

Database/Software Article

Robust High-Throughput Phenotyping with Deep Segmentation Enabled by a Web-Based Annotator

Jialin Yuan ¹, Damanpreet Kaur ¹, Zheng Zhou,¹ Michael Nagle ¹,
Nicholas George Kiddle,¹ Nihar A. Doshi,² Ali Behnoudfar ¹, Ekaterina Peremyslova,¹
Cathleen Ma,¹ Steven H. Strauss ¹ and Fuxin Li ¹

¹Oregon State University, Corvallis, OR, USA

²University of Southern California, Los Angeles, CA, USA

Correspondence should be addressed to Jialin Yuan; yuanjial@oregonstate.edu, Damanpreet Kaur; kaurd@oregonstate.edu, and Fuxin Li; fuxin.li@oregonstate.edu

Received 16 September 2021; Accepted 17 March 2022; Published 19 May 2022

Copyright © 2022 Jialin Yuan et al. Exclusive Licensee Nanjing Agricultural University. Distributed under a Creative Commons Attribution License (CC BY 4.0).

The abilities of plant biologists and breeders to characterize the genetic basis of physiological traits are limited by their abilities to obtain quantitative data representing precise details of trait variation, and particularly to collect this data at a high-throughput scale with low cost. Although deep learning methods have demonstrated unprecedented potential to automate plant phenotyping, these methods commonly rely on large training sets that can be time-consuming to generate. Intelligent algorithms have therefore been proposed to enhance the productivity of these annotations and reduce human efforts. We propose a high-throughput phenotyping system which features a Graphical User Interface (GUI) and a novel interactive segmentation algorithm: Semantic-Guided Interactive Object Segmentation (SGIOS). By providing a user-friendly interface and intelligent assistance with annotation, this system offers potential to streamline and accelerate the generation of training sets, reducing the effort required by the user. Our evaluation shows that our proposed SGIOS model requires fewer user inputs compared to the state-of-art models for interactive segmentation. As a case study of the use of the GUI applied for genetic discovery in plants, we present an example of results from a preliminary genome-wide association study (GWAS) of *in planta* regeneration in *Populus trichocarpa* (poplar). We further demonstrate that the inclusion of a semantic prior map with SGIOS can accelerate the training process for future GWAS, using a sample of a dataset extracted from a poplar GWAS of *in vitro* regeneration. The capabilities of our phenotyping system surpass those of unassisted humans to rapidly and precisely phenotype our traits of interest. The scalability of this system enables large-scale phenomic screens that would otherwise be time-prohibitive, thereby providing increased power for GWAS, mutant screens, and other studies relying on large sample sizes to characterize the genetic basis of trait variation. Our user-friendly system can be used by researchers lacking a computational background, thus helping to democratize the use of deep segmentation as a tool for plant phenotyping.

1. Introduction

Advances in high-throughput genome sequencing and computation have enabled the investigation of genomic variation within large populations [1–4]. In genome-wide association studies (GWAS), statistical models are used to model relationships between genotypes and phenotypes. When significant relationships are located, these can lead to the discovery of genetic markers and/or genes responsible for the phenotype of interest [3, 5]. However, genomic sequencing capacity and precision have outstripped the capability to produce high-quality phenomic data at scale. Phenotyping often requires individual

inspection of large numbers of samples, with statistical power depending on the sample size; for instance, as the sample size of a GWAS or mutant screen increases, the representation of rare genetic markers increases along with the power of models to discover them [6–8]. Sample size can be limited by the time cost of phenotyping, which is affected by challenges in both the accuracy and precision of measuring complex phenotypes as are common in studies of plant growth and development [9, 10]. The constraint of sample size therefore limits the insights that can be gained from these studies. Furthermore, phenotyping by humans commonly involves summarization of phenotypes in a manner that limits detail and subjects data

to potential biases of individuals performing phenotyping. For example, each sample may be given an ordinal or binary “score” representing a general summary of their complex trait or traits, which would be difficult to quantitatively measure in a simple and objective manner. For studies similar to our case example involving plant regeneration, common traits studied include callus size and quality [11], rates of callus differentiation into shoot [12] and the numbers of shoots produced [13]. These are very difficult to accurately measure without an annotation and imaging system, and their measurement will often compromise the required sterility of *in vitro* cultures during sequential analyses of growth. Comparable levels of complexity can be seen in populations of growing cancerous cells imaged with microscopy [14, 15].

While deep learning has demonstrated unprecedented abilities for high-throughput and precise phenotyping of complex plant traits including plant stress, disease, and development of specific tissues, these methods generally require large training sets that are labor-intensive to generate [16–19]. The amount of labor needed for phenotyping can be greatly reduced by supervised deep learning, although not eliminated altogether due to an inherent reliance on training data, which requires manual inspection and labeling of some number of samples.

Multiple approaches have been employed to alleviate the need for large training sets or to reduce the time cost of generating them. The number of training samples required for desired accuracy can be reduced, for example, by employing transfer learning. With this ubiquitous approach, users begin with a model trained for a similar task (often using a gold-standard training set) and retrain the model using a training set for their specific task. This re-training process tends to require relatively few annotated examples, whereas many more would be needed if the model were not already trained for a similar task [20]. Additionally, data augmentation approaches, such as those involving image rotation or other linear transformations, can be used to produce multiple training samples from each image [21]. Generation of training labels is particularly laborious when research designs require precise segmentation of image boundaries rather than simple classification of images. In a recent study, deep segmentation of healthy and cancerous tissues relied on a training set with precisely drawn lines separating healthy and cancerous tissue; the labor-intensive task of producing this training set was crowdsourced among 25 medical doctors and students [22].

As an alternative to manual drawing of segment boundaries, several established methods and tools for annotation depend on drawing of boxes or polygons [23–26], presenting a challenge for precisely labeling curved edges of objects such as leaves and other plant tissues. Our system features a user-friendly web-based annotator tool (IDEAS, Intelligent DEep Annotator for Segmentation) that applies deep segmentation to detect boundaries of objects given positive and negative cues by the user and a semantic prior map, thus enabling rapid and reliable labeling of objects with pixel-level precision (Figure 1).

Deep learning requires a massive amount of data to train a model. Generation of large datasets by unassisted humans

may prove time-prohibitive, especially for complex tasks. The time cost can be mitigated by the use of intelligent annotation algorithms to simplify image annotation tasks. Much research has been dedicated to building intelligent algorithms for automatic image annotation and building image annotation tools [27]. Nevertheless, these algorithms fail to predict well for complex structures and require a considerable amount of user effort to correct the incorrect model predictions. Hence, there is a need for an efficient tool for image segmentation using an intelligent algorithm that requires minimal user interaction and can label datasets belonging to any distribution of data. Our system, IDEAS, facilitates labeling highly precise curved objects with minimal labor and works well for segmenting previously unseen classes.

For the backend algorithm used in IDEAS, we propose to utilize semantic probability maps as a prior that acts as an additional guide to our interactive object segmentation model. The availability of strong prior information helps to guide the neural network in producing accurate boundaries and reduces the level of user interaction needed in providing positive and negative markers. Semantic segmentation has made tremendous progress in the last few years and is now among the fastest and most mature algorithms in deep learning. With state-of-the-art architectures such as DeepLab [28], PSPNet [29], and RefineNet [30] delivering strong performance on challenging benchmarks, obtaining semantic results for images is no longer difficult. Considering the potential for transfer learning to reduce the amount of training data needed, we hypothesize that the user effort required in providing clicks to guide segmentation can be reduced if segmentation is guided by a prior generated by a model trained on a partial dataset.

Leveraging semantic information to guide the main task has been previously followed in object detection [31, 32], dense object reconstruction [33] and interactive co-segmentation [34] tasks. However, to the best of our knowledge, our method is the first to apply this strategy for interactive object segmentation. We performed a benchmark experiment to evaluate the time cost advantage of using a semantic prior map along with the user click inputs, and found that use of the semantic prior map provides a substantial improvement over the use of user inputs alone.

The nature of our overall system as a web-based tool offers advantages for collaborative annotation and system administration. Multiple users can log into the system online from personal computers, working on a server with IDEAS installed and managed centrally by an administrator. Users benefit from the flexibility of being able to work from any computer with internet access, while the computationally intensive task of generating labels is assigned to a GPU on the server. Furthermore, users without a computational background benefit from an ability to apply the system without installing and updating software or otherwise interacting with the command line, as these tasks are handled centrally by an administrator.

After completion of a training set with IDEAS, a neural network is trained and used for semantic segmentation of large numbers of images that could not be practically




FIGURE 1: Annotation of an example image of a poplar stem tip undergoing regeneration: (a) Image to be annotated. (b) Polygon-based annotation from LabelMe, which was generated using 364 clicks precisely on the boundaries. (c) Annotation from IDEAS, which was generated based on 25 positive marks and 49 negative marks. This comparison demonstrates how IDEAS can be used to produce highly precise boundaries for objects with reduced effort.

annotated by humans. We use the state-of-the-art PSPNet architecture [29], notable for a global pyramid pooling module that enhances model accuracy. Finally, traits computed from segmented images provide quantitative measures of growth of distinct plant tissues. These measures include the relative areas and counts of specific tissues, thus providing an array of statistics biologists may choose from to represent tissue development in downstream statistical models.

2. Materials and Methods

IDEAS (Figure 2) features separate front-end and back-end components. The front-end, accessible to the user through a web browser, contains all user-facing, interactive modules. These modules cooperate with each other to support the workflow and are implemented through HTML, JavaScript, and jQuery. On the back-end, the Python Flask web framework is used to receive requests from the front-end and return the generated object mask for the object being annotated. *IDEAS is deployed at* <https://ideas.eecs.oregonstate.edu> and is accessible to the public.

The development of IDEAS enables the collection of ground truth segmentation masks, to be used as training labels, for various types of plant images. A semantic segmentation model can then be trained using the collected annotations to compute masks for large sets of unseen images. Below, we introduce features of the front-end, back-end, and auxiliary functions on IDEAS.

2.1. IDEAS: Front-End Interactive Modules. The front-end interface of IDEAS features seven different modules (Figure 3):

① **Object and Class Panels:** Our system allows the user to customize the identity of objects and the classes they include. Figure 4(a) shows the unfolded class panel and object panel. The class panel displays a list of classes to be annotated. Using the object panel, the user can define objects and their nested classes by either 1) adding multiple classes to an object when defining the object on the object panel, or 2) using the button ‘add to’ on the class panel to add classes to a specified object one-by-one. The intuitive design of object and class hierarchy is influenced by applications that involve analysis of multiple, distinct parts of given objects (e.g. distinct plant tissues such as callus or shoot, growing in each individual sample on a petri dish (Figure 4(b))).

② **Toolkit:** Point-and-click tools for user annotation of images include: 1) Two types of pens, for drawing positive

and negative strokes inside and outside of a desired segment, respectively, which are used by the back-end to generate the object mask; 2) A bounding box tool to specify a local window around a specific object the user desires to interact with, accelerating computation time by limiting mask computation to pixels within the box. An alert is generated recommending the use of this tool when annotating images larger than 512×512 ; and 3) drop-down menus allowing the user to select line width for marks drawn and to choose between manual annotation and a deep algorithm for intelligent boundary detection.

③ **Mask Upload:** This tool allows the user to upload a semantic mask to provide prior information to the classifier, which can reduce the number of user clicks required for strong performance in our deep interactive model, SGIOS. If no mask is uploaded, the algorithm utilizes a blank mask and produces results equivalent to those from interactive segmentation without a semantic prior map.

④ **Canvas:** The central editor, implemented as an HTML5 canvas element, provides a space for the user to apply tools from the toolkit (rectangle and positive/negative strokes) on a given image. Zooming functions can be used to assist in fine-scale drawing or to provide a broad overview of an image. The user begins interaction with the canvas first by selecting a nested class for an object of interest (from the object panel), then draws positive and negative strokes and clicks the ‘process’ button in the toolkit module. For quick and efficient labeling, these strokes are usually drawn as short lines. Each pixel in a stroke is considered an example of a pixel inside (positive strokes) or outside (negative strokes) the segment of interest. The user-provided marks are sent to the back-end, which detects boundaries and returns a mask to be superimposed over the segment of interest in the graphical front-end. This superimposed mask appears with a white line on boundaries, with inner pixels visualized with a user-selected level of opacity (set using the opacity slider), allowing the user to see both masks and underlying objects. If the mask does not adequately fit the desired segment, the user can provide additional positive and negative strokes and again click the ‘process’ button to generate a new mask.

⑤ **History Panel:** A history panel is integrated into the annotator to help the user recognize and correct errors. Users can view independent steps of annotation in this panel and correct mistakes with undo and redo functions. The history panel also provides a ‘clear’ button to erase all operations and return to the initial state. Three additional

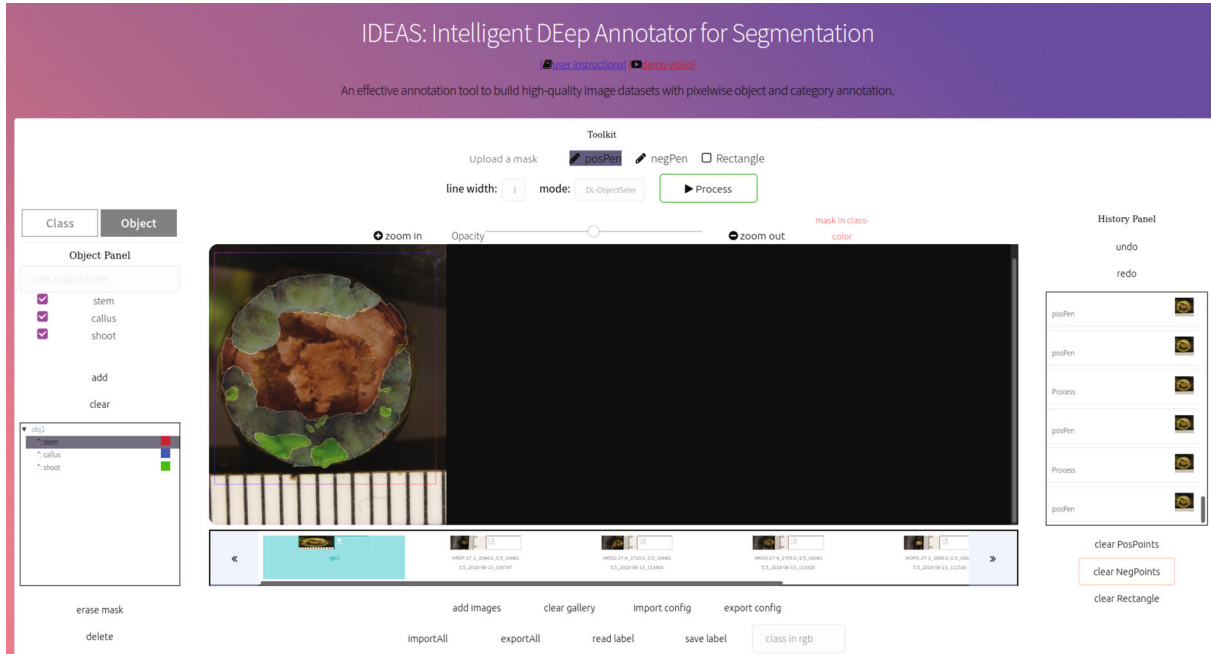


FIGURE 2: Screenshot of IDEAS with one annotated image in a web browser.




FIGURE 3: Front-end modules in IDEAS.

buttons allow the user to undo specific toolkit actions. These include 'clearPositivePoints' and 'clearNegativePoints' buttons to erase all 'posPen' and 'negPen' marks, as well as a 'clearRectangle' button to remove the bounding box set with the rectangle tool.

⑥ **File Gallery:** A file gallery is displayed underneath the canvas, providing an overview of all uploaded images. Images can be uploaded by two methods: first, through the image upload prompt seen when logging into IDEAS, or alternatively, using the 'addImage' button to add images as annotation is ongoing. Using either method, images can be uploaded either independently or in batches and their

thumbnails are displayed in the file gallery. To switch the image being annotated in the 'Canvas' panel, the user can simply click the thumbnail of another image. To remove all uploaded images from the gallery, the user clicks the 'clear gallery' button.

⑦ **Importing and Exporting:** Upon completing annotation for a given image, the user can save the labels (generated mask) as a PNG image file with labels recorded in any of three formats: with segments colored by either their class or object, or both. An example is shown in Figure 4(c)–4(e). The customized class and object hierarchy information can be exported as a .txt file and re-imported during later

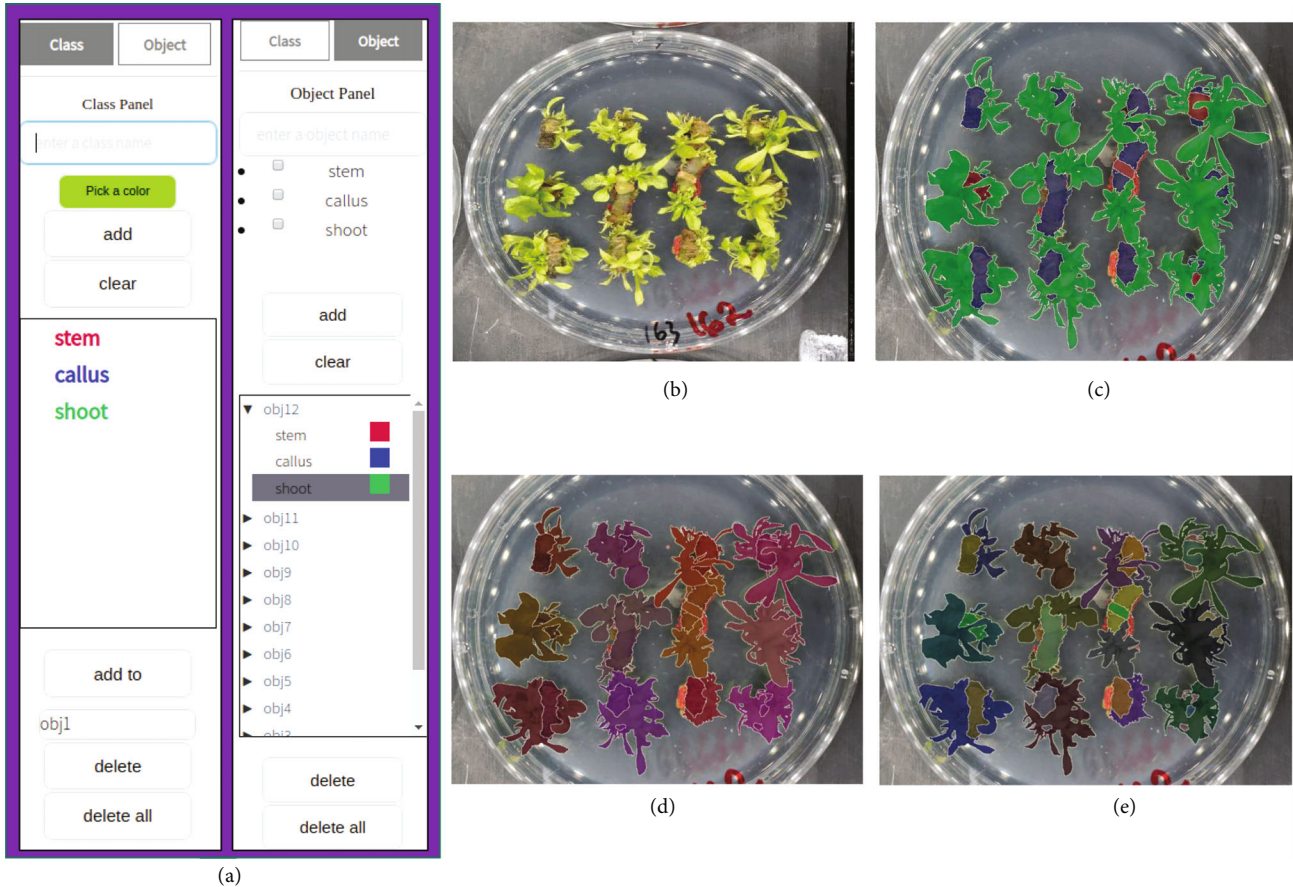


FIGURE 4: One annotation example: (a) Customized configuration of class and object hierarchy. (b) RGB image. (c-e) RGB images with superimposed labels for class (c), object (d), and nested object-class (e).

annotation sessions, a quick and simple means of maintaining consistent settings over the course of an experiment. Furthermore, the user has an option to import or export a .txt file containing the entire annotation environment, including all user actions and generated masks. This feature enables the user to restore the environment to update annotations at a later time.

2.2. IDEAS: Back-End Deep Interactive Algorithm. When the user clicks the ‘process’ button, the front-end collects user inputs including positive and negative strokes, the prior mask, and the bounding box, along with the image and any pre-existing labels, and sends this data to the back-end. The back-end then generates a mask using one of two algorithms, as selected via the ‘mode’ drop-down menu under toolkit. In *manual* mode, the pointer is used to mark every pixel of the desired segment. Annotating whole images by this method can be difficult and time-consuming, as the user is required to draw a label covering the segment of interest and has no assistance in accurately drawing boundaries. Contrarily, the *DL-ObjectSelect* mode which uses SGIOS provides a rapid means of automatically segmenting the selected object when the user provides a small number of positive and negative strokes. We demonstrate that this latter method is fast and efficient, minimizing the required

inputs and time commitment from the user. Thus, the *manual* mode is rarely used in practice.

SGIOS (Semantic-guided Interactive Object Segmentation). Segmentation annotation of images presents a challenge in that every pixel must be labeled for each training image, which could be a time-consuming process, especially in plants where the boundaries usually cannot be expressed with simple geometric shapes. Intelligent algorithms are needed to reduce the human effort required to produce these labels. Here, we propose an interactive segmentation algorithm utilizing a semantic probability map as prior information to guide the segmentation while the user provides inputs of positive and negative clicks. The use of prior information is a strategy to guide the deep convolutional network and provide context for the object of interest, reducing the number of user inputs required for desired performance.

Figure 5 illustrates the pipeline for computing object masks after an object of interest is labeled. The image, user-interaction pairs (positive and negative marks), and semantic information are used as inputs for a fully convolutional network (FCN), which is trained with the bootstrapped cross-entropy loss function to predict a binary mask of the object. On the user interaction guided map, we calculate the pixel value $u_{x,y}$ at location (x,y) as the

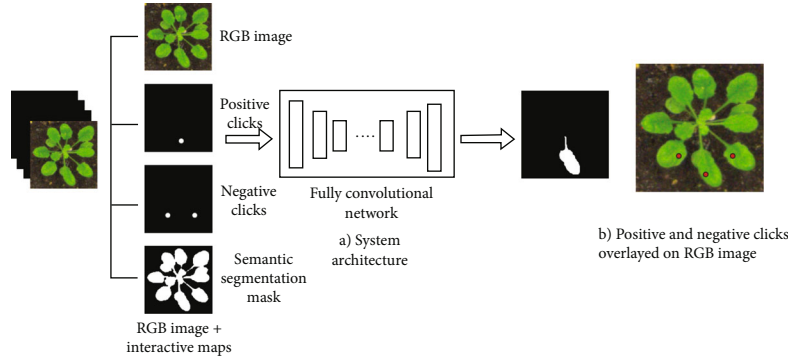


FIGURE 5: Pipeline of the Semantic-Guided Interactive Object Segmentation Algorithm.

minimum Euclidean distance between (x, y) and the set of marks. In other words, given a set of points $P = \{(i, j)\}$, where (i, j) is the point location, $u_{x,y} = \min_{(i,j) \in P} (i-x)^2 + (j-y)^2$.

Most CNN networks [35] for interactive object selection are trained non-iteratively, despite the often iterative nature of refining results with additional clicks by the user. Following the iterative training strategy proposed by ITIS [36], we choose to train our model iteratively. Simulated initial user clicks are updated with the corrective clicks based on the mask obtained in the previous iteration, and this information is sent to the FCN as an input to predict the refined mask. Interactive clicks are further simulated until the desired mask is finally obtained. Below we describe the key components in training our interactive model, SGIOS, with simulated user clicks.

- (i) Initial click sampling: To make our model flexible and not depend on the location or the number of clicks, we randomly sample n positive and m negative clicks until a total of 20 clicks are placed. To sample the positive clicks, we begin by identifying points in the center of the object and randomly sample a point from these. All positive sampled points lie on the object of interest and are d step away from the previously sampled positive points. We first sample negative clicks on other objects close to or touching our object of interest, and the remaining clicks are sampled around the object. When the number of negative clicks sampled is zero, we send a blank mask for the negative clicks and a positive click encoded mask with clicks only on the object of interest.
- (ii) Iterative click sampling on incorrect prediction: Our iterative click sampling strategy is similar to that proposed by ITIS [36], with modifications in the erroneous region selection and click placement on each iteration. Below, we describe the steps for generating the correction clicks while training:
 - (a) The output prediction from the last step is compared with the ground truth mask for the object of interest, and the incorrect pixels are

grouped into multiple clusters using the connected components.

- (b) Instead of selecting the largest cluster on each iteration as is done with ITIS [36], we sample R regions where the clicks should be placed from the distribution of all regions. The distribution of regions is weighted by the probability proportional to the number of pixels in each region.
 - (c) K points are randomly sampled from the R selected erroneous regions, where $K \geq R$ and the total number of clicks is thresholded to 20.
 - (d) The selected clicks are encoded similarly as the initial clicks, using a Gaussian distribution. An example of iterative click sampling is shown in Figure 6.
- (iii) Semantic Prior Maps: The semantic prior maps were generated from an offline semantic segmentation model. For example, in our case study, we used a Deeplabv3+ [37] model trained on plant images for segmenting a image into three categories (stem, callus, and shoot) along with background. The multi-channel output before the final sigmoid layer in the semantic segmentation model was output as the semantic probability maps featuring one channel for each class. Since the semantic prior maps are floating-point numbers on each pixel, we saved them in TIFF format. As a note, the users can use a pre-trained semantic model that was trained for the categories of interest to generate these prior maps. IDEAS does not support online semantic model training or generating TIFF files.
 - (a) The appropriate semantic channel is selected based on the semantic category of the pixels of positive clicks. With the assumption that the user will always place at least one positive click on the object of interest, we obtain the semantic probability map corresponding to the category with maximal probability at the given pixel locations. Using this approach, we can adapt




FIGURE 6: Iterative click sampling on incorrect predictions. (a) Initial prediction when two clicks are placed on the object of interest. (b) Prediction after an additional click is sampled in the right error region. (c) Prediction after one click is sampled in each of the remaining two error regions.

to semantic segmentation tasks with any number of categories. With inclusion of the semantic prior map, we are often able to obtain adequate results for interactive object segmentation with only a single positive click on the object.

- (b) When no semantic prior map is provided, interactive object segmentation proceeds without this information, using a blank mask. This approach is most appropriate when producing a training set for an initial model, which can be used for producing the semantic prior maps for the remaining dataset.

2.3. *IDEAS: Auxiliary Functions.* Beyond the fundamental workflow introduced above, we emphasize the following features in IDEAS to enhance precision, usability and annotation speed:

- (i) Positive/Negative Strokes: We incorporated a feature to allow the user to provide groups of positive or negative markers as strokes as an alternative to providing single markers one click at a time. Rather than only clicking to provide individual points as positive and negative marks, the user can move the mouse while holding the click, drawing strokes to mark multiple pixels along a path. The ability of the user to quickly provide many marks is valuable because segmentation accuracy depends on the number of marks drawn.
- (ii) Quality improvement by adding more ‘posPen’ or ‘negPen’ strokes: We enhanced the semantic-guided interactive object segmentation by adding a feature that enables the user to improve the returned object mask by adding additional positive and negative strokes and again pressing the ‘process’ button. This ‘refinement’ process is integrated into IDEAS. As shown in Figure 7, the accuracy of object mask boundaries improves as additional positive/negative strokes are given by the user. This provides an easy and robust means for the user to produce annotations with desired boundaries precisely labeled.

- (iii) Lock/Unlock the labeled objects: The user can lock or unlock each single labeled nested class for an object by double-clicking it on the object panel. As Figure 8 shows, the labeled objects cannot be modified while locked. Unlocking is necessary before the given object can be edited again. This locking function allows the user to annotate objects without affecting previously drawn labels of other objects. This is particularly useful when small objects are annotated prior to larger, nearby objects. The former objects can be locked as additional objects are labeled, then unlocked and edited again at a later time.

3. Results

In this section, we first showed the efficiency and accuracy of SGIOS on the Pascal VOC [38] dataset and the Computer Vision Problems in Plants Phenotyping (CVPPP) leaf segmentation challenge (LSC) dataset [39, 40]. Then we showed the performance of IDEAS by conducting a preliminary study with the help of an expert biologist, who helped perform annotations using our GUI. Finally, we presented an application of IDEAS for high-throughput phenotyping as part of a GWAS to identify genes that control regeneration in poplar.

3.1. *Semantic-Guided Interactive Object Segmentation. Implementation detail.* We used the model architecture proposed by ITIS [36] with the modification of adding a semantic probability map input. To benchmark this approach, we simulated the user clicks in training using the previously described strategies for initial click sampling and iterative click sampling on incorrect predictions.

Training SGIOS on the augmented Pascal VOC dataset. We used the augmented Pascal VOC [38] dataset with additional annotation from the Semantic Boundaries Dataset (SBD) [41], which consists a total of 10,582 images for around 25,000 objects belonging to 20 different classes in the training subset. The validation subset consists of 3,500 object instances and was used for testing. We fine-tuned the SGIOS model for 50 epochs with a training batch size of 5 and an initial learning rate of $1e-4$. The learning rate varied between the epochs using an exponential decay schedule with a decay rate of 0.9 every 5 epochs and was




FIGURE 7: Demo of the “refinement” process: (a) The label from given ‘posPen’ and ‘negPen’ strokes is errant along the boundary. (b) User adds more strokes and the precision of the predicted labels is highly improved. (c) Finalized label for the object of interest.




FIGURE 8: Demo of the lock/unlock functions. (a) The classes *callus* (blue) and *shoot* (green) have been already labeled. (b) If *callus* and *shoot* are locked, annotation of the new object, *stem* (red) cannot change their labels. (c) If *callus* and *shoot* are unlocked, annotation of *stem* changes their labels.

capped to $5e - 7$. To make the model flexible and avoid over-fitting to the semantic prior map, we randomly reset this prior map to a blank mask containing all zeros in the training stage.

Comparison with the state of the art models. We used the same methodologies used in prior work [36, 42] to evaluate our model’s effectiveness on Pascal VOC.

- (i) *The average number of clicks required to reach the intersection over union (IoU) threshold 85% on Pascal VOC.* For each object, we compute the number of clicks used until obtaining a segmentation with intersection-over-union (IoU) greater than or equal to 85%. If the desired IoU was not reached for any object within 20 clicks, we thresholded the number of clicks to 20. The average number of clicks was computed over all objects in the validation dataset. Table 1(a) reports the result, SGIOS achieves the best number of clicks at 3.1.
- (ii) *The mIoU with different numbers of clicks.* Figure 9 shows the mIoU using different numbers of clicks over the objects in the validation dataset. This shows the benefit of using a semantic prior map is most apparent when a smaller number of clicks is provided. With a single click, SGIOS reaches an IoU of 74.7% and outperforms all other models by a considerable margin, demonstrating that the use of SGIOS reduces the number of clicks required for strong performance, thus mitigating the time cost of annotation. Figures S1 to S3 show several example predictions obtained on Pascal VOC using a single click, two clicks, and three clicks.

Testing on the LSC dataset. We tested our SGIOS model on plant segmentation using the LSC dataset [39, 40]. We used 276 LSC images with object mask labels out

of the 347 images provided in the original dataset. We considered every leaf to be a single object and obtained 3152 object instances from the 276 images. Table 1(b) demonstrates that using the semantic prior map provides a marginal benefit. The LSC dataset is very challenging for instance segmentation because the leaves are in very close proximity and often obscure one another. Figures S4 and S5 show some example predictions obtained on this dataset using two and three clicks, respectively. Please note that we used the model trained on Pascal VOC to test on this dataset, and did not train on this dataset.

Besides, we tested the performance of our model in the absence of semantic prior map by passing a blank mask to the semantic mask channel input. The average number of clicks required to reach the threshold of 85% for this subset is 3.5. This experiment confirms that the model is flexible to the presence of the semantic information and can perform reasonably well even if the semantic information is not present.

3.2. A Study on Annotation Using IDEAS. We conducted a preliminary study with the help of an expert biologist who helped perform annotations using our GUI. The study was performed to evaluate the advantage of using the semantic prior map for interactive image segmentation task. In particular, the study was designed to determine whether the use of a semantic prior map helps to save time to annotate an image.

Design of the study. A set of ten *in vitro* poplar tissue culture images were selected for annotation using IDEAS. These images were randomly divided into two counterbalanced groups as follows:

- (i) Group 1 - The operator annotates five images without using the semantic prior map first and then annotates the same set of images using the semantic prior map.

TABLE 1: (a) Comparison of interactive models to find the number of clicks required to reach a threshold IoU (85%) on Pascal VOC. (b) Number of clicks required by SGIOS to reach a threshold IoU (85%) on the LSC dataset with and without using the semantic prior map.

(a)	
<i>Method</i>	Pascal VOC @85%
Graph cut [43]	15.0
Geodesic matting [44]	14.7
Random walker [45]	11.3
iFCN [35]	6.8
RIS-Net [46]	5.1
DEXTR [47]	4.0
ITIS [36]	3.8
SGIOS (ours)	3.1

(b)	
<i>Method</i>	LSC @85%
SGIOS (without)	16.14
SGIOS (with)	15.40

- (ii) Group 2 - The annotator annotates five images using the semantic prior map first and then annotates the same set of images without using the semantic prior map.

The study took place in four sessions, with two sessions for each group. Each session took close to two hours to complete and was conducted over two days. The scheme was designed to control for the risk of bias from the learning effect when the operator annotates a given image twice and may recall details from the first annotation. Annotations were performed by an expert in plant issue culture. The participant was briefly introduced to how to use the system and was given time to practice using the system for a day before beginning the study. Time taken to annotate each individual explant on each plate was recorded.

Collected images. The images used were sourced from an ongoing study of *in vitro* regeneration in poplar (Figure 4), a departure from our use of *in planta* samples presented for the earlier GWAS demonstration (Figures 1–3). The complete dataset features a total of 1,278 genotypes, each represented by four plates with twelve stem explants. Two plates of each were subjected to an indirect regeneration treatment (callus induction media followed by shoot induction media) and a direct regeneration treatment (shoot induction with no pre-incubation on callus induction media).

Prepare the offline model for computing semantic prior map. We first randomly sampled 100 images from the collected dataset and annotated them on IDEAS with no semantic prior map, in which the default SGIOS model

trained on Pascal VOC is used. We then fine-tuned the model DeepLabv3+ [28] for semantic segmentation. We use the data augmentation strategies (including random rotation, random flipping, random cropping) to improve the performance. The model was fine-tuned for 25,000 iterations using an SGD optimizer with a learning rate of $1e-3$ and batch size of 4. Afterwards, this trained model was used to generate semantic prior maps to assist further annotation using SGIOS, thus enabling accelerated expansion of the training set beyond the initial 100 images.

Timing of annotation. We recorded the time taken to annotate individual explants on the plate with and without semantic prior maps. The study consisted of annotating ten plates where each plate had 12 explants, yielding a total of 120 paired data points. We next performed statistical analysis to inform whether the difference in annotation speed results from an advantage of SGIOS, or is due to random chance. We converted the recorded time into seconds and checked the difference in time for normality using the Shapiro-Wilk test. By observing the histogram of the difference in Figure 10 and the p -value from the Shapiro test equal to $1.868e-12$, we concluded that the data violates the normality assumption. Thus, we conducted a non-parametric test, the Wilcoxon signed-rank test, on the paired data. The Wilcoxon signed-rank test indicated that significantly less time was needed for annotation when semantic information was used, ($V = 1954.5$, p -value = $9.168e-06$, Wilcoxon signed-rank test). Boxplots for both groups are shown in Figure 11. The mean time taken to annotate an explant using a semantic map is 19.99% (27.66 +/-8.15 secs, $n=120$) less than the time taken to annotate an explant without using it.

3.3. A High-Throughput Phenotyping System with IDEAS. We applied our system, IDEAS, as part of a GWAS of *in planta* regeneration in *Populus trichocarpa* (poplar). In this experiment, we applied our phenotyping system toward discovering genes that contribute to variable regeneration response. Plant materials phenotyped included 1,206 clones that represent genetic diversity from California to British Columbia and have been propagated at a field trial location in Corvallis, OR. Plant cuttings from a wide variety of natural genotypes [48] underwent a regeneration-promoting hormone treatment at the cut surface while rooting in water, then were imaged to provide RGB data for analysis. A subset of these images was annotated using IDEAS to build a training set. A semantic segmentation model was trained and then used to segment the remaining, unseen images. Using image processing techniques, we calculated statistics summarizing the segmentation outputs. These statistics were used to represent phenotypes in GWAS, revealing genetic markers associated with these phenotypes.

Phenotype Data Collection and Summarization.

- (i) *Imaging.* The stem cuttings were collected from the field and incubated at 4 ° C for 2-4 weeks. Dormant stem cuttings were placed in 50 mL Falcon tubes with water for five weeks. Once-per-week from the second week onward, the cut tip of each stem was treated with a droplet containing the plant growth




FIGURE 9: Mean IoU vs Number of clicks on Pascal VOC (Thresholded to 20).




FIGURE 10: Histogram of difference in time with and without using the semantic mask, with time shown in seconds.

regulator thidiazuron at 0.5 mg/mL in water. Stem tips were imaged at weekly timepoints by conventional photography using a consumer-grade camera held over plants by a mount. As phenotyping all cuttings at once would have been impractical with limited resources, cuttings were divided into seven groups of up to 400 cuttings each, with groups undergoing this process one after another. This process was performed for up to 400 cuttings at a time. The first group was the only one to include images at the first week and did not include images for fourth and fifth weeks. The third group was missing data

for the fourth week. From the fourth group onward, plants were studied in replicates of two and the computed values used for GWAS are the average of values for each replicate.

- (ii) *Annotating*. We collected 4,896 phenotype images as described, then randomly selected 249 images and annotated them using IDEAS. Figure 12 shows some examples of the annotated labels. Images were annotated to divide plant tissue into segments yet to undergo regeneration (stem) or one of two stages of regenerated tissue (callus or shoot). These




FIGURE 11: Time taken to annotate with and without the semantic mask (in seconds).




FIGURE 12: Examples of annotated ground truth labels overlapped with the RGB image: callus is in red, shoot is in green, and stem is in blue.

annotated images were then randomly split into the validation set and the training set with 24 and 225 images, respectively. In Table 2, we summarized the content of the labeled dataset.

Plant Segmentation. We then trained a semantic segmentation model with the annotated images for computing biological traits. The details of this step are below.

(i) *PSPNet* [29]. PSPNet is among the state-of-the-art architectures used for semantic segmentation. We adopted PSPNet as the deep segmentation model and used *ResNet-50* as its backbone network in our work and employed the following training protocol to train the model to segment the images: We set the initial learning rate to 0.002 and decreased it using a polynomial learning rate schedule with

TABLE 2: Distribution of annotated images.

Items	Analysis result
Distribution over weeks 1 ~ 5	12.8%, 19.0%, 27.7%, 16.9%, 23.6%
Avg. percentage of annotated pixels	Shoot: 0.14%, callus: 2.59%, stem: 1.94%
Avg. number of isolated clusters	Shoot: 0.14, callus: 3.00, stem: 3.72

power at 0.9 in each epoch. The training set was augmented with 1) random scaling to different sizes with scales 0.55 ~ 1.0, 2) flipping horizontally with probability at 0.5, (3) random rotation with angle between $-10^\circ \sim 10^\circ$. We trained the model for 60 epochs with batch size of 4 and weight decay of $4e-4$. Note that PSPNet can be replaced with another semantic segmentation model without affecting the phenotyping system pipeline.

- (ii) *Metric.* To verify the performance on semantic segmentation, we apply the IoU metric, which is computed as $Pred \cap GT / Pred \cup GT$. Here, we report the IoU for each class as well as the mean IoU over all classes.
- (iii) *Experiment 1.* Our research dataset of images from regenerating plants features a relatively small number of classes (stem, callus and shoot) found in a limited number of contexts, unlike highly complex datasets that dominate deep neural network research. Considering this together with reported successes of random forest in segmenting plant tissues [49, 50], we compared the performance of random forests and PSPNet [29] in segmenting our plant images. PSPNet [29] outperformed random forests by a margin of approximately 30% IoU for each class of interest in the validation dataset (Table 3). Furthermore, PSPNet yields a smaller difference between training and validation IoU than was seen with random forests, suggesting that the deep architecture of PSPNet [29] better captures high-level features and avoids overfitting.
- (iv) *Experiment 2.* To gain insight into an appropriate number of training samples for use in our workflow, we explored the relationship between the number of training examples and the quality of segmentation. We trained six different models using 20, 40, 60, 80, 100, and 120 images randomly selected from our training set and tested each model’s performance using the same validation set of 20 images. As shown in Figure 13, the Mean IoU increased significantly as the training sample size increased to 40, 60, and 80. As the sample size increases to 100 and further, additional increases in Mean IoU become markedly diminished but may remain significant and valuable for applications that benefit from high precision. In summary, the deep learning models’ performance relies heavily on the size of the training set, and the inclusion of additional images can almost always be expected to improve the quality of segmentation

results. These results demonstrate the importance of a substantial training set for semantic segmentation, underscoring the value in developing efficient and intelligent interactive image segmentation algorithms and annotation tools such as IDEAS.

Biological traits. After the semantic segmentation model was trained on the collected training set, we used it to predict labels given unseen images. Next, we calculate the relative area, which represents the development of specific tissues in each image, and used these in our downstream models for GWAS. Computation of relative area is shown in eq. (2):

$$\text{Area}_c = \sum_{x,y} 1_{\{c\}}(I(x,y)) \quad (1)$$

$$\text{RArea}_c = \frac{\text{Area}_c}{\sum_c \text{Area}_c} \quad (2)$$

Where I is the semantic segmentation from the PSPNet [29] and c is the category of the tissue for which relative area is being calculated.

Association mapping. We present results from GWAS for a key trait of interest, the area of shoot grown at the final timepoint (fifth week). In this analysis, we apply a mixed linear model approach that is commonly used for GWAS of continuous phenotypes and assumes normality of residuals. To avoid a severe violation of this assumption, genotypes labeled by the model as having zero values (no shoot) were dropped and the remaining phenotype data underwent a natural logarithm transformation. The population analyzed in this case study consists of 326 clones for which genotype data was available and phenotype data is nonzero.

Genotype data used for analysis was obtained from the Bioenergy Science Center at Oak Ridge National Laboratory (<https://cbi.ornl.gov/gwas-dataset/>) and filtered for minor allele frequency and missing rates using PLINK [51]. The filtered dataset used for analysis features approximately 5.3 million single-nucleotide polymorphisms (SNPs), each of which are polymorphic in at least 5% of genotypes and non-missing in at least 90%.

We employed Genome-wide Efficient Mixed Model Association (GEMMA) [52] as a GWAS method. Mixed models are built for each SNP, explaining the phenotype as a function of the given SNP and covariates. To control for potential confounding factors of cutting size and population stratification, we included covariates of stem diameter, a kinship matrix and three principle components derived from SNPs. Multiple-testing correction performed using the

TABLE 3: Segmentation IoU score (in %) on the training and validation dataset.

Method	Dataset	Background	Stem	Callus	Shoot	mIOU
RF [50]	train.	94.62%	69.83%	38.75%	40.82%	61.00%
RF [50]	val.	93.83%	57.67%	20.92%	39.36%	52.95%
PSPNet [29]	train.	99%.35	89.76%	73.96%	74.36%	84.36%
PSPNet [29]	val.	99.36%	88.14%	61.95%	67.40%	79.21%




FIGURE 13: The Mean IoU for different number of training examples.

Benjamini-Hochberg method was used to calculate a p -value threshold with a false-discovery rate of 10% [53]. Putative associations with p -values less than or equal to this threshold are reported in Figure 14 and Table 4.

Interrogation of putative associations. To determine genes implicated by these genetic markers, we consulted a reference *Populus trichocarpa* genome [54] annotation (version 3.1 available at <https://phytozome-next.jgi.doe.gov>) to determine relative positions of these genetic markers to gene transcripts. To identify homologs of implicated genes in better-characterized plant species, and thus gain insight into the function of these genes, we referred to a database of Smith-Waterman alignments of predicted peptides between *P. trichocarpa* and the model plant *Arabidopsis thaliana* (<https://phytozome-next.jgi.doe.gov>).

The two candidate genes implicated by lowest p -values are related to *Arabidopsis* genes with known roles in regeneration (Figure 14 and Table 4). *CHALCONE SYNTHASE* (*CHS*) encodes a protein essential for a rate-limiting step of the biosynthesis of anthocyanin, which may influence shoot regeneration by dual mechanisms of auxin transport regulation [55] and light stress protection. Loss-of-function mutations of *CHS* are reported to be deficient in shoot regeneration, with a light-dependent effect [56]. *SALT-AND-DROUGHT-INDUCED RING FINGER 1* (*SDIR1*) is an E3 ubiquitin ligase critical for regulation of protein degradation downstream of the hormone abscisic acid.

Loss-of-function and overexpression lines of *SDIR1* display enhanced and inhibited levels of *in vitro* seedling germination, respectively [57]. Supplementation of *in vitro* media with abscisic acid has been reported to enhance *in vitro* regeneration in diverse plants, particularly via the route of somatic embryogenesis [58].

Our third candidate (Potri.005G004700) is a gene of unknown function that invites further review by biologists aiming to characterize the genetic basis of plant regeneration. Smith-Waterman alignments reveal no *Arabidopsis* homolog of this gene.

Table 4 shows quantitative trait loci (QTLs) with their p -values calculated by GEMMA. Accession IDs of nearby gene candidates are reported, along with the distance of the QTL from each gene. The most similar *Arabidopsis* homolog is shown for each candidate gene, along with the similarity as calculated by Smith-Waterman alignment. Chr14:2487581 appears between two genes and both of these possible causative genes are listed. Potri.005G004700 is implicated by two significant QTLs.

4. Discussion

While supervised deep learning has demonstrated great power for diverse image tasks, the potential for this technology to be applied to specific tasks is constrained by the abilities of users to produce sufficient training data. Moreover,




FIGURE 14: This Manhattan plot shows the significance of genetic markers in models explaining the shoot regeneration phenotype as a function of each marker. The blue line represents a threshold for a false discovery rate of 10%, calculated by the Benjamini-Hochburg method [53].

TABLE 4: QTLs, implicated poplar gene candidates and related Arabidopsis homologs.

QTL	p-value	Gene candidate	Proximity (bp)	Arabidopsis gene	Similarity
Chr5:284409	5.1e-08	Potri.005G004700	692	None	N/A
Chr5:284964	6.5e-08	Potri.005G004700	137	None	N/A
Chr9:4459368	4.7e-08	Potri.009G034700	1,283	AT3G55530: SALT- AND DROUGHT-INDUCED RING FINGER1	85.4%
Chr12:9509286	1.1e-07	Potri.012G070801	8,961	None	N/A
Chr12:15334435	4.5e-08	Potri.012G138800	1,155	AT5G13930: CHALCONE SYNTHASE	83.0%
Chr14:2487581	5.5e-08	Potri.014G029200; Potri.014G029300	1,924; 2,556	AT1G07700: Thioredoxin superfamily protein; AT5G42890: STEROL CARRIER PROTEIN 2	88.7%; 81.3%

for plant biology and other fields in which researchers commonly lack a computer science background, an essential prerequisite for the broad dissemination of deep learning is the development of generalizable workflows with user-friendly, high-level interfaces for annotating training data and deploying models. This combination of challenges presents a need for innovation both in algorithm and interface design. To this end, we created a novel annotation system tailored to the needs of plant biologists engaged in phenotyping of plant tissues during any stage of growth or regeneration.

IDEAS enables high-throughput measurement of plant tissues, and can be easily applied to diverse tissue types by user annotation of features of interest. A major advantage of IDEAS is the capability to accurately define tissue object boundaries that are highly complex and thus ignored or misidentified by common annotation methods involving polygons or bounding boxes. Manual pixel-scale annotation can produce highly precise boundaries, but the labor-intensiveness of this process presents an obstacle to generating sizable training sets. Our semantic-guided interactive object segmentation algorithm provides high precision at high speed, delivering substantial reductions to the labor cost of generating high-quality training sets for complex traits.

Worth emphasizing is that interactive segmentation is a more complex problem than semantic segmentation, especially because user annotation involves separate instances of a given type of object (class). Thus, the potential for the

semantic prior maps to improve interactive segmentation has several noteworthy limitations. First, in cases where there exist adjacent or touching objects of the same class, the separation of these objects by the user is particularly difficult, and may not be improved by the use of a semantic prior map lacking instance information. Second, the performance of semantic segmentation tends to be better for larger objects than for smaller objects. Hence, for small objects, a semantic prior map may provide little or no benefit to the system. These limitations considered, the effective use of a semantic prior map to guide interactive segmentation must include an ability to apply the prior only when it offers a performance benefit. In a preliminary study, we examined the use of prior information and conducted a statistical analysis to determine whether using a semantic prior map helps to accelerate annotation. The semantic prior map was found to confer a statistically significant increase in annotation speed of approximately 20 percent. Finally, we demonstrated our system using the case study of a GWAS of *in planta* regeneration in poplar. Top associations from this GWAS implicate genes with Arabidopsis homologs that have known roles in regeneration pathways. Regeneration entails hormone-driven processes [59] involving hormones including auxin and abscisic acid, which are regulated by genes including CHS [55] and SDIR1 [57, 58], respectively. The agreement between our results and established models of these genes' roles in regeneration suggests that our phenotyping system is effective in capturing the shoot regeneration phenotype in our case study. Furthermore, the

appearance of an unknown gene as a putative association provides an example of how this system can contribute to genetic discovery.

5. Conclusion

We developed a robust deep segmentation phenotyping system utilizing a web-based annotator, IDEAS, for generating ground truth datasets. Using the semantic-guided interactive object segmentation backend, IDEAS provides an accelerated means of labeling objects at pixel-scale with precise boundaries. Using labels generated by the annotator, researchers can train a deep model for semantic segmentation, deploy the model to make predictions over a large dataset and compute statistics summarizing segments of biological interest. Downstream, GWAS revealed genetic markers associated with traits phenotyped by computer vision. Our system can be used by plant biologists who are interested in complex segmentation-based traits, for which generation of large training sets may otherwise be time-prohibitive.

Conflicts of Interest

There is conflict of interest regarding this article with Apple Inc., Microsoft Corporation Inc., and Parametrix.ai Inc..

Authors' Contributions

JY and DK designed the image analysis pipeline and experimental procedures and performed evaluation of the ideas relating to the sections involving computer vision techniques. MN provided important help in the experimental section on GWAS study and writing of this paper. ZZ, NGK, and NAD assisted in the GUI design. EP assisted in collecting data with IDEAS. CM, SS, and FL supervised the project and designed the research. Jialin Yuan and Damanpreet Kaur contributed equally to this work.

Acknowledgments

General: We greatly appreciate the support by Drs. Jerry Tuskan and Wellington Muchero of BESC for their collaboration and advice during research grant development and project execution. We thank the members of the GREAT TREES (Genetic Research on Engineering and Advanced Transformation of Trees) Research Cooperative at Oregon State University for its long term investment in our transformation and regeneration studies. We thank Middleton Spectral Vision (Wisconsin) for their customized imaging system used to capture regeneration data, and high quality imaging system support.

Poplar SNP dataset: Support for the Poplar GWAS dataset is provided by the U.S. Department of Energy, Office of Science Biological and Environmental Research (BER) via the Bioenergy Science Center (BESC) under Contract No. DE-PS02-06ER64304. The Poplar GWAS Project used resources of the Oak Ridge Leadership Computing Facility and the Compute and Data Environment for Science at

Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. We thank the National Science Foundation Plant Genome Research Program for funding "Analysis of Genes Affecting Plant Regeneration and Transformation in Poplar," IOS #1546900, NSF IIS #1911232, and IIS #1909038. The research was also partially supported by Amazon Research Award, DARPA Contract N66001-19-2-4035, HR001120C0011 and a gift from Kuaishou Inc..

Supplementary Materials

Figure S1. Example predictions on Pascal VOC dataset on a single click. Figure S2. Example predictions on Pascal VOC dataset on two clicks. Figure S3. Example predictions on Pascal VOC dataset on three clicks. Figure S4. Example predictions on Leaf Segmentation Challenge (LSC) dataset on two clicks. Figure S5. Example predictions on Leaf Segmentation Challenge (LSC) dataset on three clicks (*Supplementary Materials*)

References

- [1] 1000 Genomes Project Consortium, "A global reference for human genetic variation," *Nature*, vol. 526, no. 7571, pp. 68–74, 2015.
- [2] G. Tuskan, G. Slavov, S. DiFazio et al., "Populus resequencing: towards genome-wide association studies," *BMC Proceedings*, vol. 5, no. 7, 2011.
- [3] W. Yang, Z. Guo, C. Huang et al., "Combining high-throughput phenotyping and genome-wide association studies to reveal natural genetic variation in rice," *Nature Communications*, vol. 5, no. 1, p. 5087, 2014.
- [4] M. Rahaman, D. Chen, Z. Gillani, C. Klukas, and M. Chen, "Advanced phenotyping and phenotype data analysis for the study of plant growth and development," *Frontiers in Plant Science*, vol. 6, p. 619, 2015.
- [5] D.-C. Trinh, J. Lavenus, T. Goh et al., "PUCHI regulates very long chain fatty acid biosynthesis during lateral root and callus formation," *Proceedings of the National Academy of Sciences*, vol. 116, no. 28, pp. 14325–14330, 2019.
- [6] X. Cao, L. Xing, H. He, and X. Zhang, "Views on GWAS statistical analysis," *Bioinformatics*, vol. 16, no. 5, pp. 393–397, 2020.
- [7] W. D. Dupont and W. D. Plummer Jr., "Power and sample size calculations for studies involving linear regression," *Controlled Clinical Trials*, vol. 19, no. 6, pp. 589–601, 1998.
- [8] A. D. Skol, L. J. Scott, G. R. Abecasis, and M. Boehnke, "Joint analysis is more efficient than replication-based analysis for two-stage genome-wide association studies," *Nature Genetics*, vol. 38, no. 2, pp. 209–213, 2006.
- [9] Z. Li, R. Guo, M. Li, Y. Chen, and G. Li, "A review of computer vision technologies for plant phenotyping," *Computers and Electronics in Agriculture*, vol. 176, no. 6, article 105672, 2020.
- [10] J. A. Atkinson, M. P. Pound, M. J. Bennett, and D. M. Wells, "Uncovering the hidden half of plants using new advances in root phenotyping," *Current Opinion in Biotechnology*, vol. 55, pp. 1–8, 2019.

- [11] G. A. Tuskan, R. Mewalal, L. E. Gunter et al., “Defining the genetic components of callus formation: a GWAS approach,” *PLoS One*, vol. 13, no. 8, article e0202519, 2018.
- [12] Q. Zhang, S. Zhifang, Y. Guo, S. Zhang, L. Jiang, and W. Rongling, “Genome-wide association studies of callus differentiation for the desert tree, *Populus euphratica*,” *Tree Physiology*, vol. 40, no. 12, pp. 1762–1777, 2020.
- [13] R. Lardon, E. Wijnker, J. Keurentjes, and D. Geelen, “The genetic framework of shoot regeneration in *Arabidopsis* comprises master regulators and conditional fine-tuning factors,” *Communications Biology*, vol. 3, no. 1, pp. 1–13, 2020.
- [14] L. Keren, M. Bosse, D. Marquez et al., “A structured tumor-immune microenvironment in triple negative breast cancer revealed by multiplexed ion beam imaging,” *Cell*, vol. 174, no. 6, pp. 1373–1387.e19, 2018.
- [15] S. E. Ahmed Raza, K. Abduljabbar, M. Jamal-Hanjani et al., “Deconvolving convolutional neural network for cell detection,” in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 891–894, Venice, Italy, 2019.
- [16] S. T. Namin, M. Esmailzadeh, M. Najafi, T. B. Brown, and J. O. Borevitz, “Deep phenotyping: deep learning for temporal phenotype/genotype classification,” *Plant Methods*, vol. 14, no. 1, pp. 1–14, 2018.
- [17] J. R. Ubbens and I. Stavness, “Deep plant phenomics: a deep learning platform for complex plant phenotyping tasks,” *Frontiers*, vol. 8, article 1190, 2017.
- [18] E. David, S. Madec, P. Sadeghi-Tehran et al., “Global Wheat Head Detection (GWHD) dataset: a large and diverse dataset of high-resolution RGB-labelled images to develop and benchmark wheat head detection methods,” *Plant Phenomics*, vol. 2020, article 3521852, 12 pages, 2020.
- [19] Y. Jiang and C. Li, “Convolutional neural networks for image-based high-throughput plant phenotyping: a review,” *Plant Phenomics*, vol. 2020, article 4152816, 22 pages, 2020.
- [20] F. Zhuang, Z. Qi, K. Duan et al., “A comprehensive survey on transfer learning,” *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2020.
- [21] C. Shorten and T. M. Khoshgoftaar, “A survey on image data augmentation for deep learning,” *Journal of Big Data*, vol. 6, no. 1, p. 60, 2019.
- [22] M. Amgad, H. Elfandy, H. Hussein et al., “Structured crowdsourcing enables convolutional segmentation of histology images,” *Bioinformatics*, vol. 35, no. 18, pp. 3461–3467, 2019.
- [23] B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman, “LabelMe: a database and web-based tool for image annotation,” *International Journal of Computer Vision*, vol. 77, no. 1–3, pp. 157–173, 2008.
- [24] J. Brooks, “COCO Annotator,” 2019, <https://github.com/jsbrooks/coco-annotator/>.
- [25] A. Dutta and A. Zisserman, “The VIA Annotation Software for Images, Audio and Video,” in *In Proceedings of the 27th ACM International Conference on Multimedia*, pp. 2276–2279, Nice France, 2019.
- [26] “VGG Image Annotator,” <http://www.robots.ox.ac.uk/vgg/software/via/via-1.0.6.html>.
- [27] “Sergi Caelles DEXTR AnnoTool,” <https://github.com/karanshr/DEXTR-AnnoTool>.
- [28] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Computer Vision – ECCV 2018. ECCV 2018*, vol. 11211 of Lecture Notes in Computer Science, pp. 801–818, Springer, Cham.
- [29] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2881–2890, Honolulu, HI, USA, 2017.
- [30] G. Lin, A. Milan, C. Shen, and I. Reid, “Refinenet: Multi-path refinement networks for high-resolution semantic segmentation,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1925–1934, Honolulu, HI, USA, 2017.
- [31] T. V. Nguyen and L. Liu, “Salient object detection with semantic priors,” 2017, <https://arxiv.org/abs/1705.08207>.
- [32] G. M. Garcia, F. Husain, H. Schulz, S. Frintrop, C. Torras, and S. Behnke, “Semantic segmentation priors for object discovery,” in *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 549–554, Cancun, Mexico, 2016.
- [33] S. Y. Bao, M. Chandraker, Y. Lin, and S. Savarese, “Dense object reconstruction with semantic priors,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1264–1271, Portland Oregon, USA, 2013.
- [34] Z. Kamranian, A. R. Naghsh Nilchi, A. Monadjemi, and N. Navab, “Iterative algorithm for interactive co-segmentation using semantic information propagation,” *Applied Intelligence*, vol. 48, no. 12, pp. 5019–5036, 2018.
- [35] N. Xu, B. Price, S. Cohen, J. Yang, and T. S. Huang, “Deep interactive object selection,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 373–381, Las Vegas, NV, USA, 2016.
- [36] S. Mahadevan, P. Voigtlaender, and B. Leibe, “Iteratively Trained Interactive Segmentation,” 2018, <https://arxiv.org/abs/1805.04398>.
- [37] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2017.
- [38] M. Everingham, S. M. A. Eslami, L. van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes challenge: a retrospective,” *International Journal of Computer Vision*, vol. 111, no. 1, pp. 98–136, 2015.
- [39] H. Scharr, M. Minervini, A. Fischbach, and S. A. Tsafaris, “Annotated image datasets of rosette plants,” in *European Conference on Computer Vision*, pp. 6–12, Zürich, Suisse, 2014.
- [40] M. Minervini, M. M. Abdelsamea, and S. A. Tsafaris, “Image-based plant phenotyping with incremental learning and active contours,” *Ecological Informatics*, vol. 23, pp. 35–48, 2014.
- [41] B. Hariharan, P. Arbelaez, L. Bourdev, S. Maji, and J. Malik, “Semantic contours from inverse detectors,” in *2011 International Conference on Computer Vision*, Barcelona, Spain, 2011.
- [42] S. Majumder and A. Yao, “Content-aware multi-level guidance for interactive instance segmentation,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11602–11611, Long Beach, CA, USA, 2019.
- [43] Y. Y. Boykov and M.-P. Jolly, “Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images,” in *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, Vancouver, BC, Canada, 2001.
- [44] X. Bai and G. Sapiro, “Geodesic matting: A framework for fast interactive image and video segmentation and matting,”

- International Journal of Computer Vision*, vol. 82, no. 2, pp. 113–132, 2009.
- [45] L. Grady, “Random walks for image segmentation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 28, no. 11, pp. 1768–1783, 2006.
- [46] J. H. Liew, Y. Wei, W. Xiong, S.-H. Ong, and J. Feng, “Regional interactive image segmentation networks,” in *2017 IEEE international conference on computer vision (ICCV)*, pp. 2746–2754, Venice, Italy, 2017.
- [47] K.-K. Maninis, S. Caelles, J. Pont-Tuset, and L. Van Gool, “Deep extreme cut: From extreme points to object segmentation,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 616–625, Salt Lake City, UT, USA, 2018.
- [48] H. B. Chhetri, D. Macaya-Sanz, D. Kainer et al., “Multitrait genome-wide association analysis of *Populus trichocarpa* identifies key polymorphisms controlling morphological and physiological traits,” *New Phytologist*, vol. 223, no. 1, pp. 293–309, 2019.
- [49] U. Lee, S. Chang, G. A. Putra, H. Kim, and D. H. Kim, “An automated, high-throughput plant phenotyping system using machine learning-based plant segmentation and image analysis,” *PLoS One*, vol. 13, no. 4, article e0196615, 2018.
- [50] M. M. Rahaman, M. A. Ahsan, and M. Chen, “Data-mining techniques for image-based plant phenotypic traits identification and classification,” *Scientific Reports*, vol. 9, no. 1, p. 19526, 2019.
- [51] S. Purcell, B. Neale, K. Todd-Brown et al., “PLINK: A Tool Set for Whole-Genome Association and Population-Based Linkage Analyses,” *The American Journal of Human Genetics*, vol. 81, no. 3, pp. 559–575, 2007.
- [52] X. Zhou and M. Stephens, “Genome-wide efficient mixed-model analysis for association studies,” *Nature Genetics*, vol. 44, no. 7, pp. 821–824, 2012.
- [53] Y. Benjamini and Y. Hochberg, “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.
- [54] G. A. Tuskan, S. Difazio, S. Jansson et al., “The genome of black cottonwood, *Populus trichocarpa* (Torr. & Gray),” *Science*, vol. 313, no. 5793, pp. 1596–1604, 2006.
- [55] D. E. Brown, A. M. Rashotte, A. S. Murphy et al., “Flavonoids act as negative regulators of auxin transport In Vivo in Arabidopsis,” *Plant Physiology*, vol. 126, no. 2, pp. 524–535, 2001.
- [56] B. Nameth, S. J. Dinka, S. P. Chatfield et al., “The shoot regeneration capacity of excised Arabidopsis cotyledons is established during the initial hours after injury and is modulated by a complex genetic network of light signalling,” *Plant, Cell & Environment*, vol. 36, no. 1, pp. 68–86, 2013.
- [57] Y. Zhang, C. Yang, Y. Li et al., “SDIR1 is a RING finger E3 ligase that positively regulates stress-responsive abscisic acid signaling in Arabidopsis,” *The Plant Cell*, vol. 19, no. 6, pp. 1912–1929, 2007.
- [58] M. K. Rai, N. S. Shekhawat, A. K. Gupta, M. Phulwaria, K. Ram, and U. Jaiswal, “The role of abscisic acid in plant tissue culture: a review of recent progress,” *Plant Cell, Tissue and Organ Culture*, vol. 106, no. 2, pp. 179–190, 2011.
- [59] K. Sugimoto, H. Temman, S. Kadokura, and S. Matsunaga, “To regenerate or not to regenerate: factors that drive plant regeneration,” *Current Opinion in Plant Biology*, vol. 47, pp. 138–150, 2019.