

# BOLLWM: A REAL-WORLD DATASET FOR BOLLWORM PEST MONITORING FROM COTTON FIELDS IN INDIA

**Jerome White Chandan Agrawal Anmol Ojha**  
**Apoorv Agnihotri Makkunda Sharma Jigar Doshi**  
agri-ai@wadhwaniai.org  
Wadhvani Institute for Artificial Intelligence

## ABSTRACT

This paper presents a dataset of agricultural pest images captured over five years by thousands of small holder farmers and farming extension workers across India. The dataset has been used to support a mobile application that relies on artificial intelligence to assist farmers with pest management decisions. Creation came from a mix of organized data collection, and from mobile application usage that was less controlled. This makes the dataset unique within the pest detection community, exhibiting a number of characteristics that place it closer to other non-agricultural object detection datasets. This not only makes the dataset applicable to future pest management applications, it opens the door for a wide variety of other research agendas.

## 1 INTRODUCTION

Pest management plays a critical role for farmers around the world. This is especially true for smallholder farmers in low- and middle-income countries where failure to adhere to best practices can have devastating consequences (Guha & Das, 2022). Artificial intelligence can help, but requires the right kind of data to be effective—data that is not only sufficient in quantity for meaningful training, but also accurately representative of the deployment scenario.

This paper introduces an open source dataset that addresses that gap.<sup>1</sup> The dataset consists of images captured during the development and deployment of a mobile phone application (app) that assists farmers with pest management decisions (Dalmia et al., 2020). Collection has been ongoing since 2018, spanning several cotton farms and cotton seasons across India. The app provides recommendations for dealing with pink and American bollworms, two types of pests that have plagued the country’s cotton farmers for decades (Mohan et al., 2017; Fand et al., 2019; Vonzun et al., 2019; Najork et al., 2021). Most images in the set consist of one or more of these pests, and come with annotated bounding boxes delineating their location.

As compared to other pest-specific datasets, our contribution is notable for its scale and its authentic representation of field conditions (Sections 2 and 3). It also has characteristics that make it comparable to other object detection datasets not in the agriculture domain (Section 4). Finally, as a result of collection-by-deployment, there are a number of images that do not contain pests (Section 5). Their inclusion not only opens the door for interesting research directions, but facilitates model development that is robust and practical.

## 2 RELATED WORK

There has been a steady stream of deep learning work focused on the detection of pests in digital images (Prajapati et al., 2016; Borowiec et al., 2022; Suto, 2022; Júnior & Rieder, 2020; Barbedo, 2020; Li et al., 2021; Verma, 2022; Mekha & Parthasarathy, 2022). A side effect of this effort has been the presentation of several pest-related datasets. Nam & Hung (2018), for example, use a database of hundreds of images from a confectionery factory; Xie et al. (2018) build a set of

---

<sup>1</sup><https://github.com/WadhvaniAI/pest-management-opendata>

several thousand images from farms around China. The data presented in this work most closely resembles efforts to detect bollworms (Bakkay et al., 2018; Ding & Taylor, 2016) caught using pheromone traps (Sun et al., 2018; Ding & Taylor, 2016; Kalamatianos et al., 2018) installed on cotton fields (Sarangdhar & Pawar, 2017; Parikh et al., 2016; Revathi & Hemalatha, 2012). In many of these cases, the datasets are difficult to utilize because they are not available publicly (Nam & Hung, 2018; Ding & Taylor, 2016; Xie et al., 2018). That barrier aside, many were collected in controlled settings, either from internet searches (Liu et al., 2016; Ahmad et al., 2022; Wu et al., 2019), or using purpose-built trap imaging devices (Wang et al., 2020). In both cases, such sets are of limited value to our solution, as images from our users are often of more diverse quality than those found on the internet, or come with more noise than those taken with dedicated hardware.

Recent work has sought to fill this gap by introducing publicly available datasets designed for building pest detection models. IP102 (Wu et al., 2019) is a collection of several thousand images across more than 100 species. The scale of this dataset—both with respect to images and pest classes—is comparable to what we present. Unlike our images however, which were collected by farmers and farm extension workers in cotton fields, their set was collated from images found on the internet. AgriPest (Wang et al., 2021) is a collection consisting of thousands of images across various crop and pest types. As with our dataset, their images are also collected in actual field conditions; however they do not contain the same crop or pest types we do. Pest24 (Wang et al., 2020) is a data set of almost 30,000 images containing almost 40 different pest species—the species in our dataset are a subset of theirs. Like our work, their dataset was collected on actual fields. However, because they use highly specialized monitoring devices, images are much more homogeneous than ours. We provide a closer comparison to IP102 and AgriPest in Section 4.

### 3 DATASET CREATION

Our dataset consists of images containing bollworms caught in pheromone traps on cotton farms across India. A pheromone trap is a class of pest traps in which a pheromone-emitting lure attracts male adult bollworms (Figure 3a). Each lure is designed to attract a specific species; for our collection those were pink and American bollworms, which we sometimes refer to as PBW and ABW, respectively. This section describes the two efforts that were undertaken to build our dataset.

#### 3.1 DATA FROM DEDICATED COLLECTION

Data collection began with the explicit goal of building a pest detection model to support a mobile app for pest management. The intent was for farmers to take photos of bollworms caught in pheromone traps around a field. The app would then count relevant pests in the images and provide spray recommendations based on mappings between counts and agriculturally established best-practices. Having this solution clarity allowed us to design the data collection activity with protocols to build an appropriate set. To facilitate collection, partner organizations with expertise in cotton farming and having long-standing relationships with individual farmers were identified and engaged to assist. The organizations agreed that while making their regularly planned field visits, they would also collect images of trap catch.

Coordinated data collection began in early October 2018. A total of 26 farms were chosen in and around Wardha, a district in eastern Maharashtra (India). Each farm was outfitted with two pest traps specific to bollworms. Images were captured using ODK (Hartung et al., 2010). Forms were designed that allowed extension workers to first specify the lure type, then to take a photo of its contents. Workers were asked to empty the contents of the trap onto a plain white sheet of paper and take photos using reasonable photography techniques—to not engage the flash, and to ensure bollworms were in focus. They were further instructed to ensure that the photo was taken from a distance that minimized the space between the contents and the edge of the frame. Because photos were taken at a range of distances, the data set contains a corresponding range of pest scales.

Once Wardha’s kharif season was finished, coordinated collection restarted in Nannilam, a town in Tamil Nadu approximately 1,300 kilometers south. Unlike Wardha, Nannilam grows cotton during the summer season, a time that is inactive for farmers in Wardha. This allowed us to continue collection relatively uninterrupted. Forty four farms were enrolled during this phase. Workers were

Table 1: Descriptive statistics of comparative datasets. See Sections 4 and A.1.1 for further details.

Dataset	Images		Boxes				Classes	
	Total	Res. (MP)	Total	Per Image	Size	O/L	Total	Per Image
BOLLWM	19 536	9.85±0.04	575 939	29.48±0.49	0.21%	0.64	2	1.05±0.00
BOLLWM/PBW	16 176	9.62±0.04	536 653	33.18±0.58	0.15%	0.65		
BOLLWM/ABW	2 316	11.49±0.08	13 838	5.97±0.21	2.29%	0.50		
IP102	15 177	0.33±0.01	17 839	1.18±0.01	37.18%	0.05	97	1.00±0.00
AgriPest	28 169	3.33±0.02	144 798	5.14±0.07	0.28%	0.18	13	1.00±0.00
CARPK	989	0.92±0.00	42 274	42.74±0.47	0.46%	0.97	1	
SKU-110K	8 807	7.86±0.02	1 299 450	147.55±0.47	0.29%	0.90	1	
COCO/small	63 911	0.27±0.00	371 655	5.82±0.02	0.30%	0.45	80	2.05±0.01
COCO/med.	86 362	0.27±0.00	307 732	3.56±0.01	2.98%	0.43	80	1.85±0.01

requested to visit fields with a higher, more regular cadence than had been done during the previous season.

In both phases, quality assurance was maintained in two ways. First, regular feedback was provided through weekly conference calls and field visits. During conference calls, a sample subset of images were reviewed by the authors with the extension workers and appropriate feedback on picture quality given. Second, mobile chat groups consisting of extension workers, managers of the extension workers, and the authors of this paper were maintained. The group was used to quickly respond to general questions, provide feedback, and motivate extension workers to take high quality pictures. This was particularly important since data collection was not the primary focus of their field visits.

### 3.2 DATA FROM APPLICATION DEPLOYMENT

Using the data gathered from dedicated collection, object detection models were trained to recognize pink and American bollworms, facilitating development of our mobile app. During the 2019 kharif season, we deployed the app throughout Maharashtra, including to the same community of farmers and farming extension workers in Wardha with whom we interacted a year prior.

This phase marked a turning point in data collection: from that season onward images came exclusively from the app, which was deployed to a wider audience of users as the years went on. Prior to each deployment, users were educated on how to use the app, including instructions for how to empty traps and take photos such that the chances of model accuracy are highest. However, as usage scaled it became difficult to ensure the same level of protocol adherence as when collection was organized—the primary goal of the user base was pesticide decision making rather than positive feedback on image quality. This had noticeable implications on the range of data we received (Section 5), and on modeling techniques we pursued (Dalmia et al., 2020; White et al., 2022).

## 4 DATASET COMPARISONS

Characteristics of our dataset allow it to fit at the intersection of other popular object detection focused datasets, both within the pest detection space and otherwise. Our complete dataset (BOLLWM) consists of 47,851 images. The analysis in this section, presented in Table 1, is based on a subset of that total meant specifically for training. Similar training subsets—“train” and “validation”—were considered for third party datasets to keep comparisons fair. As such, numbers for datasets that are not ours may differ from their previously published results. In addition, our data was restricted to images that contained pests.

### 4.1 PEST DETECTION DATASETS

Rows IP102, AgriPest, and BOLLWM represent datasets devoted to pest detection. We have a similar number of images as IP102 (*images*→*total*) and a similar average pest size as AgriPest (*boxes*→*size*), likely because both we and AgriPest contain images taken in similar field conditions. IP102’s larger pest size is likely an artifact of it being a curated set of internet images,

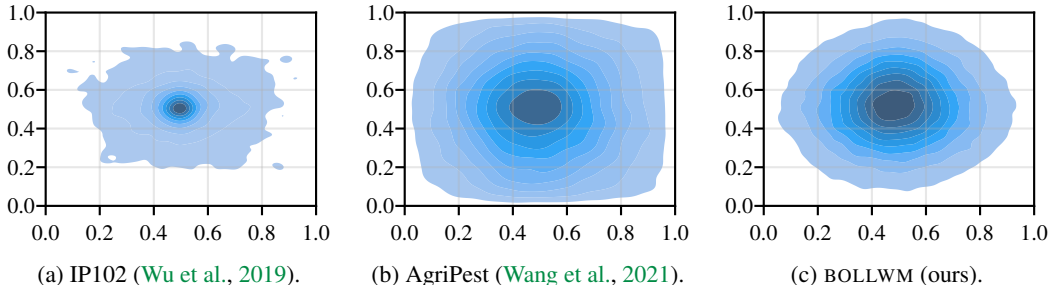


Figure 1: Kernel density estimations of pest centroids. See Sections 5 and A.1.2 for further details.

which often fill the frame the crop of the pest. Our image resolution ( $images \rightarrow res.$ ) and total number of boxes ( $boxes \rightarrow total$ ) are much larger, in both cases by several multiples. We are also the only set with images that contain multiple classes ( $classes \rightarrow per\ image$ ).

Our dataset stands out for its pest density. The number of boxes per image ( $boxes \rightarrow per\ image$ ) and the number of images with at least one overlapping box ( $boxes \rightarrow O/L$ ) are proxies for this. The almost 30 boxes per image exhibited by our set is almost five times higher than the others. The standard error around that average is also much higher, meaning that there are some very tightly packed pest clusters in some of our images. Within the pest detection community, it has been found that this characteristic makes object detection more difficult (Wang et al., 2020).

Pest positions across the datasets also differ. Figure 1 presents density plots of bounding box centroids for each dataset. For each image box centroids were extracted and converted to positions relative to that image shape. Each panel in Figure 1 is the aggregation of these points by dataset. Lighter hues in a panel represent positions with fewer centroids; darker regions represent the opposite. In each dataset, pests tend to fall more heavily into the center of images. This is most acute in IP102. AgriPest tends to spread pests more equally around the image than other datasets. This is likely because many of their pests are found on the leaves and stems of plants, which tend to fill the frame. Our dataset has the widest distribution about the center.

#### 4.2 GENERAL OBJECT DETECTION DATASETS

Some of the differentiating characteristics in our dataset are comparable to other datasets not dedicated to pest detection. Table 1 quantifies these differences through rows in the “general” domain. Those datasets include the Car Parking Lot Dataset (CARPK) (Hsieh et al., 2017), SKU-110K (Goldman et al., 2019), and Microsoft’s Common Objects in Context (COCO) (Lin et al., 2014).

CARPK is a collection of car images taken by drones flown over parking lots. It was designed to facilitate research in object counting, but has also been used to explore methods on densely packed scenes (dos Santos de Arruda et al., 2022; Goldman et al., 2019). SKU-110K was designed specifically for detection in densely packed scenes. It is a collection of images of retail shelf displays, which require scores of intersecting bounding boxes to properly annotate. From Table 1, both datasets contain far more boxes per image than ours on average ( $boxes \rightarrow per\ image$ ); and images are much more likely to contain an overlapping pair ( $boxes \rightarrow O/L$ ). However, with respect to these metrics, our dataset is closer to these specialized corpora than to those in the pest detection space. Our dataset also contains more than one class, adding additional responsibility to models that want to count well in our domain.

COCO is a well established, multi-purpose dataset that is commonly used for object detection research. Although it is a more general purpose dataset, its popularity makes it convenient for couching comparisons in a well known context. In this case, we look specifically at its “small” and “medium” scale subsets, defined by the COCO maintainers as objects with area less-than 32 square pixels, and 96 square pixels, respectively. For either subset, the percentage of images containing overlapping boxes is lower than ours, but closer relatively than to the other pest detection sets. Our subset of ABW images has characteristics of both scale subsets: the number of boxes per image is in line with COCO-small, while the relative size of each box is comparable to COCO-medium.

## 5 IMAGE TYPES

Section 3 described the two phases of data collection. In the first, enumerators were explicitly tasked with taking photos that met our protocol. In the second collection became a side effect of a wider pest management application deployment. As mentioned, this had implications on the type of data added to our corpus.

Across phases there were changes to image quality, and to the diversity of pest configurations. Of note were the number of images that were submitted that did not have pests, and those that had pests, but on backgrounds other than white sheets of paper. Figure 2 visualizes this progression across all data splits. During the initial kharif and summer seasons—late 2018 and early 2019, respectively—almost all images contained pests; noted as the “positive” set in the figure. Once collection came from deployment, late 2019 onward, the number of “negative” images—images not containing pests—began to rise. Of the almost 48,000 images in our dataset, approximately 41 percent (19,558) are negatives. Section A.2 provides examples of all image types.

Negative images are a natural occurrence for two reasons. The first is that there are genuine cases in which no pests are caught. There are also cases when users want to maintain their cadence of app usage, and knowingly submit no-pest images to do so. This set of negatives are generally white sheets of paper, as we request, that are either empty or filled with naturally occurring no-pest farm artifacts such as pieces of soil. The other reason for this negative set are users experimenting with the app. As with any application, it is natural for users to experiment before using it in its proper context to get a feel for its usage or to demonstrate the flow to show others. In such cases, images will not contain any evidence of pheromone trap contents, and instead be of common household or farm-related items.

Others looking to build their own pest management solutions based on uncontrolled user interaction will have to be resilient to this situation. Models that assume a strict domain are unlikely to perform well under such conditions. When our application is unable to correctly distinguish negative images from positive, for example, users lose trust and are hesitant to rely on the solution for such a critical component of their livelihoods. Inclusion of this negative set allows others to build models that are robust to a variety of input types.

## 6 CONCLUSION

Pest management is one of the most important responsibilities of a farmer. Good pest management involves consistent monitoring of pest populations, with appropriate action taken when those populations exceed a threshold (Kogan, 1998; Mancini et al., 2009; Rajashekhar et al., 2022). Although there are ongoing attempts to educate farmers on proper pest detection and intervention, successful programs require sizable on-ground support to scale and continued follow-up to be effective (Mancini et al., 2007). In the absence of such established, formalized, and accepted practices, pest detection can be ad hoc and pesticide usage excessive (Devi, 2010; Ranganathan, 2018); something that is not only expensive, but also dangerous to the environment and to human health (Donthi, 2021).

There is scope for artificial intelligence to mitigate this cycle. To do so, however, requires the right kind of data. This paper has presented a dataset that, due to the circumstances of its acquisition, represents a corpus that is closer to the ground reality than any other dataset in the open domain. The dataset is publicly available for researchers to develop computer vision techniques in a real-world setting, or for practitioners to explore new solutions to pest management that rely on artificial intelligence.

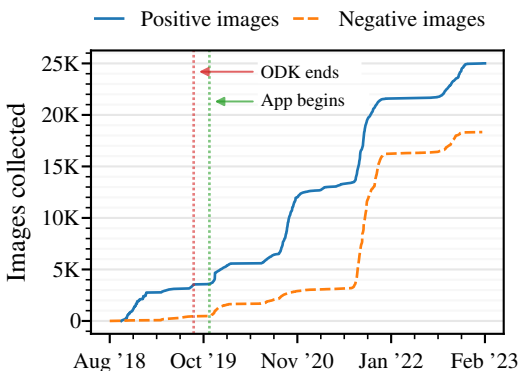


Figure 2: Progression of image collection.

#### ACKNOWLEDGMENTS

The authors thank Balasaheb Dhame, Sonali Ghike, Rajesh Jain, Raghavendra Udupa, Srinivasa Rao, Kalyani Shastry, Sonu Vijan, and Dhruvin Vora for assistance at various stages of collection and curation. On-ground support was facilitated by Ambhuj Cement Foundation, Deshpande Foundation, Lupin, STL Garv, Syngenta, Tata Trust, and Welspun Foundation.

Portions of this work were made possible with support by the convening sponsor, “FAIR Forward—Artificial Intelligence for All,” which is implemented by Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH on behalf of the German Ministry for Economic Cooperation and Development (BMZ).

#### REFERENCES

- Ifthikhar Ahmad, Yayun Yang, Yi Yue, Chen Ye, Muhammad Hassan, Xi Cheng, Yunzhi Wu, and Youhua Zhang. Deep learning based detector yolov5 for identifying insect pests. *Applied Sciences*, 12(19), 2022. doi: 10.3390/app121910167.
- Mohammed Chafik Bakkay, Sylvie Chambon, Hatem A. Rashwan, Christian Lubat, and Sébastien Barsotti. Support vector machine (SVM) recognition approach adapted to individual and touching moths counting in trap images. *CoRR*, abs/1809.06663, 2018. doi: 10.48550/arXiv.1809.06663.
- Jayne Garcia Arnal Barbedo. Detecting and classifying pests in crops using proximal images and machine learning: A review. *AI*, 1(2):312–328, 2020. doi: 10.3390/ai1020021.
- Marek L. Borowiec, Rebecca B. Dikow, Paul B. Frandsen, Alexander McKeeken, Gabriele Valentini, and Alexander E. White. Deep learning as a tool for ecology and evolution. *Methods in Ecology and Evolution*, 13(8):1640–1660, 2022. doi: 10.1111/2041-210X.13901.
- Aman Dalmia, Jerome White, Ankit Chaurasia, Vishal Agarwal, Rajesh Jain, Dhruvin Vora, Balasaheb Dhame, Raghu Dharmaraju, and Rahul Panicker. Pest management in cotton farms: An ai-system case study from the global south. In *International Conference on Knowledge Discovery and Data Mining*, pp. 3119–3127. ACM, 2020. doi: 10.1145/3394486.3403363.
- Indira Devi. Pesticides in agriculture—a boon or a curse? A case study of Kerala. *Economic and Political Weekly*, 45(26–27), June 2010.
- Weiguang Ding and Graham Taylor. Automatic moth detection from trap images for pest management. *Computers and Electronics in Agriculture*, 123:17–28, 2016. ISSN 0168-1699. doi: 10.1016/j.compag.2016.02.003.
- Narasimha Reddy Donthi. Pesticide poisoning in India. *Economic and Political Weekly*, 56(45–46), November 2021.
- Mauro dos Santos de Arruda, Lucas Prado Osco, Plabiany Rodrigo Acosta, Diogo Nunes Gonçalves, José Marcato Junior, Ana Paula Marques Ramos, Edson Takashi Matsubara, Zhipeng Luo, Jonathan Li, Jonathan de Andrade Silva, and Wesley Nunes Gonçalves. Counting and locating high-density objects using convolutional neural network. *Expert Systems with Applications*, 195:116555, 2022. doi: 10.1016/j.eswa.2022.116555.
- Babasaheb B. Fand, V. S. Nagrare, S. P. Gawande, D. T. Nagrale, B. V. Naikwadi, Vrushali Deshmukh, Nandini Gokte-Narkhedkar, and V. N. Waghmare. Widespread infestation of pink bollworm, *pectinophora gossypiella* (saunders) (lepidoptera: Gelechiidae) on Bt cotton in central India: a new threat and concerns for cotton production. *Phytoparasitica*, 47(3):313–325, 2019. doi: 10.1007/s12600-019-00738-x.
- Eran Goldman, Roei Herzig, Aviv Eisenschtat, Jacob Goldberger, and Tal Hassner. Precise detection in densely packed scenes. In *Conference on Computer Vision and Pattern Recognition*. IEEE/CVF, June 2019.
- Pradyut Guha and Tiken Das. Farmers’ suicides in India. *Economic and Political Weekly*, 57(5):13, 2022.

- Carl Hartung, Adam Lerer, Yaw Anokwa, Clint Tseng, Waylon Brunette, and Gaetano Borriello. Open data kit: Tools to build information services for developing regions. In *International Conference on Information and Communication Technologies and Development*, number 18, pp. 18:1–18:12. ACM, 2010. doi: 10.1145/2369220.2369236.
- Meng-Ru Hsieh, Yen-Liang Lin, and Winston H. Hsu. Drone-based object counting by spatially regularized regional proposal network. In *International Conference on Computer Vision (ICCV)*, pp. 4165–4173. IEEE, October 2017. doi: 10.1109/ICCV.2017.446.
- Telmo De Cesaro Júnior and Rafael Rieder. Automatic identification of insects from digital images: A survey. *Computers and Electronics in Agriculture*, 178:105784, 2020. doi: 10.1016/j.compag.2020.105784.
- Romanos Kalamatianos, Ioannis Karydis, Dimitris Doukakis, and Markos Avlonitis. Dirt: The dacus image recognition toolkit. *Journal of Imaging*, 4(11), 2018. ISSN 2313-433X. doi: 10.3390/jimaging4110129.
- Marcos Kogan. Integrated pest management: Historical perspectives and contemporary developments. *Annual Review of Entomology*, 43(1):243–270, 1998. doi: 10.1146/annurev.ento.43.1.243.
- Wenyong Li, Tengfei Zheng, Zhankui Yang, Ming Li, Chuanheng Sun, and Xinting Yang. Classification and detection of insects from field images using deep learning for smart pest management: A systematic review. *Ecological Informatics*, 66:101460, 2021. doi: 10.1016/j.ecoinf.2021.101460.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), *European Conference on Computer Vision*, pp. 740–755. Springer International Publishing, 2014.
- Ziyi Liu, Junfeng Gao, Guoguo Yang, Huan Zhang, and Yong He. Localization and classification of paddy field pests using a saliency map and deep convolutional neural network. *Scientific Reports*, 6(20410), 2016. ISSN 2045-2322. doi: 10.1038/srep20410.
- Francesca Mancini, Ariena van Bruggen, and Janice Jiggins. Evaluating cotton integrated pest management (IPM) farmer field school outcomes using the sustainable livelihoods approach in India. *Experimental Agriculture*, 43(1):97–112, 2007. doi: 10.1017/S001447970600425X.
- Francesca Mancini, Janice Jiggins, and Michael O’Malley. Reducing the incidence of acute pesticide poisoning by educating farmers on integrated pest management in South India. *International Journal of Occupational and Environmental Health*, 15(2):143–151, 2009. doi: 10.1179/oeht.2009.15.2.143.
- Jose Mekha and V. Parthasarathy. An automated pest identification and classification in crops using artificial intelligence—a state-of-art-review. *Automatic Control and Computer Sciences*, 56(3): 283–290, 2022. doi: 10.3103/S0146411622030038.
- Komarlingam S Mohan et al. An area-wide approach to pink bollworm management on Bt cotton in India—a dire necessity with community participation. *Curr. Sci*, 112(10):2016–2020, 2017.
- Katharina Najork, Susheel Gadela, Padmarao Nadiminti, Sreeramulu Gosikonda, Raghava Reddy, Ejnavarzala Haribabu, and Markus Keck. The return of pink bollworm in India’s Bt cotton fields: Livelihood vulnerabilities of farming households in Karimnagar district. *Progress in Development Studies*, 21(1):68–85, 2021. doi: 10.1177/14649934211003457.
- Nguyen Tuan Nam and Phan Duy Hung. Pest detection on traps using deep convolutional neural networks. In *Proceedings of the 2018 International Conference on Control and Computer Vision*, pp. 33–38. ACM, 2018. doi: 10.1145/3232651.3232661.
- Aditya Parikh, Mehul Raval, Chandrasinh Parmar, and Sanjay Chaudhary. Disease detection and severity estimation in cotton plant from unconstrained images. In *International Conference on Data Science and Advanced Analytics*, pp. 594–601. IEEE, October 2016. doi: 10.1109/DSAA.2016.81.



- Bhumika Prajapati, Vipul Dabhi, and Harshadkumar Prajapati. A survey on detection and classification of cotton leaf diseases. In *International Conference on Electrical, Electronics, and Optimization Techniques*, pp. 2499–2506, March 2016. doi: 10.1109/ICEEOT.2016.7755143.
- Mandla Rajashekhar, Thalla Prabhakar Reddy, Keerthi Manikyanahalli Chandrashekara, Banda Rajashekar, Mogili Jagan Mohan Reddy, Kommagoni Ramakrishna, Eetela Satyanarayana, Adhi Shankar, Afifa Jahan, Parimal Kumar, and Lakkakula Shrivika. Evaluation of integrated pest management module for pink bollworm, *pectinophora gossypiella* (saunders) and its economic analysis under farmer’s field conditions. *International Journal of Pest Management*, 0(0):1–9, 2022. doi: 10.1080/09670874.2022.2096269.
- Thiagu Ranganathan. Pesticide usage by cotton farmers in India. *Economic and Political Weekly*, 53(19), December 2018.
- P. Revathi and M. Hemalatha. Advance computing enrichment evaluation of cotton leaf spot disease detection using image edge detection. In *International Conference on Computing, Communication and Networking Technologies*, pp. 1–5, July 2012. doi: 10.1109/ICCCNT.2012.6395903.
- Adhao Asmita Sarangdhar and V. R. Pawar. Machine learning regression technique for cotton leaf disease detection and controlling using IoT. In *International conference of Electronics, Communication and Aerospace Technology*, volume 2, pp. 449–454, April 2017. doi: 10.1109/ICECA.2017.8212855.
- Yu Sun, Xuanxin Liu, Mingshuai Yuan, Lili Ren, Jianxin Wang, and Zhibo Chen. Automatic in-trap pest detection using deep learning for pheromone-based *dendroctonus valens* monitoring. *Biosystems Engineering*, 176:140–150, 2018. ISSN 1537-5110. doi: 10.1016/j.biosystemseng.2018.10.012.
- Jozsef Suto. Codling moth monitoring with camera-equipped automated traps: A review. *Agriculture*, 12(10), 2022. doi: 10.3390/agriculture12101721.
- Amit Verma. SVM, CNN and VGG16 classifiers of artificial intelligence used for the detection of diseases of rice crop: A review. In Subarna Shakya, Valentina Emilia Balas, Sinchai Kamolphiwong, and Ke-Lin Du (eds.), *Sentimental Analysis and Deep Learning*, pp. 917–931. Springer Singapore, 2022.
- Seraina Vonzun, Monika M. Messmer, Thomas Boller, Yogendra Shrivasa, Shreekant S. Patil, and Amritbir Riar. Extent of bollworm and sucking pest damage on modern and traditional cotton species and potential for breeding in organic cotton. *Sustainability*, 11(22), 2019. doi: 10.3390/su11226353.
- Qi-Jin Wang, Sheng-Yu Zhang, Shi-Feng Dong, Guang-Cai Zhang, Jin Yang, Rui Li, and Hong-Qiang Wang. Pest24: A large-scale very small object data set of agricultural pests for multi-target detection. *Computers and Electronics in Agriculture*, 175:105585, 2020. doi: 10.1016/j.compag.2020.105585.
- Rujing Wang, Liu Liu, Chengjun Xie, Po Yang, Rui Li, and Man Zhou. Agripest: A large-scale domain-specific benchmark dataset for practical agricultural pest detection in the wild. *Sensors*, 21(5), 2021. ISSN 1424-8220. doi: 10.3390/s21051601.
- Jerome White, Pulkit Madaan, Nikhil Shenoy, Apoorv Agnihotri, Makkunda Sharma, and Jigar Doshi. A case for rejection in low resource ML deployment. In *Conference on Neural Information Processing Systems, NeurIPS. Workshop on Challenges In Deploying And Monitoring Machine Learning Systems*, 2022. doi: 10.48550/arXiv.2208.06359.
- Xiaoping Wu, Chi Zhan, Yukun Lai, Ming-Ming Cheng, and Jufeng Yang. IP102: A large-scale benchmark dataset for insect pest recognition. In *IEEE CVPR*, pp. 8787–8796, 2019.
- Chengjun Xie, Rujing Wang, Jie Zhang, Peng Chen, Wei Dong, Rui Li, Tianjiao Chen, and Hongbo Chen. Multi-level learning features for automatic classification of field crop pests. *Computers and Electronics in Agriculture*, 152:233–241, 2018. ISSN 0168-1699. doi: 10.1016/j.compag.2018.07.014.



## A APPENDIX

### A.1 EXTENDED CAPTIONS

#### A.1.1 DESCRIPTIVE STATISTICS

**Table 1** provides descriptive statistics of our dataset alongside other object detection datasets:

**Image columns** Total images in the dataset along with average image resolution (image height multiplied by image width).

**Box columns** Total number of bounding boxes and the average number of bounding boxes per image. Size is the percentage of area each box consumes on average, relative to its parent image resolution. Overlap (*O/L*) represents the fraction of images in which at least one pair of boxes is overlapping.

**Class columns** Total number of bounding box classes, along with the average number of classes per image.

Slashes (“/”) in dataset names denote data subsets. For example,  $x/y$  denotes dataset  $x$  restricted to subset  $y$ .

Uncertainty—values after each plus-minus—are standard errors about the mean. Uncertainty was omitted from size due to all values being much lower-than zero. All statistics are based on data from respective train and validation sets.

#### A.1.2 RELATIVE BOX POSITIONS

**Figure 1** presents kernel density estimation plots of bounding box centroids. Centroids coordinates are normalized by the dimensions of their respective parent image. The  $x$ - and  $y$ -axis in each panel thus represents relative image width and height, respectively. Density estimates for each panel are independent; no normalization within or between datasets were taken. For this reason, color hues should not be compared across figures.

### A.2 IMAGE EXAMPLES

Figures 3, 4, and 5, and show examples of various image types in our dataset. Images have been scaled for brevity. **Figure 3a** is an example of a pheromone trap. Pests in figures 4, and 5 have been highlighted with their respective bounding boxes. Such boxes are not present in the actual raw data; they been added here to aid reader understanding.

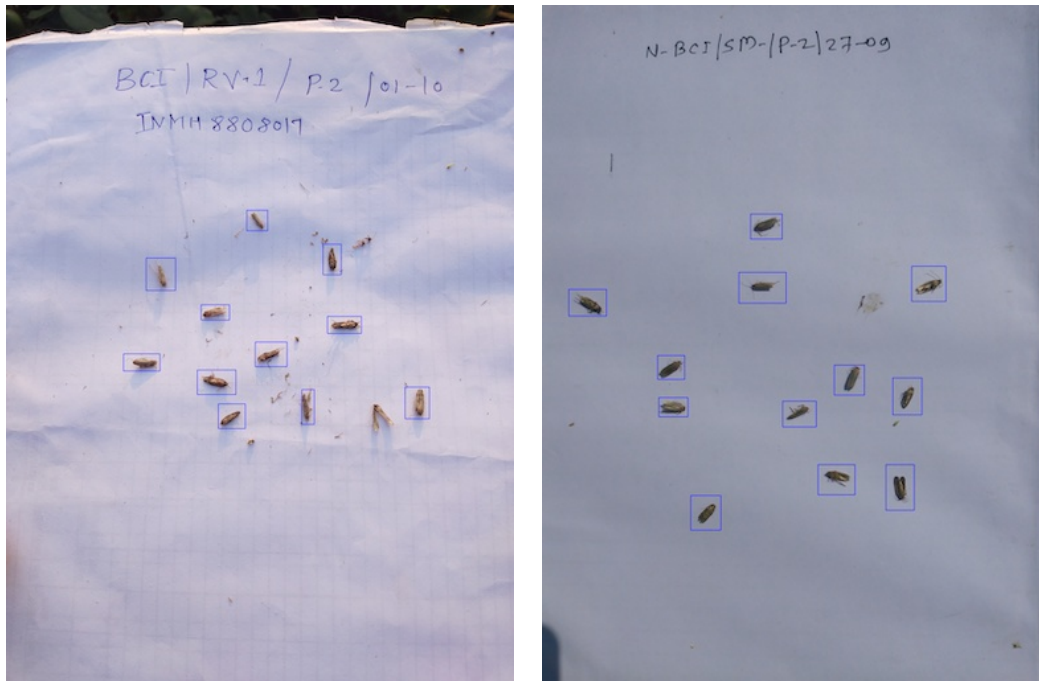


(a)



(b)

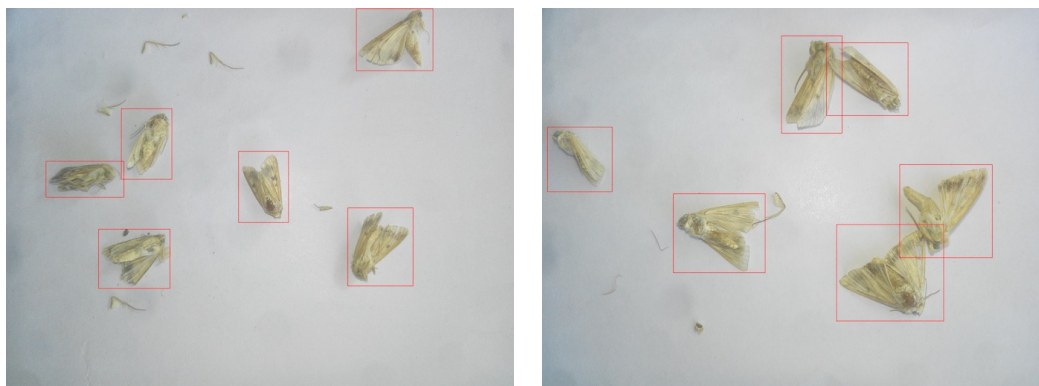
Figure 3: Examples of negative images in the dataset.



(a)

(b)

Figure 4: Examples of PBW images in the dataset.



(a)

(b)

Figure 5: Examples of ABW images in the dataset.