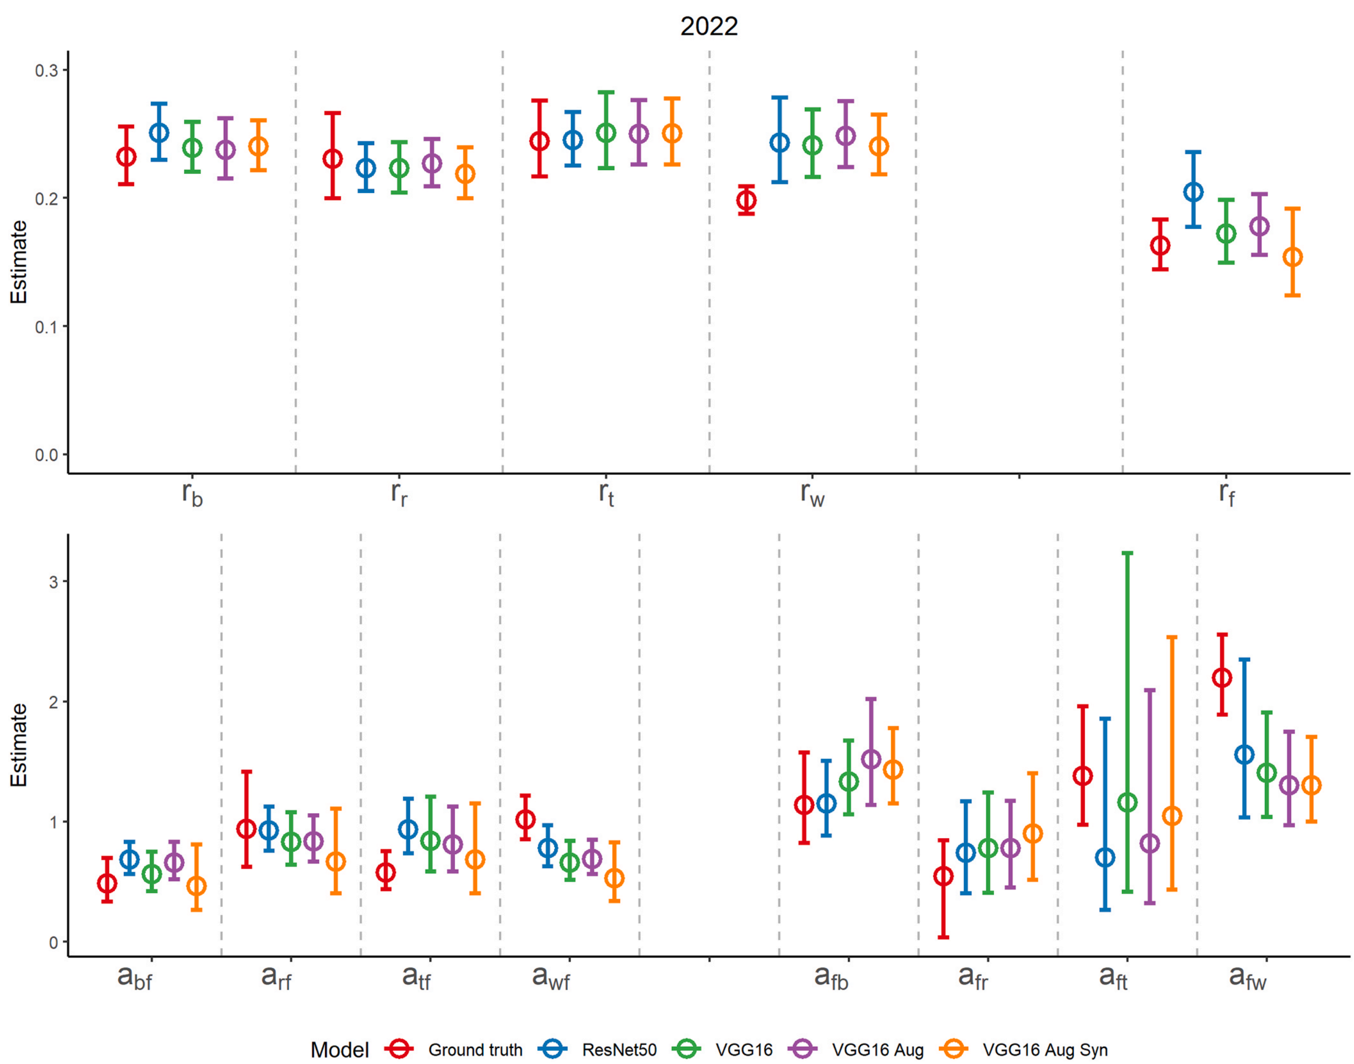

Fig. 7. CNN-model Lotka-Volterra parameters for different species combinations. Parameter estimates of the Lotka-Volterra model (Eq. 1) fitted to the 2022 ground-truth data and CNN-derived data from four CNN models (ResNet-50, VGG-16, VGG-16 with augmentation (VGG16 Aug Syn), and VGG-16 with augmentation and synthetic data (VGG16 Aug Syn). Parameters shown are species-specific intrinsic growth rates for cereals (r_b for barley, r_r for rye, r_t for triticale, r_w for wheat), a single intrinsic growth rate for faba bean (r_f), and species-specific interspecific competition coefficients (e.g., a_{bf} and a_{fb} for the effect of barley on faba bean and faba bean on barley, respectively). All parameters were estimated from 1:1 row intercrop data of faba bean with each cereal species. Error bars represent 95 % confidence intervals.

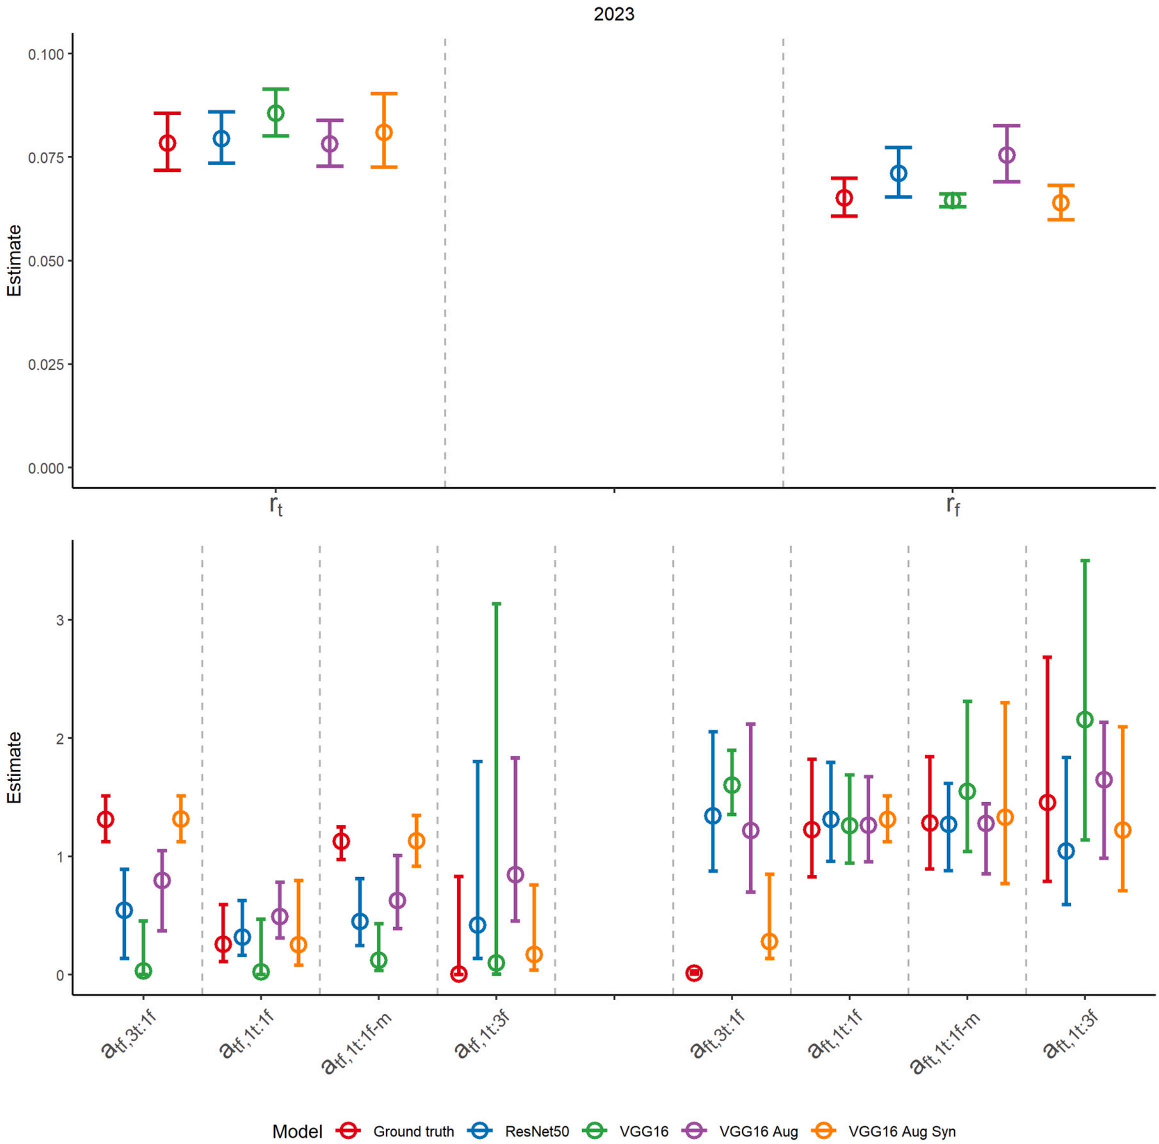

Fig. 8. CNN-model Lotka-Volterra parameters for different spatial arrangements. Parameter estimates of the Lotka-Volterra model (Eq. 1) fitted to 2023 ground-truth data and CNN-derived data from four CNN models (ResNet-50, VGG-16, VGG-16 with augmentation (VGG16 Aug), and VGG-16 with augmentation and synthetic data (VGG16 Aug Syn). Parameters shown are the intrinsic growth rates for triticale (r_t) and faba bean (r_f), and treatment-specific interspecific competition coefficients capturing effects of different spatial arrangements: a_{tf} and a_{ft} for the effect of triticale on faba bean and faba bean on triticale, respectively, in 1:1 row intercrops (1t:1 f), 1:1 within-row mixtures (1t:1f-m), and unbalanced row intercrops at 1:3 (1t:3 f) and 3:1 (3t:1 f) ratios. The treatment-specific competition coefficients reflect how spatial configuration affects competitive interactions between species. Error bars represent 95 % confidence intervals.

masking the distinctively lower growth rate that distinguished wheat from other cereals in the ground-truth analysis (Fig. 7). This overestimation appeared to be partially compensated by an underestimated competition coefficient for faba bean on cereal (a_{fw}).

The 2023 experiments revealed model-specific differences in parameter estimation (Fig. 8). For the 3 T:1 F and 1 T:1F-M treatments, where triticale competition was stronger, ResNet50, VGG16, and VGG16 Aug significantly underestimated the a_{tf} parameter compared to ground truth, whereas VGG16 Aug Syn produced estimates consistent with the ground-truth values. Similarly, all models overestimated the 3 T:1 F a_{ft} parameter, though VGG16 Aug Syn again provided the most accurate estimate. Across all comparisons, CNN-derived parameters often exhibited wider confidence intervals than ground-truth estimates, reflecting greater uncertainty in parameter estimation when using CNN-derived rather than manually annotated canopy cover data.

4. Discussion

This paper addressed two questions: what competitive dynamics emerge in cereal–faba bean intercrops, and whether CNN-based segmentation accuracy is sufficient for agroecological analysis. We found distinct competitive hierarchies among cereals with implications for intercrop design, and that CNN performance requirements depend on analytical objectives. We discuss the mechanisms underlying these competitive differences and the relationship between segmentation quality and ecological inference reliability.

4.1. *Segmented canopy cover reveals competitive dynamics in cereal–legume intercrops*

Distinguishing canopy cover by species, rather than measuring total cover alone, reveals the underlying competitive interactions between crops and weeds in intercropping systems. Real-time identification of species-specific canopy cover enables precision agriculture applications and informed management decisions. Moreover, tracking temporal dynamics of species-specific cover throughout the growing season, rather than relying solely on endpoint measurements like biomass or yield, captures the competitive processes that determine intercrop performance. This temporal perspective informs design choices including species selection, spatial arrangement, and management timing. Our combined approach of analysing canopy cover trajectories and fitting Lotka–Volterra competition models quantified these dynamics, revealing which species dominate, under what conditions coexistence occurs, and how spatial configuration and sowing ratios modulate competitive outcomes. While Lotka–Volterra models have long been used in ecology, their application to non-destructively derived canopy cover data for quantifying competitive parameters and guiding intercrop design represents a novel use of this framework in agricultural systems.

The canopy dynamics reveal a clear hierarchy in cereal competitive ability against faba bean, with rye as the most suppressive, followed by barley and triticale, and wheat as the least suppressive. The Lotka-Volterra parameters identify two distinct mechanisms underlying these differences. First, intrinsic growth rates (r_c) separate wheat from the other cereals: barley, triticale, and rye shared similarly high growth rates, while wheat's lower intrinsic growth rate constrained its competitive impact. Second, per-unit-cover competitive effects varied among the faster-growing cereals, with rye exerting stronger suppression per unit cover than barley or triticale, which in turn were slightly weaker competitors than faba bean on a per-cover basis. Notably, wheat achieved higher cover proportions than faba bean despite lower competitiveness per unit of projected area, primarily because wheat established earlier and thus began with greater initial cover. This timing advantage allowed wheat to capture canopy space before faba bean could respond, highlighting the importance of emergence time in competitiveness (Kropff and Spitters, 1991; Cousens et al., 1987). These patterns align with previous findings from the same experimental systems, where rye-based intercrops showed superior suppression of both weeds and faba bean compared to wheat-based systems (Kottelenberg et al., 2025a; 2025b).

In triticale-faba bean intercropping, species ratio and spatial design influenced canopy cover dynamics. In the 1 T:3 F row treatment, cereal and legume remained balanced for most of the season, only at the end did faba bean overtake, a pattern explained by faba bean's lower intrinsic growth rate ($r_f < r_t$) but superior competitive ability ($a_{ft} > a_{tf}$) (Fig. 5). The 1 T:1 F row treatment showed similar dynamics. In contrast, treatments with either greater within-row mixing (1 T:1F-M) or higher cereal proportions (3 T:1 F) shifted the competitive balance toward triticale, as reflected in altered competition coefficients that favored cereal dominance. These patterns are consistent with previous studies which have demonstrated that increasing cereal proportion intensifies competitive pressure on faba bean (Dhima et al., 2014; Aynehband and Behrooz, 2011). Here, we are able to quantify these differences to inform the design of intercrops with more balanced competitive dynamics. If competitiveness is too unequal, the weaker species may be over-suppressed, resulting in significant yield losses (Gu et al., 2025; Monti et al., 2016; Hauggaard-Nielsen et al., 2008). Beyond system design factors examined here, in future research this approach could quantify how competition dynamics respond to environmental conditions (e.g., nutrient availability, weather variability) or management practices (e.g., seeding date, irrigation), providing insights for optimising intercrop systems across diverse contexts.

Competition in intercrops has previously been analysed by destructive biomass and yield measurements collected over time (Wu et al., 2023; Zhang et al., 2015; Andersen et al., 2007). Our approach offers both practical and analytical advantages. Practically, canopy cover measurements are non-destructive, inexpensive, and can be collected repeatedly throughout the season without damaging plants, enabling higher temporal resolution at lower cost than destructive biomass sampling. Analytically, the Lotka–Volterra model applied to canopy cover data separates intrinsic growth rates from per-unit-cover competitive effects, providing mechanistic insight into how competitive imbalances develop.

Beyond the Lotka–Volterra model, canopy cover time series could support a range of ecological analyses. For instance, ground cover proportions at each measurement date could be used to estimate plant height and leaf nitrogen content (Lu et al., 2021), while temporal changes in canopy cover could serve as input for light competition or radiation interception models using canopy cover as a proxy for the leaf area index (Raj et al., 2021; Slattery and Ort, 2021). Such applications illustrate that CNN-derived canopy segmentation provides a flexible data source for quantifying plant interactions and canopy dynamics in cropping systems.

4.2. The complexity of CNN models needed depends on the sensitivity of the analysis

While convolutional neural networks offer automation potential for canopy segmentation, their utility depends on the analytical objective. For visual assessment of canopy dynamics, even simple CNN architectures and methodologies proved adequate, capturing temporal patterns comparable to ground-truth data. However, deriving Lotka–Volterra competition parameters imposed stricter accuracy requirements. Simpler models and those trained without data augmentation failed to preserve subtle ecological differences, such as treatment-specific variation in competitive effects or differences in intrinsic growth rates among cereals (Fig. 7). Our most methodologically complex model, trained with both augmentation and synthetic data, performed better, preserving some ecologically important patterns (Fig. 8), but still missed distinctions present in ground-truth data, particularly parameters that showed compensatory relationships (e.g., low intrinsic growth offset by reduced competitive pressure). This compensation illustrates that multiple parameter combinations can produce similar canopy cover dynamics (Fig. 4; Suppl. Fig. B.9), creating uncertainty about which parameter set reflects biological reality. All CNN-derived parameters showed wider confidence intervals than ground-truth estimates, reflecting the inherent uncertainty of using CNN-derived rather than manually annotated canopy cover data.

These findings suggest that standard CNN architectures and methodologies alone are insufficient for detailed Lotka–Volterra modelling, but enhanced training strategies incorporating data augmentation and synthetic data substantially improve parameter estimation. Note that the methodological complexity examined here did not include architectural aspects of the model such as network depth or layer configuration. Further improvements to CNN performance may require more sophisticated CNN architectures, enriched training datasets, higher measurement frequency, or alternative sensing methods such as LiDAR or multispectral imaging to achieve the precision needed for capturing subtle competitive dynamics.

Interestingly, while our CNN models with data augmentation and synthetization achieved higher IoUs, these gains did not translate into substantial improvements of the pixel classification RMSE (Table 1). This discrepancy suggests that while enhanced segmentation accuracy refines individual pixel classifications, overall pixel counts for each category remains balanced even with minor classification errors. We hypothesize that CNN misclassifications mostly occur at leaf edges, with errors distributed roughly equally between neighbouring categories. As a result, total category counts remain relatively accurate even though local classification errors are frequent. Supporting this hypothesis, precision, which reflects a model's ability to avoid false positives, and recall, which reflects its ability to avoid false negatives, tend to show similar values across models, years, and categories (Suppl. Table B.1). This balance between false positives and false negatives likely contributes to the stability of pixel count proportions. This would also explain the difficulty in classifying the weed category. Weed plants are generally smaller than faba bean plants, resulting in a relatively high proportion of edges in this category. While the better-performing CNN models have minimal impact on average canopy cover per category, they appear to improve consistency by reducing the noise in these predictions (Suppl. Table B.2). From a modelling perspective, this distinction matters because canopy cover dynamics are presented as averages per time-point, whereas the Lotka–Volterra model uses individual observations as input. As a result, all of our CNN models produce visually near-perfect canopy cover dynamics, but only the highest-performing model (VGG16 Aug Syn) approaches the ground truth in estimating Lotka–Volterra competition parameters (Figs. 7 and 8). These findings highlight the importance of defining performance requirements relative to research goals rather than defaulting to exhaustive model refinement.

Our CNN models demonstrated strong performance particularly considering the smaller dataset used for training. Despite being trained on just 397 original images, our models achieved mean IoUs ranging from 0.900 to 0.926 across different years, underscoring the effectiveness of our approach. Direct comparison with published models is difficult given differences in image acquisition methods and data collection costs. However, our performance compares favourably with values reported from larger datasets (Kim and Park, 2022; Khan et al., 2020), and may benefit from working with faba bean rather than morphologically more complex legumes such as the pea used by Munz and Reiser (2020). However, weed detection remained a challenge, with our models achieving a maximum weed IoU of 0.373. This lower performance likely resulted from the limited number of weeds in our training data, which created class imbalance and caused the model to focus primarily on crop categories. This discrepancy highlights the difficulties associated with detecting smaller plants and the potential benefits of larger, more balanced datasets. Even though with data synthetization we included relatively more weed plants, the improvement in weed IoU was modest. Furthermore, as discussed above, the low performance of our models on weed segmentation could be the results of the relatively small weed plants in our data, resulting in a high proportion of edges. Additionally, our smaller dataset may have introduced more noise in the segmentations, which in turn impaired the stability of Lotka–Volterra model parameter estimates. While we investigated model refinement to improve segmentation quality, we did not explore the effect of dataset size. Indeed, more data generally results in higher CNN performance (Huang et al., 2025; Khalid et al., 2020), making it uncertain if any of our CNN models would perform sufficiently for the Lotka-Volterra parameter estimates with larger data quantities.

Simplifying the segmentation task to three classes (faba bean, weed, and background) rather than four (explicitly separating cereal and soil) was a deliberate choice to improve model stability. Increasing segmentation class number increases classification complexity and typically reduces pixel-wise accuracy due to greater class overlap and edge ambiguity (Hatkar and Ahmed, 2025; Luo et al., 2019). Because cereal cover could be indirectly derived from the remaining categories, explicitly annotating it would not necessarily improve model accuracy but would increase computational and methodological complexity without guaranteed analytical gain. However, this simplification comes with a potential trade-off: cereal canopy cover estimates depend on the accuracy of the other categories, so any misclassification of faba bean or weed pixels may indirectly affect the derived cereal cover. This approach prioritises stable model training and efficient use of annotated data, though future work could test whether training more on sole-crop images (with weeds annotated if present) and reducing the intercrop images or on a reduced set of measurement dates provides comparable performance. Such alternatives might further reduce annotation effort, but they would possibly limit the model's ability to generalise across diverse canopy mixtures and growth stages.

4.3. CNN complexity should match analytical sensitivity in agroecological research

Our case study demonstrates that CNNs, including off-the-shelf architectures of moderate complexity, can effectively support agroecological research by automating species-specific canopy segmentation. This is particularly useful for tasks such as visualizing canopy structure and estimating pixel proportions. The increasing accessibility of CNNs, driven by advances in deep learning and their uptake in fields like precision agriculture (e.g., crop classification, disease detection, and yield estimation; Coulibaly et al., 2022), opens new opportunities for their adoption in agroecology. By showing that relatively simple models can reproduce key ecological patterns from canopy cover images, our study highlights a low-barrier entry point for integrating deep learning into ecological research and precision crop management.

However, our findings also reveal clear limitations. While CNN-based segmentations are sufficient for qualitative assessments, their prediction noise limit more complex tasks such as those estimating competition parameters. Although our best-performing model likely does not reflect the full potential of deep learning, it showed only modest improvement in pixel-level classification RMSE compared to the lowest-performing model, leading to somewhat better Lotka–Volterra parameter estimates but still not always aligning with ground-truth values. This highlights the importance of matching model complexity to analytical goals. As increasingly refined segmentation models become available (Bhatti et al., 2024; Yang et al., 2024; Sahin et al., 2023; Kim and Park, 2022), we encourage interdisciplinary collaboration to ensure that CNN performance aligns with the precision demands of ecological inference. In doing so, agroecologists can benefit from machine learning advances without defaulting to unnecessarily complex models. Our work contributes to bridging the gap between deep learning innovation and its practical deployment in sustainable, precision-managed cropping systems.

We propose a two-step approach for incorporating CNNs into agroecological research. First, manually annotated images should be used to train a standard off-the-shelf model. Second, the resulting segmentation data should be evaluated within the context of the intended agroecological model or analysis. Only if model performance proves insufficient should researchers invest in more data accumulation, more complex architectures, or synthetic image generation. This approach ensures that model refinement is driven by clearly defined analytical needs, rather than pursued by default, making deep learning more approachable and enabling agroecologists to integrate these tools even when technical expertise or resources are limited.

5. Conclusion

This study demonstrates a novel approach that combines species-specific canopy cover measurements with Lotka–Volterra competition modelling to analyse cereal–legume intercrops. By linking segmented canopy cover data directly to a mechanistic competition framework, we provide a new way to quantify and interpret species interactions throughout the first part of the growing season. Comparing Lotka–Volterra parameters derived from these measurements provides insights into competitive dynamics that average canopy cover trends alone cannot reveal. We show that simple, off-the-shelf CNN models can accurately capture average canopy cover dynamics over time, but that reliable estimation of Lotka–Volterra parameters demands either more advanced models, data enrichment techniques, or larger datasets. CNN complexity should therefore be tailored to the precision requirements of the intended agroecological analysis. Notably, strong visual and quantitative insights can still be achieved using accessible models trained on limited data, lowering the barrier for applying deep learning in agroecology. We recommend a goal-oriented, staged approach to CNN integration, starting with standard models and scaling up only as analytical requirements dictate, thus ensuring efficient use of resources while expanding the role of deep learning in precision-managed sustainable cropping systems.

CRediT authorship contribution statement

Rick van Essen: Writing – review & editing, Methodology. **Lammert Bastiaans:** Writing – review & editing, Supervision, Methodology, Conceptualization. **David Kottelenberg:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Conceptualization. **Jacob C. Douma:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Gert Kootstra:** Writing – review & editing.

Declaration of Competing Interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: David Kottelenberg reports financial support was provided by Nederlandse organisatie voor wetenschappelijk onderzoek (NWO). If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Xueying Liang for her contributions to CNN model preparation and Gijs Theuws for his work on data synthetization, and both for their efforts in data annotation.

Availability of data and code

The datasets generated and/or analysed during the current study along with the code used to analyse the data are available in a public GitHub, https://github.com/David-BK/Intercropping—Canopy-Cover-Segmentation-Analyses, and through a public dataset at doi:10.17026/LS/LFNB4K.

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.fcr.2025.110236.

Data availability

Data is made publically available through data repositories, as indicated in the manuscript where relevant.

References

Abdalla, A., Cen, H., Wan, L., Rashid, R., Weng, H., Zhou, W., He, Y., 2019. Fine-tuning convolutional neural network with transfer learning for semantic segmentation of ground-level oilseed rape images in a field with high weed pressure. Comput. Electron. Agric. 167. https://doi.org/10.1016/j.compag.2019.105091.

Andersen, M.K., Hauggaard-Nielsen, H., Weinier, J., Jensen, E.S., 2007. Competitive dynamics in two- and three-component intercrops. J. Appl. Ecol. 44 (3), 545–551. https://doi.org/10.1111/j.1365-2664.2007.01289.x.

Asad, M.H., Bais, A., 2020. Weed detection in canola fields using maximum likelihood classification and deep convolutional neural network. Inf. Process. Agric. 7 (4), 535–545. https://doi.org/10.1016/j.inpa.2019.12.002.

Aynehband, A., Behrooz, M., 2011. Evaluation of cereal-legume and cereal-pseudocereal intercropping systems through forage productivity and competition ability. Am. Eurasia J. Agric. Environ. Sci. 10, 675–683.

Bedoussac, L., Journet, E., Hauggaard-Nielsen, H., Naudin, C., Corre-Hellou, G., Prieur, L., Jensen, E., Justes, E., 2014. Eco-functional Intensification by Cereal-Grain Legume Intercropping in Organic Farming Systems for Increased Yields, Reduced Weeds and Improved Grain Protein Concentration. In: Bellon, S., Penvern, S. (Eds.), Organic Farming, Prototype for Sustainable Agricultures. Springer, Dordrecht, pp. 47–63. https://doi.org/10.1007/978-94-007-7927-3.

Bhatti, M.A., Syam, M.S., Chen, H., Hu, Y., Keung, L.W., Zeeshan, Z., Ali, Y.A., Sarhan, N., 2024. Utilizing convolutional neural networks (CNN) and U-Net architecture for precise crop and weed segmentation in agricultural imagery: A deep learning approach. Big Data Res. 36, 100465. https://doi.org/10.1016/j.bdr.2024.100465.

Bolker, B.M., 2008. Ecological Models and Data in R. illustrated edn. Princeton University Press, Oxford.

Bybee-Finley, K.A., Ryan, M.R., 2018. Advancing intercropping research and practices in industrialized agricultural landscapes. Agriculture 8 (6). https://doi.org/10.3390/agriculture8060080.

Champ, J., Mora-Fallas, A., Goeau, H., Mata-Montero, E., Bonnet, P., Joly, A., 2020. Instance segmentation for the fine detection of crop and weed plants by precision agricultural robots. Appl. Plant Sci. 8 (7), e11373. https://doi.org/10.1002/aps3.11373.

Chen, G., Kong, X., Gan, Y., Zhang, R., Feng, F., Yu, A., Zhao, C., Wan, S., Chai, Q., 2018. Enhancing the systems productivity and water use efficiency through coordinated soil water sharing and compensation in strip-intercropping. Sci. Rep. 8 (1), 10494. https://doi.org/10.1038/s41598-018-28612-6.

Cheriere, T., Lorin, M., Corre-Hellou, G., 2020. Species choice and spatial arrangement in soybean-based intercropping: Levers that drive yield and weed control. Field Crops Res. 256, 107923. https://doi.org/10.1016/j.fcr.2020.107923.

Coulibaly, S., Kamsu-Foguem, B., Kamissoko, D., Traore, D., 2022. Deep learning for precision agriculture: a bibliometric analysis. Intell. Syst. Appl. 16, 200102. https://doi.org/10.1016/j.iswa.2022.200102.

Cousens, R., Brain, P., O'Donovan, J.T., O'Sullivan, P.A., 1987. The use of biologically realistic equations to describe the effects of weed density and relative time of emergence on crop yield. Weed Sci. 35 (5), 720–725. https://doi.org/10.1017/S0043174500060872.

Crews, T.E., Carton, W., Olsson, L., 2018. Is the future of agriculture perennial? Imperatives and opportunities to reinvent agriculture by shifting from annual monocultures to perennial polycultures. Glob. Sustain 1. https://doi.org/10.1017/sus.2018.11.

Dhima, K.V., Vasilakoglou, I.B., Keco, R.X., Dima, A.K., Paschalidis, K.A, Gatsis, T.D., 2014. Forage yield and competition indices of faba bean intercropped with oat. Grass Forage Sci. 69 (2), 376–383. https://doi.org/10.1111/gfs.12084.

Douma, J.C., Weedon, J.T., 2019. Analysing continuous proportions in ecology and evolution: a practical introduction to beta and Dirichlet regression. Methods Ecol. Evol. 10 (9), 1412–1430. https://doi.org/10.1111/2041-210X.13234.

Ercolini, L., Grossi, N., Silvestri, N., 2022. A simple method to estimate weed control threshold by using RGB images from drones. Appl. Sci. 12 (23). https://doi.org/10.3390/app122311935.

Fawakherji, M., Potena, C., Pretto, A., Bloisi, D.D., Nardi, D., 2021. Multi-spectral image synthesis for crop/weed segmentation in precision farming. Robot. Auton. Syst. 146, 103861. https://doi.org/10.1016/j.robot.2021.103861.

Franco, J.G., King, S.R., Volder, A., 2018. Component crop physiology and water use efficiency in response to intercropping. Eur. J. Agron. 93, 27–39. https://doi.org/10.1016/j.eja.2017.11.005.

Gaba, S., Fried, G., Kazakou, E., Chauvel, B., Navas, M.-L., 2014. Agroecological weed control using a functional approach: a review of cropping systems diversity. Agron. Sustain. Dev. 34 (1), 103–119. https://doi.org/10.1007/s13593-013-0166-5.

GC, S., Zhang, Y., Koparan, C., Ahmed, M.R., Howatt, K., Sun, X., 2022. Weed and crop species classification using computer vision and deep learning technologies in greenhouse conditions. J. Agric. Food Res. 9, 100325. https://doi.org/10.1016/j.jafr.2022.100325.

Gu, C., Bastiaans, L., Anten, N.P.R., Makowski, D., van der Werf, W., 2021. Annual intercropping suppresses weeds: a meta-analysis. Agric. Ecosyst. Environ. 322. https://doi.org/10.1016/j.agee.2021.107658.

Gu, C., van der Werf, W., Bastiaans, L., 2022. A predictive model for weed biomass in annual intercropping. Field Crops Res. 277. https://doi.org/10.1016/j.fcr.2021.108388.

Gu, C., van der Werf, W., Bastiaans, L., 2025. Weed-suppressive ability and species dominance in intercropping: a meta-regression. J. Ecol. 113, 2234–2247. https://doi.org/10.1111/1365-2745.70090.

Gu, J., Wang, Z., Kuen, J., Ma, L., Shahroudy, A., Shuai, B., Liu, T., Wang, X., Wang, G., Cai, J., Chen, T., 2018. Recent advances in convolutional neural networks. Pattern Recognit. 77, 354–377. https://doi.org/10.1016/j.patcog.2017.10.013.

Hatkar T.S., Ahmed S.B. (2025). Urban scene segmentation and cross-dataset transfer learning using SegFormer. Eighth Int Conf Mach Vis Appl (ICMVA 2025) 13734. doi: 0.1117/12.3078755.

Hauggaard-Nielsen, H., Jørnsgaard, B., Kinane, J., Jensen, E.S., 2008. Grain legume–cereal intercropping: the practical application of diversity, competition and facilitation in arable and organic cropping systems. Renew. Agric. Food Syst. 23 (1), 3–12. https://doi.org/10.1017/S1742170507002025.

Huang, C.-C., Chiang, H.-F., Hsieh, C.-C., Zhu, B.-R., Wu, W.-J., Shaw, J.-S., 2025. Impact of dataset size on 3D CNN performance in intracranial hemorrhage classification. Diagnostics 15 (2). https://doi.org/10.3390/diagnostics15020216.

Jensen, E.S., Carlsson, G., Hauggaard-Nielsen, H., 2020. Intercropping of grain legumes and cereals improves the use of soil N resources and reduces the requirement for synthetic fertilizer N: a global-scale analysis. Crop Pasture Sci. 40. https://doi.org/10.1007/s13593-020-0607-x.

Khalid, O., Luca, C., Meredith, K., 2020. CNN performance dependence on linear image processing, 310-311 Electron. Imaging 32 (10), 310–311. https://doi.org/10.2352/ISSN.2470-1173.2020.10.IPAS-310.

Khan, A., Ilyas, T., Umraiz, M., Mannan, Z.I., Kim, H., 2020. CED-Net: crops and weeds segmentation for smart farming using a small cascaded encoder-decoder architecture. Electronics 9 (10). https://doi.org/10.3390/electronics9101602.

Kim, Y.H., Park, K.R., 2022. MTS-CNN: multi-task semantic segmentation-convolutional neural network for detecting crops and weeds. Comput. Electron. Agric. 199, 107146. https://doi.org/10.1016/j.compag.2022.107146.

Koshelev, I., Savinov, M., Menshchikov, A., Somov, A., 2023. Drone-aided detection of weeds: transfer learning for embedded image processing. IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens. 16, 102–111. https://doi.org/10.1109/jstars.2022.3224657.

Kottelenberg, D.B., 2024. Canopy Cover Analysis Tool. 0.6 edn. Zenodo. doi:10.5281/zenodo.13941037.
Kottelenberg D.B., Evers J.B., Anten N.P.R., Rangs M., Bastiaans L. (2025a). Effects of cereal-legume intercrop system design on weed suppression. PREPRINT (Version 2) available at Research Square. doi:10.21203/rs.3.rs-5455247/v2.
Kottelenberg D.B., Evers J.B., Anten N.P.R., Bastiaans L. (2025b). Managing species dominance in cereal-legume intercrop systems. PREPRINT (Version 1) available at Research Square. doi:10.21203/rs.3.rs-5930680/v2.
Kounalakis, T., Triantafyllidis, G.A., Nalpantidis, L., 2017. Image-based recognition framework for robotic weed control systems. Multimed. Tools Appl. 77 (8), 9567–9594. https://doi.org/10.1007/s11042-017-5337-y.
Kropff, M.J., Spitters, C.J.T., 1991. A simple model of crop loss by weed competition from early observations on relative leaf area of the weeds. Weed Res. 31 (2), 97–105. https://doi.org/10.1111/j.1365-3180.1991.tb01748.x.
Li, C., Hoffland, E., Kuyper, T.W., Yu, Y., Zhang, C., Li, H., Zhang, F., van der Werf, W., 2020. Syndromes of production in intercropping impact yield gains. Nat. Plants 6 (6), 653–660. https://doi.org/10.1038/s41477-020-0680-9.
Li, X.W., Lu, X.X., Zhang, Z.J., Huang, J., Zhang, J.M., Wang, L.K., Hafeez, M., Fernandez-Grandon, G.M., Lu, Y.B., 2021. Intercropping Rosemary (Rosmarinus officinalis) with Sweet Pepper (Capsicum annum) Reduces Major Pest Population Densities without Impacting Natural Enemy Populations. Insects 12 (1). https://doi.org/10.3390/insects12010074.
Lu, J., Cheng, D., Geng, C., Zhang, Z., Xiang, Y., Hu, T., 2021. Combining plant height, canopy coverage and vegetation index from UAV-based RGB images to estimate leaf nitrogen concentration of summer maize. Biosyst. Eng. 202, 42–54. https://doi.org/10.1016/j.biosystemseng.2020.11.010.
Luo, C., Xiojie, L., Yin, J., He, J., Gao, D., Zhou, J., 2019. How Does the Data Set and the Number of Categories Affect CNN-based Image Classification Performance? J. Softw. 14 (4), 168–181. https://doi.org/10.17706/jsw.14.4.168-181.
Ma, X., Deng, X., Qi, L., Jiang, Y., Li, H., Wang, Y., Xing, X., 2019. Fully convolutional network for rice seedling and weed image segmentation at the seedling stage in paddy fields. PLoS One 14 (4), e0215676. https://doi.org/10.1371/journal.pone.0215676.
Mesías-Ruiz, G.A., Borra-Serrano, I., Peña, J.M., de Castro, A.I., Fernández-Quintanilla, C., Dorado, J., 2024. Weed species classification with UAV imagery and standard CNN models: Assessing the frontiers of training and inference phases. Crop Prot. 182, 106721. https://doi.org/10.1016/j.cropro.2024.106721.
Monti, M., Pellicanò, A., Santonoceto, C., Preiti, G., Pristeri, A., 2016. Yield components and nitrogen use in cereal-pea intercrops in Mediterranean environment. Field Crops Res. 196, 379–388. https://doi.org/10.1016/j.fcr.2016.07.017.
Mortensen, D.A., Smith, R.G., 2020. Confronting barriers to cropping system diversification. Front. Sustain. Food Syst. Vol. 4–2020.
Munz, S., Reiser, D., 2020. Approach for image-based semantic segmentation of canopy cover in pea–oat intercropping. Agriculture 10 (8). https://doi.org/10.3390/agriculture10080354.
Olsen, A., Konovalov, D.A., Philippa, B., Ridd, P., Wood, J.C., Johns, J., Banks, W., Girgenti, B., Kenny, O., Whinney, J., Calvert, B., Azghadi, M.R., White, R.D., 2019. DeepWeeds: a multiclass weed species image dataset for deep learning. Sci. Rep. 9 (1), 2058. https://doi.org/10.1038/s41598-018-38343-3.
Patrignani, A., Ochsner, T.E., 2015. Canopeo: a powerful new tool for measuring fractional green canopy cover. Agron. J. 107 (6), 2312–2320. https://doi.org/10.2134/agronj15.0150.
Purohit, A., Pujari, P., Shah, K., Nimkar, A.V., 2024. Crop and Weed Segmentation Using ResNet-Unet Architecture, 29-30 Nov. 2024 2024 First Int. Conf. Data Comput. Commun. (ICDCC) 1–8. https://doi.org/10.1109/ICDCC62744.2024.10961765.
R Core Team (2024). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL ⟨https://www.R-project.org/⟩.
Raj, R., Walker, J.P., Pingale, R., Nandan, R., Naik, B., Jagarlapudi, A., 2021. Leaf area index estimation using top-of-canopy airborne RGB images. Int J. Appl. Earth Obs. Geoinf. 96, 102282. https://doi.org/10.1016/j.jag.2020.102282.
Revanasiddappa, B., Arvind, C.S., Swamy, S., 2020. Real-time early detection of weed plants in pulse crop field using drone with IoT. Int. J. Agric. Technol. 16 (5). https://doi.org/10.13140/RG.2.2.27656.03845.
Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. Springer International Publishing, Cham, pp. 234–241.
Sahin, H.M., Miftahushudur, T., Grieve, B., Yin, H., 2023. Segmentation of weeds and crops using multispectral imaging and CRF-enhanced U-Net. Comput. Electron. Agric. 211, 107956. https://doi.org/10.1016/j.compag.2023.107956.
Sapkota, B.B., Popescu, S., Rajan, N., Leon, R.G., Reberg-Horton, C., Mirsky, S., Bagavathiannan, M.V., 2022. Use of synthetic images for training a deep learning model for weed detection and biomass estimation in cotton. Sci. Rep. 12 (1), 19580. https://doi.org/10.1038/s41598-022-23399-z.
Slattery, R.A., Ort, D.R., 2021. Perspectives on improving light distribution and light use efficiency in crop canopies. Plant Physiol. 185 (1), 34–48. https://doi.org/10.1093/plphys/kiaa006.
Soetaert, K., Petzoldt, T., Setzer, R.W., 2010. Solving differential equations in R: package deSolve. J. Stat. Softw. 33 (9), 1–25. https://doi.org/10.18637/jss.v033.i09.
Su, D., Kong, H., Qiao, Y., Sukkarieh, S., 2021. Data augmentation for deep learning based semantic segmentation and crop-weed classification in agricultural robotics. Comput. Electron. Agric. 190, 106418. https://doi.org/10.1016/j.compag.2021.106418.
Tang, X., Zhang, C., Yu, Y., Shen, J., van der Werf, W., Zhang, F., 2020. Intercropping legumes and cereals increases phosphorus use efficiency; a meta-analysis. Plant Soil 460 (1-2), 89–104. https://doi.org/10.1007/s11104-020-04768-x.
Trinder, C.J., Brooker, R.W., Robinson, D., 2013. Plant ecology's guilty little secret: understanding the dynamics of plant competition. Funct. Ecol. 27 (4), 918–929.
Van Essen, R., Mencarelli, A., van Helmond, A., Nguyen, L., Batsleer, J., Poos, J.-J., Kootstra, G., 2021. Automatic discard registration in cluttered environments using deep learning and object tracking: class imbalance, occlusion, and a comparison to human review. ICES J. Mar. Sci. 78 (10), 3834–3846. https://doi.org/10.1093/icesjms/fsab233.
Wu, J., Bao, X., Zhang, J., Lu, B., Zhang, W., Callaway, R.M., Li, L., 2023. Temporal Stability of Productivity is Associated with Complementarity and Competitive Intensities in Intercropping. Ecol. Appl. 33 (1), e2731. https://doi.org/10.1002/eap.2731.
Yang, J., Chen, Y., Yu, J., 2024. Convolutional neural network based on the fusion of image classification and segmentation module for weed detection in alfalfa. Pest Manag. Sci. 80 (6), 2751–2760. https://doi.org/10.1002/ps.7979.
Yang, L., Xu, S., Yu, X., Long, H., Zhang, H., Zhu, Y., 2023. A new model based on improved VGG16 for corn weed identification. Front. Plant Sci. 14. https://doi.org/10.3389/fpls.2023.1205151.
Yeung, M., Sala, E., Schönlieb, C.-B., Rundo, L., 2022. Unified Focal loss: Generalising Dice and cross entropy-based losses to handle class imbalanced medical image segmentation. Comput. Med. Imaging Graph. 95, 102026. https://doi.org/10.1016/j.compmedimag.2021.102026.
Zhang, W.P., Liu, G.C., Sun, J.H., Zhang, L.Z., Weiner, J., Li, L., 2015. Growth trajectories and interspecific competitive dynamics in wheat/maize and barley/maize intercropping. Plant Soil 397, 227–238. https://doi.org/10.1007/s11104-015-2619-x.
Zou, K., Chen, X., Zhang, F., Zhou, H., Zhang, C., 2021. A Field Weed Density Evaluation Method Based on UAV Imaging and Modified U-Net. Remote Sens. 13 (2). https://doi.org/10.3390/rs13020310.