LANTERN-RD: Enabling Deep Learning for Mitigation of the Invasive Spotted Lanternfly

Srivatsa Kundurthy
The Academy for Mathematics, Science and Engineering
kundurthys@acm.org

Abstract

The Spotted Lanternfly (SLF) is an invasive planthopper that threatens the local biodiversity and agricultural economy of regions such as the Northeastern United States and Japan. As researchers scramble to study the insect, there is a great potential for computer vision tasks such as detection, pose estimation, and accurate identification to have important downstream implications in containing the SLF. However, there is currently no publicly available dataset for training such AI models. To enable computer vision applications and motivate advancements to challenge the invasive SLF problem, we propose LANTERN-RD, the first curated image dataset[1] of the spotted lanternfly and its look-alikes, featuring images with varied lighting conditions, diverse backgrounds, and subjects in assorted poses. A VGG16-based baseline CNN validates the potential of this dataset for stimulating fresh computer vision applications to accelerate invasive SLF research. Additionally, we implement the trained model in a simple mobile classification application in order to directly empower responsible public mitigation efforts. The overarching mission of this work is to introduce a novel SLF image dataset and release a classification framework that enables computer vision applications, boosting studies surrounding the invasive SLF and assisting in minimizing its agricultural and economic damage.

1 Introduction

Lycorma Delicatula, dubbed the Spotted Lanternfly (SLF), has recently spread invasively from parts of Eastern Asia to several other continents, establishing successfully in certain new environments due to reasons such as climate preference [9, 10, 25] and the abundance of sustainable plant hosts [1, 13, 12, 26]. Current regions of interest include the American Northeast and Japan, among others [27]. As a new contender in the ecosystem and a nuisance pest, the SLF poses a significant threat to local agriculture and the economy through its destruction of fruit trees, ornamental trees, timber, vineyards, and building structures [16, 18, 7], potentially leading to billions of dollars of damage. Evolving studies propose several methods for containment, from biocontrol agents [2, 4, 15] to trapping mechanisms [8, 19, 6]. Notably, local environmental authorities have mandated quarantines to slow the spread of the SLF [22], and are calling on the public to trap, report, and even take steps to exterminate sighted SLFs [21].

The expanding scope of the invasive SLF problem provides ground for computer vision applications to enhance research studies and containment efforts. For example, SLF pose estimation may boost understanding of key population-growth activities such as egg-laying and enable motion tracking to characterize population spread [17, 28]. Additionally, while ecologists studying the spatial distribution of the SLF have long relied on manual surveying methods in order to collect data [11, 5, 1, 18], advancements in AI animal detection techniques [20, 29] may improve the efficiency of such data

[1] Open-Access Dataset Link: https://github.com/srivatsa-kundurthy/LANTERN-RD-RAW

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

(a) Figured Tiger Moth (*Apantesis figurata*)

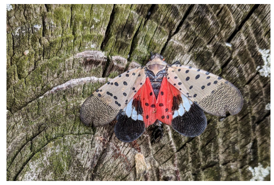

(b) Spotted Lanternfly (*Lycorma delicatula*)

(c) Pink Underwing (*Phyllodes imperialis*)

(d) Bella Moth (*Utetheisa ornatrix*)

(e) Spotted Lanternfly, wings closed

(f) White-lined Sphinx Moth (*Hyles lineata*)

Figure 1: The SLF and its look-alikes. Similar body-shape, color, and the presence of a bright pink, orange, or red underwing contribute to the visual similarity of these insects.

collection efforts. Finally, the presence of insect species that are visually similar to the SLF (Figure 1) raises additional complications for both researchers and the public, as misidentification of such insects affects the quality of collected data and results in unnecessary harm to native wildlife. For this issue, there is significant potential for AI-based classifiers to assist ecologists and promote public management efforts [3]. To enable such advancements in computer vision research for the invasive spotted lanternfly problem, there is a great need for a high-quality public dataset of SLF images.

To solve this problem, we present in this work a framework that consists of: (i) **LANTERN-RD**, a novel image dataset of the invasive SLF and visually-similar insects in order to enable computer vision research for SLF mitigation, (ii) a corresponding **baseline convolutional neural network** (CNN), and (iii) a **mobile deployment of the classifier** to empower rapid identification of the SLF against look-alikes and advocate for responsible mitigation efforts.

2 Dataset

For the purpose of empowering the development of deep learning models to assist in efforts to contain the invasive SLF, we assemble an image dataset of the spotted lanternfly and its look-alikes, as summarized in Figure 1.

2.1 Data Overview

The first iteration of LANTERN-RD includes a total of 5 insect classes: *L. delicatula*, *A. figurata*, *P. imperialis*, *U. ornatrix*, and *H. lineata*, with 1187, 1501, 672, 520, and 1970 images, respectively. The distribution is visualized in Figure 2, and the total size of the dataset is 5850 images, with each image containing exactly one insect. Sample images of the dataset are displayed in Figure 1. The data are diverse, including images of insects in different poses, assorted backgrounds, and varied lighting conditions, allowing for the training of generalized models and the generation of sub-datasets for specific tasks. The dataset is presented as several files of image URLs, grouped by class, alongside a corresponding document outlining labels.

2.2 Pipeline

In order to curate a high-quality dataset, we employ a robust pipeline to gather, clean, and compile data.

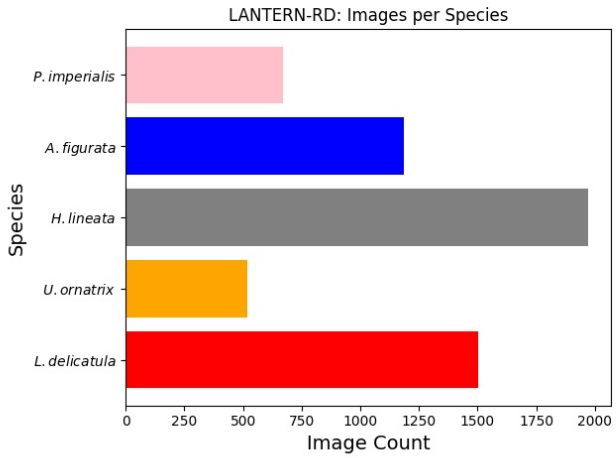

Figure 2: Distribution of Images Per Species.

Extraction. Raw image data are sourced from Bing Images and LAION [23, 24]. A number of strategic queries are utilized to capture a large volume of image URLs, which are stored separately according to class. However, the collected data are noisy and must be filtered before proceeding.

Automated Filtration. In the first step of cleaning, we use several automated techniques to eliminate extraneous data. We begin with image hashing to remove duplicate and near-duplicate images. Next, when image captions are available (as with the LAION data), we parse these captions and eliminate data containing phrases linked to insects other than the queried class. Through this process, we are able to remove a large volume of noisy data.

Additional Filtration. We parse the remaining data via several methods, removing the remaining extraneous images. For example, “Tiger Moth" also happens to be the name of a popular 20th century biplane, and images of this aircraft were highly represented in the raw data before additional cleaning. We also carefully inspect each image, applying established identification methods to verify that the insect pictured belongs to the label and ensuring that no personally identifiable information remains available.

The filtration process leaves a collection of image URLs that constitute the final, cleaned dataset. This modular pipeline allows for the efficient integration of new data sources such as community image submissions, enables expansion to additional classes of look-alike species, and takes strides to ensure that the data are clean and useful for training.

3 Baseline Experiments

Alongside the dataset, we train a baseline CNN using the VGG16 [14] model architecture (under the MIT License). The data are randomly split for training, validation, and testing, with a ratio of 60%, 20%, and 20%, respectively. Before training, the image data are augmented. Each class pictured in Figure 1 receives a numerical label, and the CNN is trained according to the categorical cross entropy loss function for 50 epochs at a rate of 1e-3. We train on Tesla K80 GPUs within Google Colaboratory.

3.1 Results

The model achieved an overall test accuracy of 97.20%. The per-class F1 scores are summarized in Table 1.

4 Mobile Implementation

We additionally propose a simple mobile application that implements the classifier trained in Section 3. Sensationalization of the spotted lanternfly in affected areas has resulted in the public becoming

Species	F1 Score
L. delicatula (Spotted Lanternfly)	0.983
A. figurata (Tiger Moth)	0.946
P. imperialis (Pink Underwing)	0.935
U. ornatrix (Bella Moth)	0.984
H. lineata (White-Lined Sphinx Moth)	0.989

Table 1: The preliminary CNN achieves the above results per-class.

more closely involved in management efforts, with programs such as New Jersey's "Stomp It Out!" [21] calling on citizens to exterminate sighted SLFs. As a result, it is imperative that the public is able to accurately identify this invasive insect against look-alike species, particularly to prevent well-intentioned citizen scientists from unnecessarily harming wildlife.

Features. The app allows users to directly capture or upload an image of an insect suspected to be an SLF. Subsequently, the picture is fed as an input to the classifier, and the user is notified of the class prediction. All user data are kept private and operations are run locally.

This app presents one useful application of LANTERN-RD in training deep learning models to assist in efforts to contain the invasive SLF. We hope that such an application motivates further advancements from the computer vision community.

5 Conclusion

In this paper, we have introduced LANTERN-RD, a curated dataset consisting of diverse images of the invasive spotted lanternfly and visually similar insects. This dataset contains 5850 images of the spotted lanternfly and four visually similar insects, and is curated via an efficient pipeline that is scalable to additional data sources and new classes. A baseline classifier trained on LANTERN-RD achieves a 97.20% test accuracy. This validates that datasets such as LANTERN-RD will enable a wide array of computer vision applications that have positive downstream impacts on efforts to contain the invasive spotted lanternfly. To explore one avenue of computer vision applications for the invasive SLF problem, we implement the preliminary classifier into a simple app designed for users to use their mobile devices to rapidly understand whether or not an insect is the SLF. This assists in research activities and boosts caution on behalf of citizen scientists. We warn against abuse of the dataset and model; researchers are urged to heed evolving guidelines on spotted lanternfly mitigation. Future work would include solicitation of community image submissions in order to expand the scale of the dataset, particularly in new and underserved insect classes, and for different stages of the SLF lifecycle. We call on ecologists and the computer vision community to come together in applying LANTERN-RD for deep learning tasks poised to increase the efficiency, reliability, and scale of efforts to contain the invasive SLF.

References

[1] Lawrence Barringer and Claire M Ciafré. Worldwide Feeding Host Plants of Spotted Lanternfly, With Significant Additions From North America. *Environmental Entomology*, 49(5):999–1011, 2020.

[2] Hannah J Broadley, Juli R Gould, Liam T Sullivan, Xiao-yi Wang, Kim A Hoelmer, Mauri L Hickin, and Joseph S Elkinton. Life History and Rearing of Anastatus orientalis (Hymenoptera: Eupelmidae), an Egg Parasitoid of the Spotted Lanternfly (Hemiptera: Fulgoridae). *Environmental Entomology*, 50(1):28–35, 2020.

[3] Thomas Y. Chen. Monarchnet: Differentiating monarch butterflies from butterflies species with similar phenotypes. *CoRR*, abs/2201.10526, 2022. URL `https://arxiv.org/abs/2201.10526`.

[4] Eric H Clifton, Ann E Hajek, Nina E Jenkins, Richard T Roush, John P Rost, and David J Biddinger. Applications of Beauveria bassiana (Hypocreales: Cordycipitaceae) to Control Populations of Spotted Lanternfly (Hemiptera: Fulgoridae), in Semi-Natural Landscapes and on Grapevines. *Environmental Entomology*, 49(4):854–864, 2020.

[5] Rachel T. Cook, Samuel F. Ward, Andrew M. Liebhold, and Songlin Fei. Spatial dynamics of spotted lanternfly, lycorma delicatula, invasion of the northeastern united states. *NeoBiota*, 70:23–42, 2021.

[6] Miriam F Cooperband, Ron Mack, and Sven-Erik Spichiger. Chipping to Destroy Egg Masses of the Spotted Lanternfly, Lycorma delicatula (Hemiptera: Fulgoridae). *Journal of Insect Science*, 18(3), 2018.

[7] Surendra K. Dara, Lawrence Barringer, and Steven P. Arthurs. Lycorma delicatula (Hemiptera: Fulgoridae): A New Invasive Pest in the United States. *Journal of Integrated Pest Management*, 6(1), 2015.

[8] Joseph A Francese, Miriam F Cooperband, Kelly M Murman, Stefani L Cannon, Everett G Booth, Sarah M Devine, and Matthew S Wallace. Developing Traps for the Spotted Lanternfly, Lycorma delicatula (Hemiptera: Fulgoridae). *Environmental Entomology*, 49(2):269–276, 2020.

[9] Melody A Keena and Anne L Nielsen. Comparison of the Hatch of Newly Laid Lycorma delicatula (Hemiptera: Fulgoridae) Eggs From the United States After Exposure to Different Temperatures and Durations of Low Temperature. *Environmental Entomology*, 50(2):410–417, 2021.

[10] Devin Kreitman, Melody A Keena, Anne L Nielsen, and George Hamilton. Effects of Temperature on Development and Survival of Nymphal Lycorma delicatula (Hemiptera: Fulgoridae). *Environmental Entomology*, 50(1):183–191, 2020.

[11] Ashley Leach and Heather Leach. Characterizing the spatial distributions of spotted lanternfly (Hemiptera: Fulgoridae) in Pennsylvania vineyards. *Scientific Reports*, 10(1), 2020.

[12] Houping Liu. Oviposition Substrate Selection, Egg Mass Characteristics, Host Preference, and Life History of the Spotted Lanternfly (Hemiptera: Fulgoridae) in North America. *Environmental Entomology*, 48(6): 1452–1468, 2019.

[13] Houping Liu. Seasonal Development, Cumulative Growing Degree-Days, and Population Density of Spotted Lanternfly (Hemiptera: Fulgoridae) on Selected Hosts and Substrates. *Environmental Entomology*, 49(5):1171–1184, 2020.

[14] Shuying Liu and Weihong Deng. Very deep convolutional neural network based image classification using small training sample size. In *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 730–734, 2015. doi: 10.1109/ACPR.2015.7486599.

[15] Robert Malek, Joe M Kaser, Hannah J Broadley, Juli Gould, Marco Ciolli, Gianfranco Anfora, and Kim A Hoelmer. Footprints and Ootheca of Lycorma delicatula Influence Host-Searching and -Acceptance of the Egg-Parasitoid Anastatus orientalis. *Environmental Entomology*, 48(6):1270–1276, 10 2019.

[16] Charles J Mason, Brian Walsh, Joseph Keller, John J Couture, Dennis Calvin, and Julie M Urban. Fidelity and Timing of Spotted Lanternfly (Hemiptera: Fulgoridae) Attack Patterns on Ornamental Trees in the Suburban Landscape. *Environmental Entomology*, 49(6):1427–1436, 2020.

[17] Alexander Mathis, Mert Yüksekgönül, Byron Rogers, Matthias Bethge, and Mackenzie W. Mathis. Pre-training boosts out-of-domain robustness for pose estimation. *CoRR*, abs/1909.11229, 2019.

[18] Kelly Murman, Gregory P Setliff, Cathryn V Pugh, Michael J Toolan, Isaiah Canlas, Stefani Cannon, Leslie Abreu, Miranda Fetchen, Longwa Zhang, Melissa L Warden, Matthew Wallace, Jacob Wickham, Sven-Erik Spichiger, Emelie Swackhamer, Daniel Carrillo, Allison Cornell, Nathan T Derstine, Lawrence Barringer, and Miriam F Cooperband. Distribution, Survival, and Development of Spotted Lanternfly on Host Plants Found in North America. *Environmental Entomology*, 49(6):1270–1281, 2020.

[19] Laura J Nixon, Heather Leach, Caitlin Barnes, Julie Urban, Danielle M Kirkpatrick, Dalton C Ludwick, Brent Short, Douglas G Pfeiffer, and Tracy C Leskey. Development of Behaviorally Based Monitoring and Biosurveillance Tools for the Invasive Spotted Lanternfly (Hemiptera: Fulgoridae). *Environmental Entomology*, 49(5):1117–1126, 2020.

[20] Mohammad Sadegh Norouzzadeh, Anh Nguyen, Margaret Kosmala, Alexandra Swanson, Meredith S. Palmer, Craig Packer, and Jeff Clune. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences*, 115(25): E5716–E5725, 2018.

[21] New Jersey Department of Agriculture. Spotted lanternfly, 2021. URL `https://www.nj.gov/agriculture/divisions/pi/prog/pests-diseases/spotted-lanternfly/`. Accessed: 2022-03-31.

[22] Pennsylvania Department of Agriculture. Spotted lanternfly quarantine, 2022. URL `https://www.agriculture.pa.gov/Plants_Land_Water/PlantIndustry/Entomology/spotted_lanternfly/quarantine/Pages/default.aspx`. Accessed: 2022-04-02.

[23] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.

[24] Christoph Schuhmann, Romain Beaumont, Cade W Gordon, Ross Wightman, mehdi cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Mitchell Wortsman, Richard Vencu, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. URL `https://openreview.net/forum?id=M3Y74vmsMcY`.

[25] Erica C Smyers, Julie M Urban, Andrew C Dechaine, Douglas G Pfeiffer, Stephen R Crawford, and Dennis D Calvin. Spatio-Temporal Model for Predicting Spring Hatch of the Spotted Lanternfly (Hemiptera: Fulgoridae). *Environmental Entomology*, 50(1):126–137, 2020.

[26] Osariyekemwen Uyi, Joseph A Keller, Anne Johnson, David Long, Brian Walsh, and Kelli Hoover. Spotted Lanternfly (Hemiptera: Fulgoridae) Can Complete Development and Reproduce Without Access to the Preferred Host, Ailanthus altissima. *Environmental Entomology*, 49(5):1185–1190, 2020.

[27] Tewodros T Wakie, Lisa G Neven, Wee L Yee, and Zhaozhi Lu. The Establishment Risk of Lycorma delicatula (Hemiptera: Fulgoridae) in the United States and Globally. *Journal of Economic Entomology*, 113(1):306–314, 2019.

[28] Hang Yu, Yufei Xu, Jing Zhang, Wei Zhao, Ziyu Guan, and Dacheng Tao. AP-10K: A benchmark for animal pose estimation in the wild. *CoRR*, abs/2108.12617, 2021.

[29] Xiaoyuan Yu, Jiangping Wang, Roland Kays, Patrick A. Jansen, Tianjiang Wang, and Thomas Huang. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences*, 115(25):E5716–E5725, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Limitations discussed throughout the work and particularly in Section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] Though we design this work for positive ecological impact, we recognize and indicate negative impacts in Section 5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The raw dataset is openly-accessible. This is work-in-progress research and additional models and data will be made open-source to the public for reproducibility and research.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Training details specified for benchmark model in Section 3.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A] We propose a benchmarked dataset in this work and do not conduct traditional experiments.

   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 2.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes] See Section 3
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] We discuss in Section 2 the pipeline used to remove such data.

5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]