Flower Classification for Kaziga 2024

- Motivation: Understanding the temporal dynamics of reproduction in cowpea may yield important breeding traits. However, manually counting flowers and other reproductive structures is labor intensive. So, we want to create a high throughput option that can capture these reproductive dynamics using imaging alone.
- Image data:
  - 40 sensing dates in 2024 for two cowpea beds with 10 plots each, representing two replicates of ten genotypes, with ~4 sensing dates per week over 10 weeks (July 1 to September 6)
  - Frames were extracted at 10 fps from 4K video taken from an iPhone mounted on a selfie stick (Kaziga)
  - Those frames were then stitched using AgRowStitch and plots were manually cropped
  - Each plot image was sliced into 640x640 px images
- Ground truth data:
  - 11 reproductive ground truth dates across the 20 sensing plots (July 11 to September 6)
  - Weekly data resolution where the counted classes were closed (yet to open) flowers, open flowers, senescent flowers (opened in the last two days), and pods (any stage)
  - Each peduncle that contained at least one of the counted classes was counted and the number of each class on the peduncle was recorded
- Considerations:
  - The uneven camera angles, the stitching process, and wind cause distortion, blurring, and a loss of some of the field of view
  - Ground truth data was not restricted by occlusion or field of view, while we expect that uneven field of view and occlusion can cause genotype-specific bias in detection probability
  - Since there were three different people doing the ground truth counts, it is likely that there is variation in the count fidelity and classification boundaries, so a degree of merging should make the results more reliable
  - Preliminary analysis suggests that occlusion is strong enough that we cannot expect to get count data that matches ground truth data and that there is genotypic bias in detectability – so we are only concerned with counts because we need sufficient counts to reproduce temporal dynamics, after that the counts should be normalized.
  - Preliminary analysis also suggested that pod and closed flower buds were difficult to detect reliably, while open flowers were relatively easy to detect, but the proportion of open flowers that was visible from imaging was relatively low due to either occlusion or restricted field of view
- Goal: Create a classification scheme that balances ease of detectability with a high number of detections so that the detections can reproduce the reproductive temporal dynamics at the genotype-scale.

Classification Scheme

- Since the process of flower opening and senescence is continuous, multiple discrete classes are difficult to reliably label. Instead, we use a single “Flower” class that encompasses a period where the color of the flower is easy to identify, which corresponds roughly to the day before a flower opens to the two days after it opens.
- This allows us to retain a tight temporal window that can be compared to the open (petals visible) and senescent (still “attached” to peduncle) ground truth counts. Importantly, we do not simply count anything that doesn’t look like a leaf because we would include senescent flowers that may have opened over a week ago, peduncles that are barren, or immature buds that are both hard to see and far from opening.
- We grade potential labels based on color, shape, and context and grant a positive ID to a flower when we are “fairly certain” of the label. It’s all hand wavy, but at the end of the day labeling is hand wavy anyway, so welcome to the club to any poor unfortunate soul using this guide.

Development and Target

- In the figures to the right, featured in the flowering descriptions made by the Pre-breeding and Genomics Team, we can see the stages of cowpea flower development.
- For our purposes, **we are targeting the stages AFTER 1F-IX and just past the final image**. This includes the stages when:
  - The banner petals change color (from green to a more white or yellow) and the bud is large and engorged
  - The petals become visible as the flower begins to open
  - The flower is fully open (C in the figure)
  - The flower closes again (beige)
  - The flower senesces (becomes yellow to dark yellow, as in the final image)
  - The flower begins to dry out and wrinkle (yellow, brown)
  - The pod has started growing and the senesced flower is being separated from the peduncle (senesced flower will now have white at the base)
  - **If this pod is visible and more than 1 mm in diameter or more than 1 cm of the pod is visible, the flower is too old to be included.**

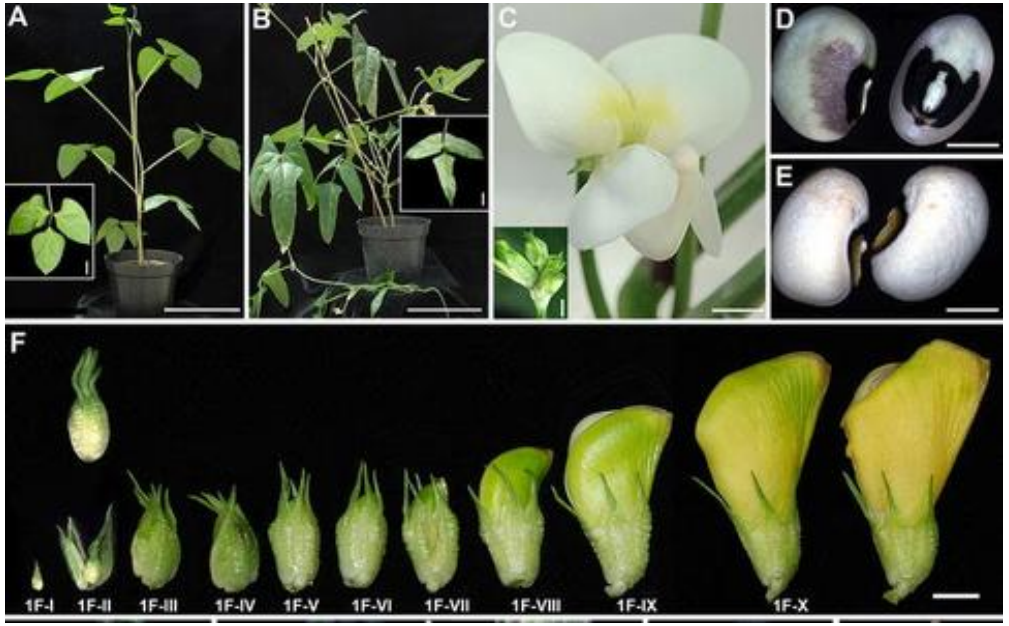

Topmost red arrows indicate that 1F-IX is the stage just prior to panel C and 1F-X is just after. Images adapted from Salinas-Gamboa et al., 2016 : 10.1007/s00497-015-0273-3

Visual Guide

Top row: The first image has deeply green banner petals, meaning it does not count even though we can just make out the petals, the other examples do count because banner petals have changed to more white, yellow, or purple.
Middle row: The banner petals have opened, revealing the petals. Flower color is always bright, but may be yellow, pink, or purple in addition to white. When viewed from the back, the banner petal color (white, lime green, pinkish) and shape can lead to an ID.
Bottom row: The flower has closed again and is in various stages of drying. The color can range from beige to gold, to brownish yellow. In the last two images the flower is too dry to count and the pod is too long and thick to count, while in the third to last image, the pod is sufficiently small and the flower has maintained its shape, so it counts.

Color

- Due to variation in flower color across genotypes and color variation due to lighting, it is difficult to make a single consistent color guide for classification. However, we use the heuristic that green objects are NOT counted.
- In the early stages of the flower bud, only the sepals are visible and these are a light green/yellow. These are both too small to count and the wrong color.
- As the flower bud develops, it gains its characteristic shape and a darker green color. These also do not count.
- The banner petals eventually change to a lighter color, though the exact color is genotype-specific. This is the first stage where we can get a positive ID. These colors include white, white-purple, white-pink, lime green, green-white, and yellow.
- When the flower opens, the petals are visible from some angles and these include white, white-pink, purple, and yellow-white. These are the easiest color IDs.
- When the flower closes and senesces, it will revert to the characteristic shape and start to dry. The colors for a positive ID are beige, yellow, gold, and yellow-brown. The flower will turn white starting at the base as it becomes dry. A senescent flower can have a positive color ID only if it is less than ¼ white.
- Color alone is insufficient for a positive ID because specular reflection can cause white and yellow areas on other structures. Leaves and pods can have overlapping colors with flowers because some flowers maintain a lime green or light-yellow color in their banner petals when open. This overlaps with some healthy leaves, senescing leaves, immature (green) pods, and mature (dried) pods.

Shape

- Since the viewing angle of flowers is variable, there are various shapes that can help with a positive ID. Refer to the Visual Guide.
- First, for both developmental and modeling reasons, we exclude objects that are too small. This should rule out flower buds that are still too immature.
- Danger: Closed leaves look very similar to closed flower buds, the tips of dried pods look very similar to senescent flowers and flowers that are opening soon, the top view of a large immature pod looks very similar to open banner petals.
- Flowers that have not opened yet have the following shapes:
  - Top view: long with a rounded tip, the visible part of the banner petals should be at least the same length as the sepals.
  - Side view: for younger IDs, the shape is similar to an axe head while the easier IDs have “the characteristic shape”.
- Flowers that are opening or are fully open have the following shapes:
  - Top view: the banner petals are no longer closed, the angle between them when fully open is 180 degrees, but the petals are often visible starting at ~45 degrees.
  - Front view: the petals are visible and as the flower opens it transitions from a clam-like shape to the full 3D shape with the banner petals open and petals extending roughly 90 degrees from the plane of the banner petals, making a chair-like shape.
  - Side view: depending on the angle and stage of opening, the banner petals should have “the characteristic shape” with perhaps parts of the petals or the opposite banner petal visible. If viewed at 90 degrees, the petals will be visible while the banner petals may be hard to see.
  - Back view: since the petal color will not be visible, these can be hard IDs, but the shape of the banner petals as they open will range from a partially open book to two clear lobes with a dividing line, reminiscent of a peach.
- Senescing flowers have the following shapes:
  - Top view: depending on the stage, this can look similar to the yet to open flowers, but the tip is generally now at a point and it may be more of a torpedo shape, there can also be a little wiggle now rather than symmetrical. The sepals should still be visible. Small amount of pod may also be visible in late stages.
  - Side view: can range from “the characteristic shape” when it is fresh to more of a bent elbow shape as the tip of the banner petals dries and shrivels, creating a right angle with the midsection of the banner petals (now concealing a young pod). When the shape is shriveled like a shrimp or worse, it EXCLUDES a positive ID.

Beautiful shape examples

“The Characteristic Shape” ™

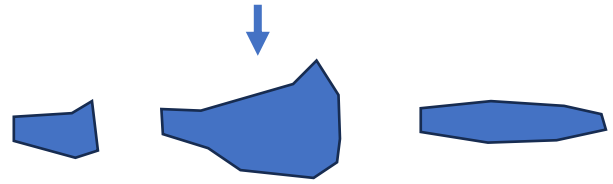

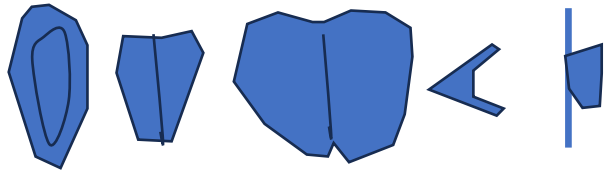

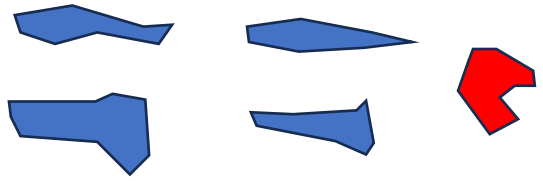

Context

- Since the images may be blurry, we have a poor viewing angle of the flower, or the flower is occluded, it can be difficult to make positive IDs.
- We can use context to help us rule in or rule out more ambiguous labels.
- If the structure is very small and without clear edges, we can exclude the structure because we assume that this is too difficult for detection.
- With one exception, all of our positive IDs should be associated with a specific peduncle. Peduncles can be identified by their long slender shape and parallel lines. They often have a white/redish/pinkish tip where reproductive structures attach. Peduncles can be also be identified by the presence of other reproductive structures connected to the tip. Long runner shoots or normal stems can be easily confused with peduncles.
- The exception to the peduncle aiding a positive ID is when a senescent flower has fallen off the pod but it has retained its color (beige, yellow, gold, less than ¼ white from the base) and/or shape. We count these senescent flowers and assume that falling off the pod is a stochastic event to help with sufficient detections and model accuracy. My expert analysis confirms that gently kicking a plant can cause senescent flowers to fall off that would otherwise be “attached” to peduncles.
- Occluded flowers or flowers lurking in shadows where the color is sufficient for a positive ID (bright white, purple, pink) are counted, especially if some of the rounded petal or banner petal shape is clearly visible.
- It is good to zoom in to confirm positive IDs, then zoom back out to confirm again to use extra context to make a decision. Look to see the full shape of the structure. Pods, in particular, may look like flowers when occluded or when viewed top down.

Can you find the “flowers” in a real image?

Purple Boxes: positive IDs
Red Boxes: DO NOT LABEL

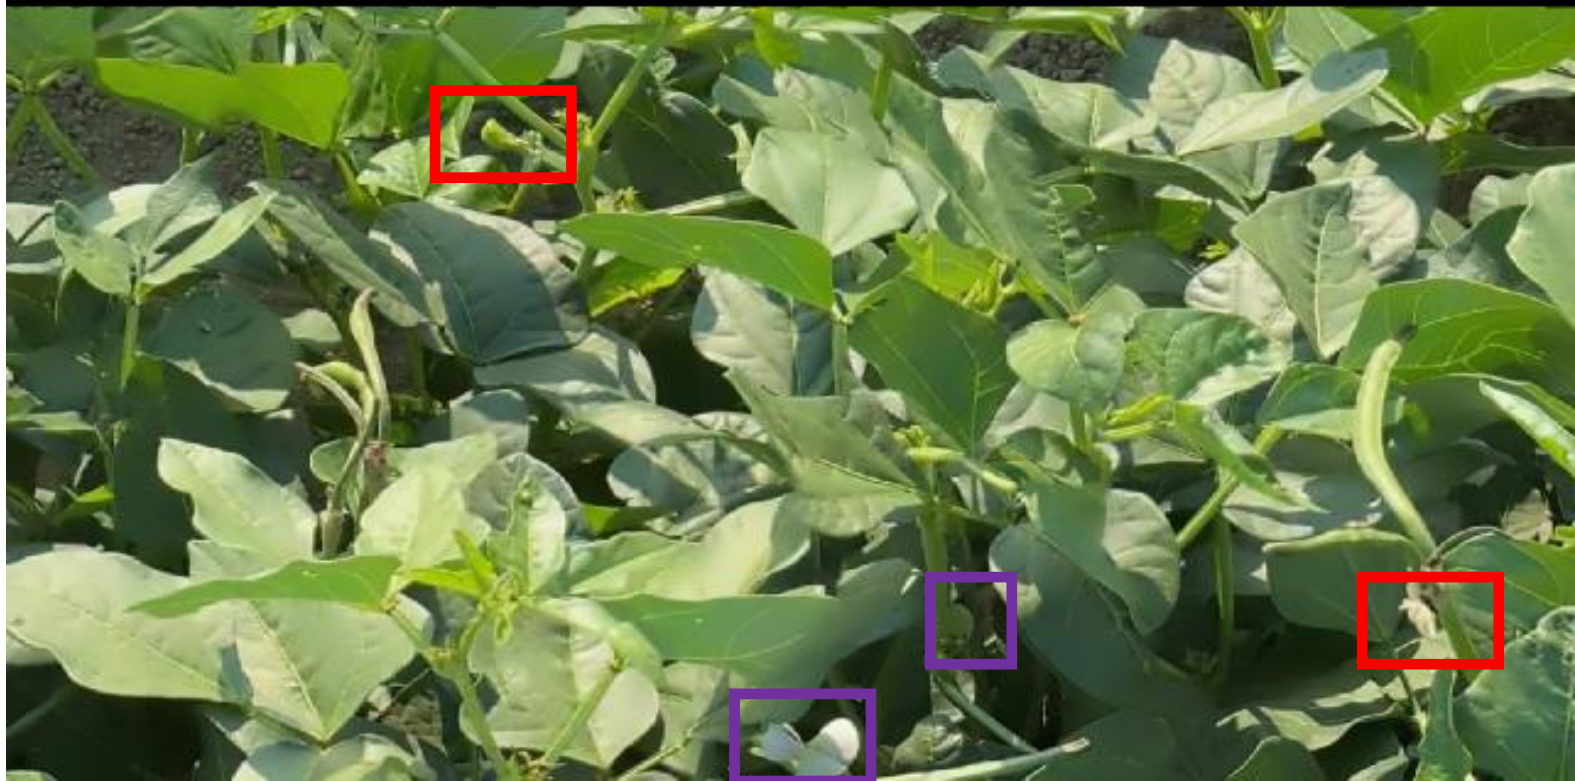

We have a flower bud that is too immature on the left and a dead and withered flower on the right that are both outside of our target.

We have a clear view of a fully open flower and a partial view of an open flower in the center.

On the left we have two senescent flowers, one still connected to the peduncle and one that is fallen on top of leaves, since they both have good color and shape, they are both counted.

On the bottom is a partially open flower and on the right is an occluded flower (white/purple banner petals visible), and a senescent flower that is right on the edge of counting (it is counted because the visible pod is sufficiently thing and short and the senescent flower is less than ¼ white but otherwise has the correct color).

The positive IDs include two senescent flowers on the ground that have good shape and color (we can see that these are actually “ghosts” due to the stitching process, but we label both). There is another withered senescent flower on the ground (> ¼ white) that is not counted. The other positive IDs are a senescent flower connected to the peduncle with good shape and color, an occluded open flower where we see enough shape and color for the positive ID, and a senescent flower that might be on a leaf that has good shape and color.

Among the negative IDs are an area of specular reflectance from a leaf in a vaguely bud shape, several dried flowers on the tips of pods, an immature flower bud, and the tip of a peduncle.

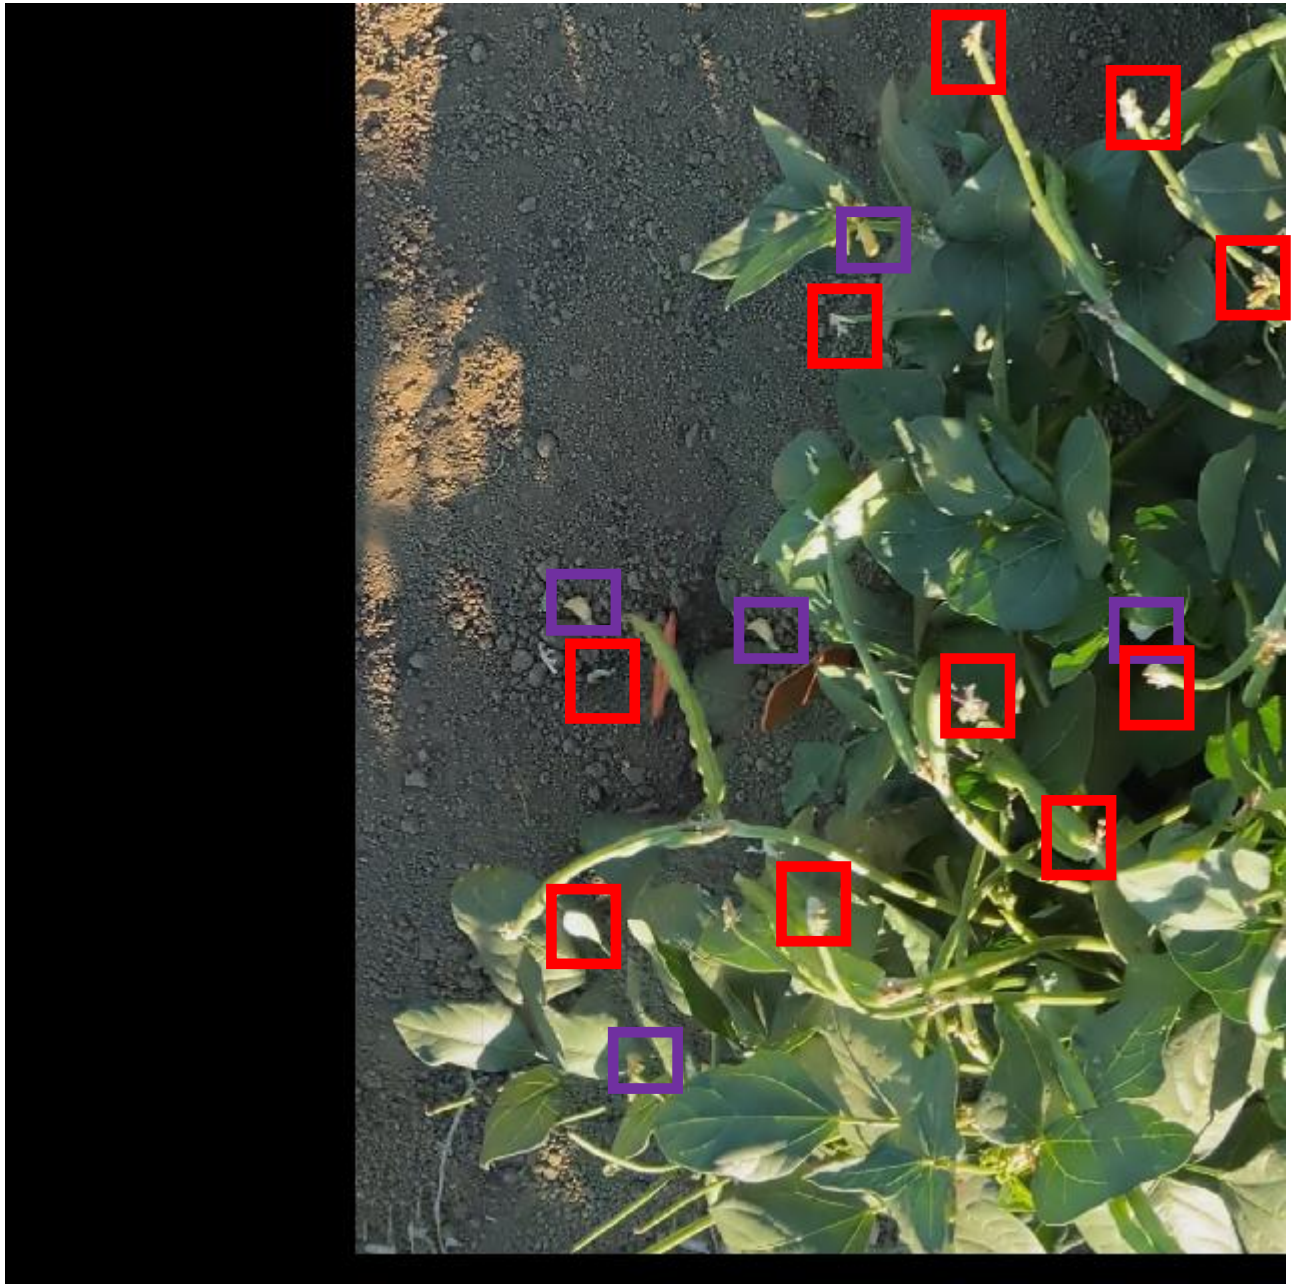

We have three positive IDs that show different angles, but with the correct shape and color for the viewing angle.

Two flowers here within the purple box, one seen from the back with mostly indirect light and another one with what looks to be a back view but with high reflectance.

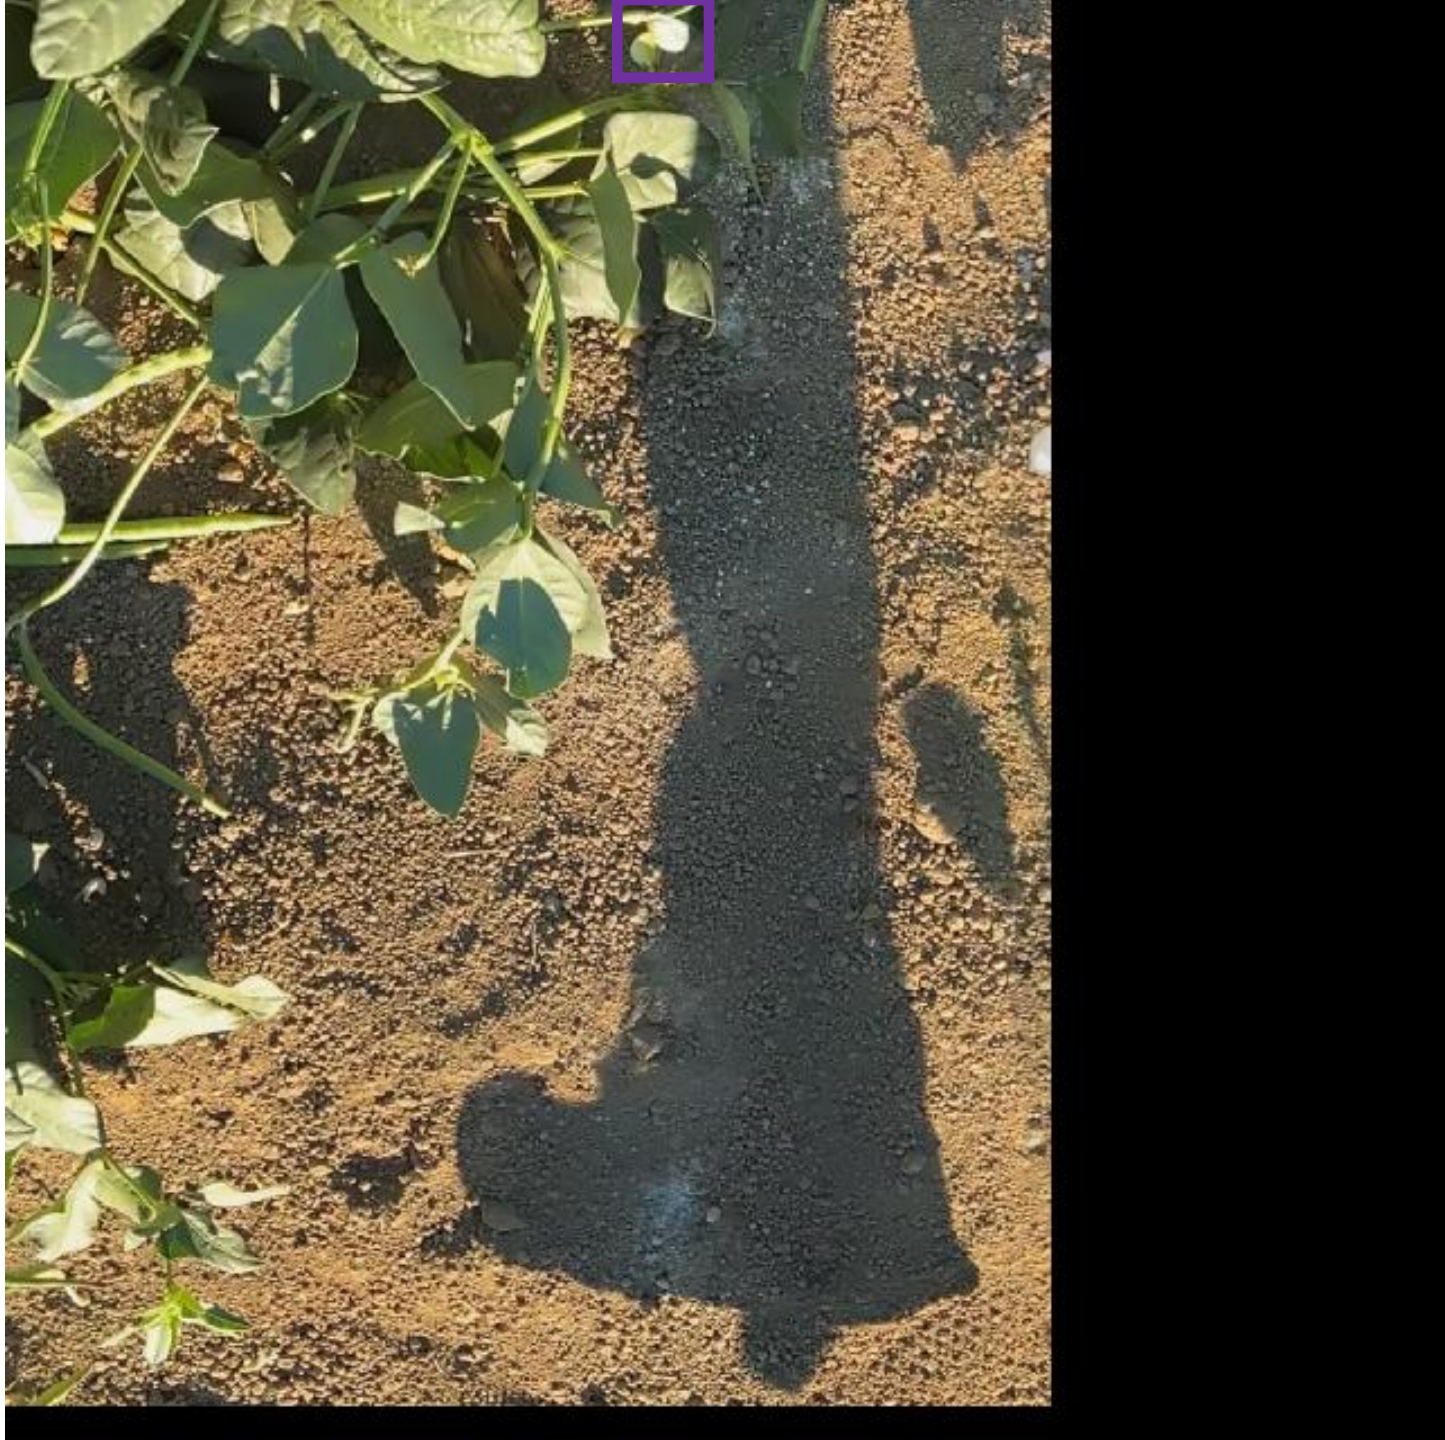

We have a clear back view of an open flower, a blurred flower with enough color for a positive ID, and a senescent flower with a thin visible pod that just passes the shape and color test.

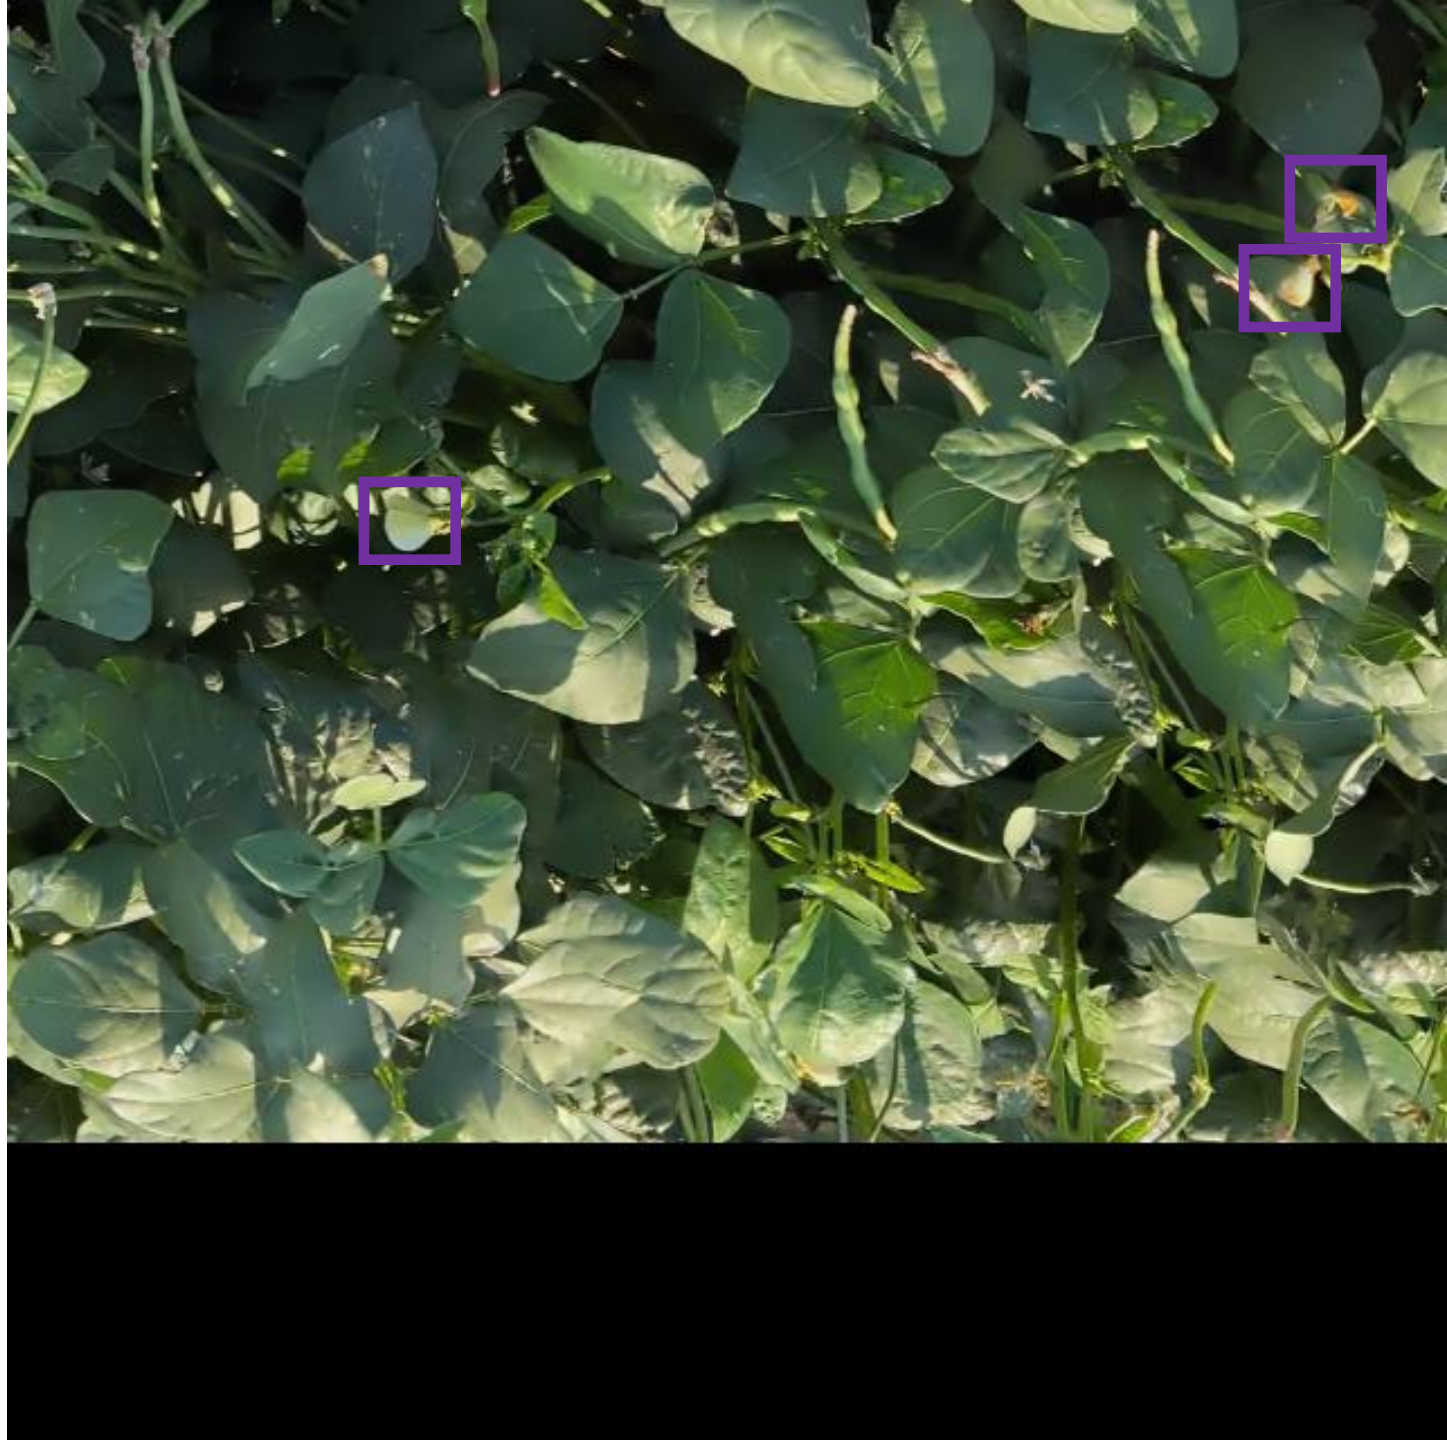

Among the negative IDs are dried pods, an immature bud, an overly withered senescent flower, and a hard to identify thing in the middle of the image.

The positive IDs include three open flowers, two senescent flowers still connected to the peduncle and one likely senescent flower that is partially occluded and on a leaf near the center of the image.

Here we have five open flowers and two senescent flowers (one on the ground with good shape and color and one on the tip of a pod that has good enough color and shape to pass since the pod is short and skinny).

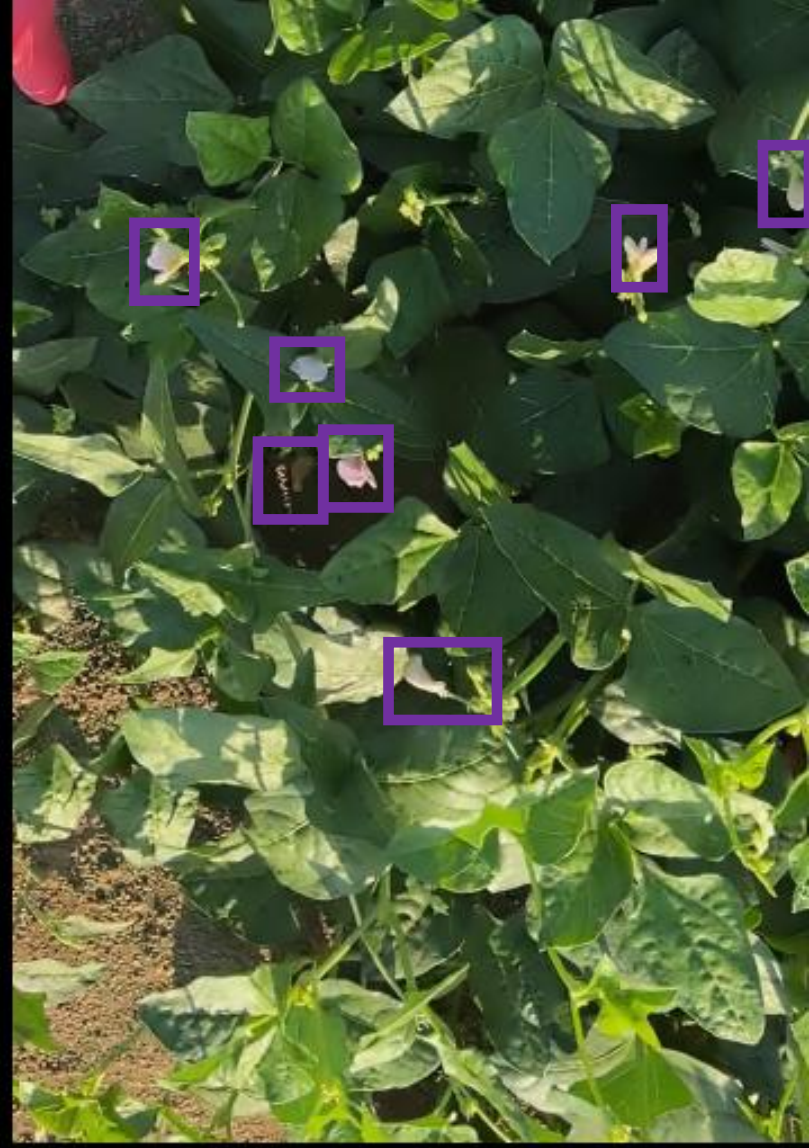

The senescent flower attached to a pod appears very white, but we still give it a positive ID since the shape is good and the pod is small, the white color can be attributed to reflection.

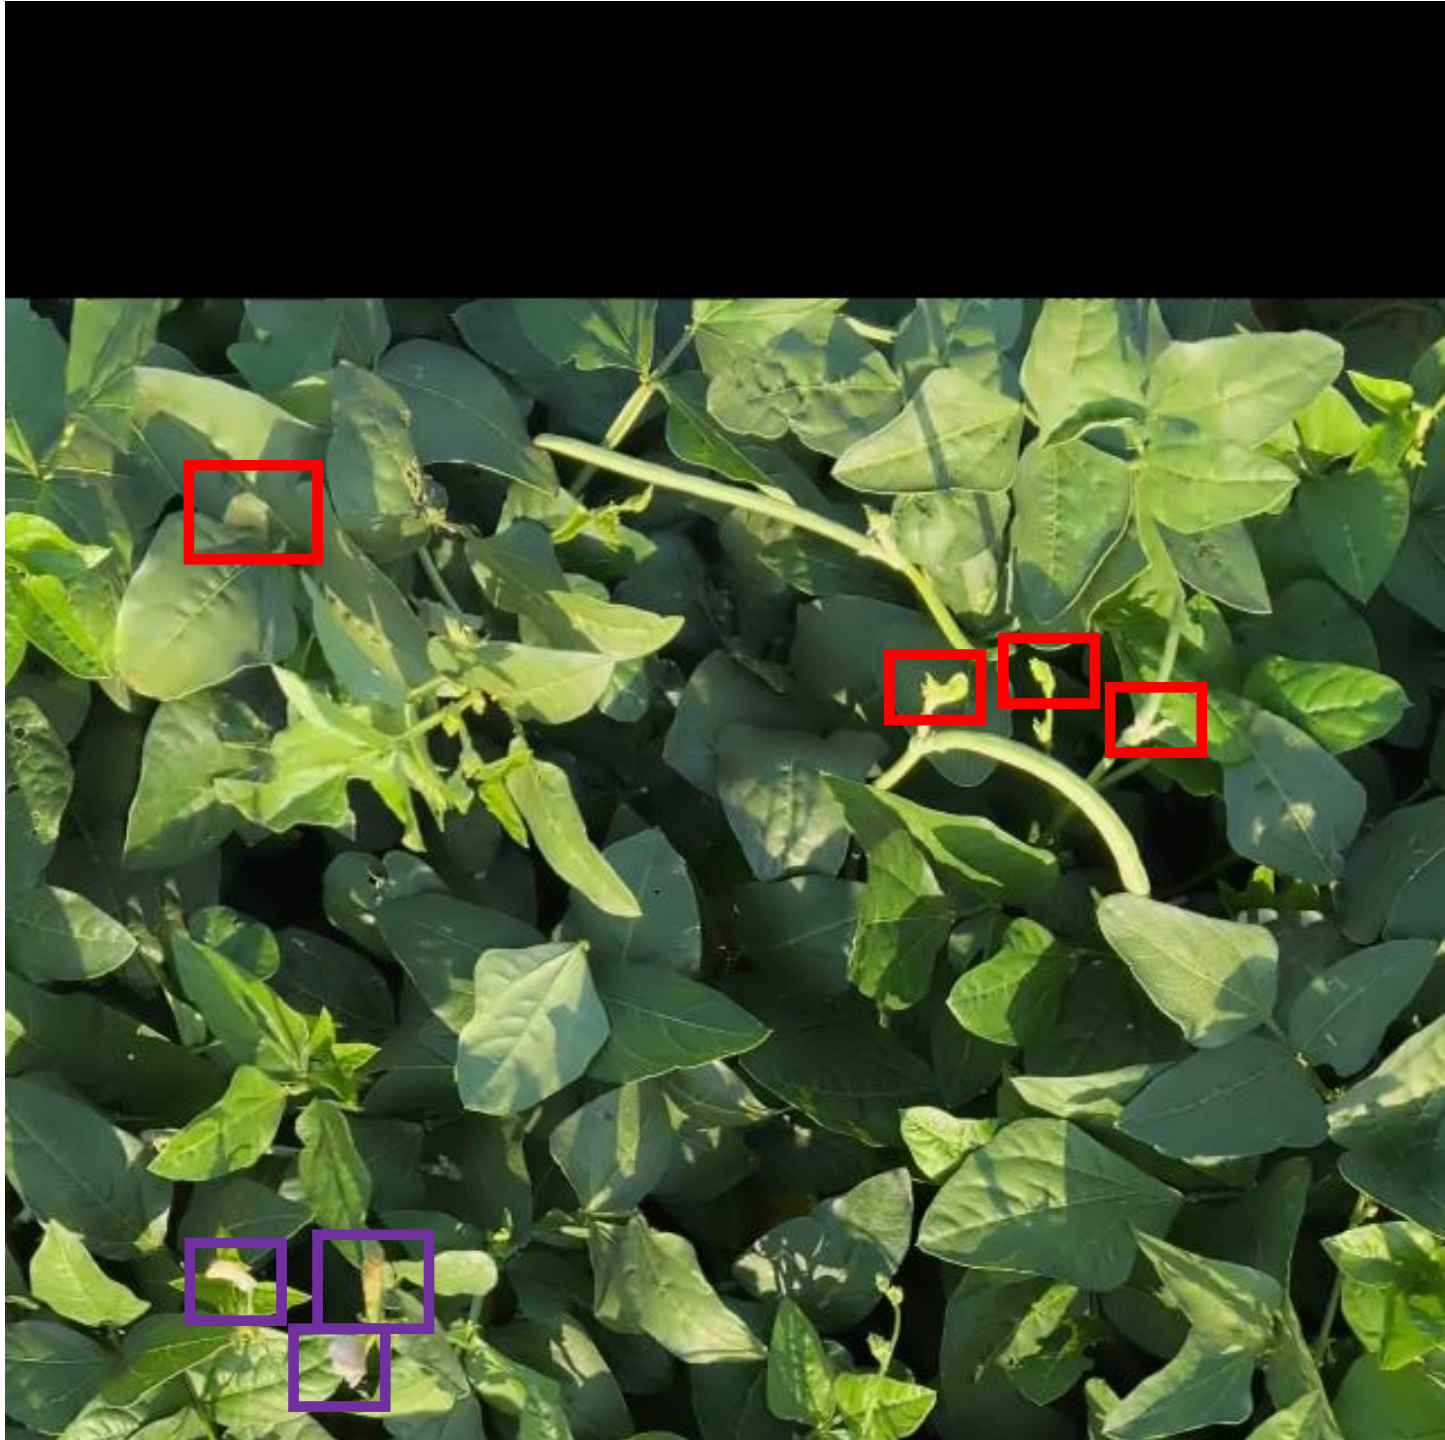

There are some yellow objects in the shadows, one where it clear it is a senescent leaf, the others we do not have enough context or shape information for a positive ID. There also some overly withered flowers that do not pass the shape and color test.

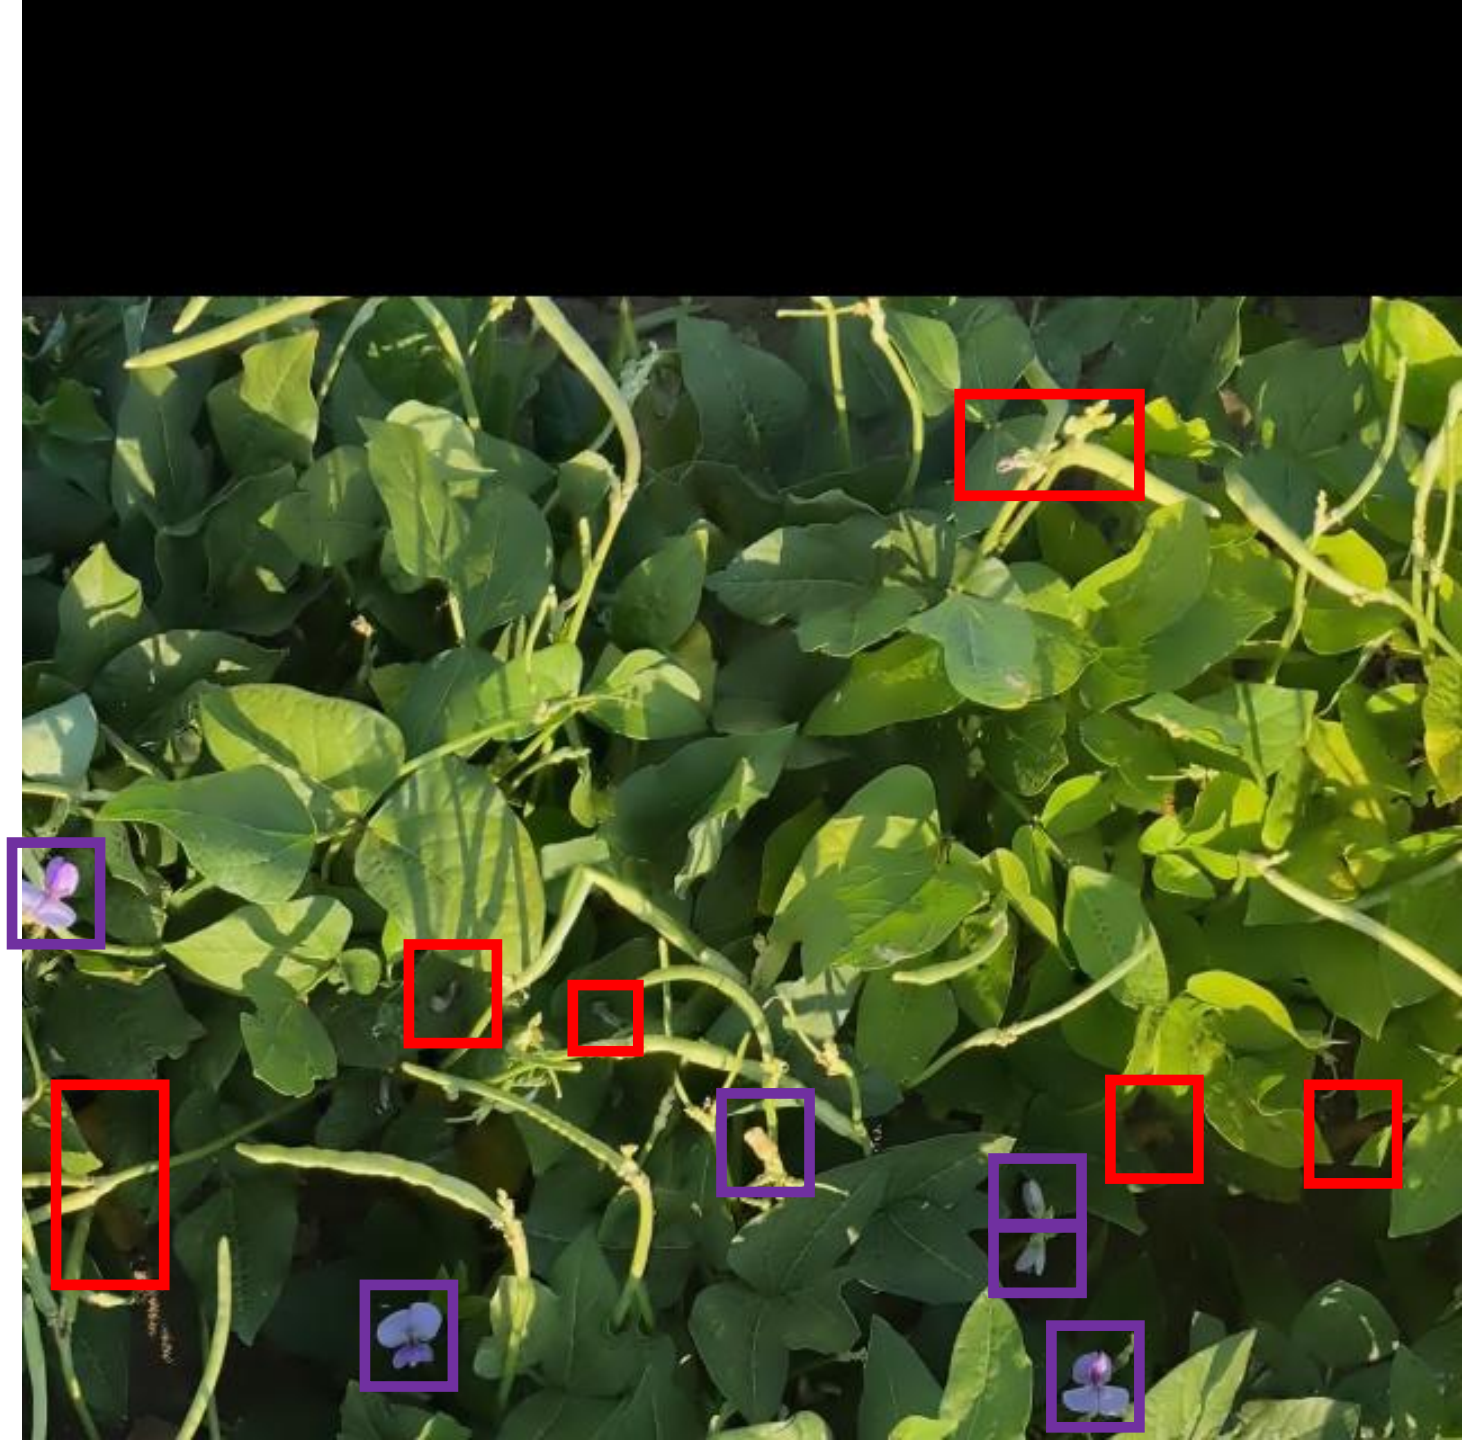

A nice view of shoots, rather than peduncles, with closed leaves at the tip.

There is an edge case that is particularly difficult. The banner petals of the open flowers in this image are a white-lime green, but the positive ID for these is easiest through shape alone. There are cases when we only see a very similar color (lime green), but it is unclear if it is a pod, an immature bud, or just the result of some reflectance. When we see an object of this color, it gets a clear positive ID when it has the open lobed shape (the banner petals are opening or opened the angle allows us to see this) and it gets an equivocal positive ID when we think the banner petals are about to open or opened at a small angle.

This is what looks to be a flower that is opening with yellow-green banner petals. The shape doesn’t look quite right for a pod, but it could just be the lighting and the pod is going away from us. I label it as a flower, but this is an example of one of the edge cases. Next to it is a similar structure but more green-white. This is potentially also a flower, but the color and shape are not good enough for a positive ID.

While still green, there is enough white/yellow to include three of these flowers that have yet to open, including one in a shadow. There are too senescent flowers on the ground with great shape and color and one that is shriveled, but just barely passes the shape test. The other senescent flowers are on peduncles and clearly pass both the color and shape test. One senescent flower is occluded but otherwise would pass the color and shape test – here I exclude it. There are also places with high reflectance that are likely just leaves and one place that looks more like the shape of a flower, but is likely a blurry leaf in the shadows.

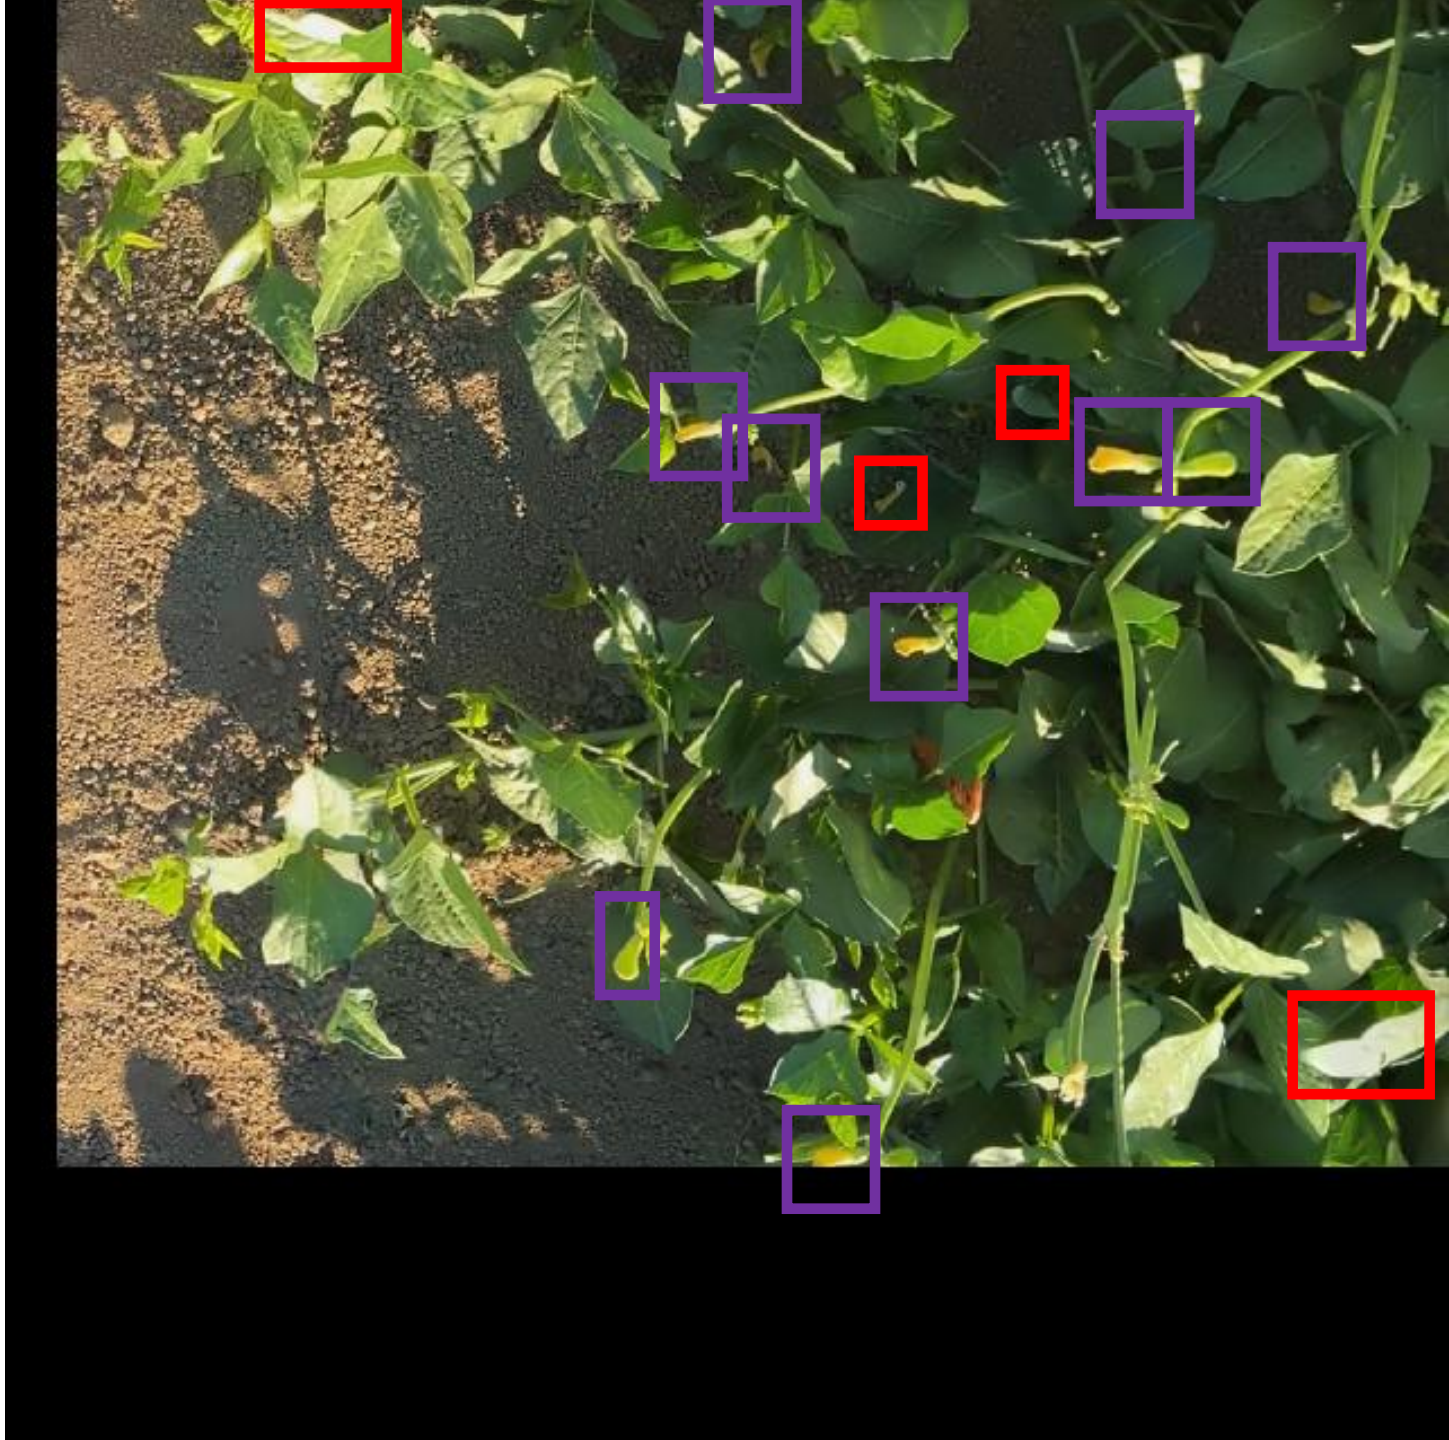

These two peduncles have three and four flowers, respectively. We see what looks to be two closed and one senescent flower on the left peduncle and one closed, one open, and two senescent flowers on the right peduncle. One senescent flower on the right peduncle is fully white, which should preclude it from a positive ID, but the pod is visible and sufficiently small that this passes with a positive ID.

The positive IDs here are mostly closed but have a nice white-purple color. A few look to be senescent as they are beige/yellow and two are hard to ID because of high reflection. We still assign positive IDs to the two flowers with high reflection based on context and shape; they are likely a closed and senescent flower.