Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

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

Original papers

Development and validation of a neural network for the automated detection of horn flies on cattle

E.T. Psota^{a,*}, E.K. Luc^b, G.M. Pighetti^b, L.G. Schneider^b, R.T. Trout Fryxell^b, J.W. Keele^c, L. A. Kuehn^c

^a University of Nebraska, Lincoln, NE, USA

^b University of Tennessee, Knoxville, TN, USA

^c USDA, ARS, U.S. Meat Animal Research Center, Clay Center, NE, USA

ARTICLE INFO

Keywords: Deep learning Fly counting Horn flies Precision livestock

ABSTRACT

When the number of horn flies that blood feed on cattle exceeds the economic threshold, they can adversely affect the health and wellbeing of their hosts. Excessive horn fly burdens also lead to reduced weight gain and, consequently, diminished profits for livestock producers. Effective management and treatment require reliable surveillance methods for estimating the degree of horn fly burden (i.e., counting the number of flies on cattle). Traditionally, these estimates are obtained through human visual estimation, either in-person or by counting images on a photo; however, these methods are costly both in terms of time and labor and are prone to subjectivity and bias. In contrast, automated methods can expedite the counting process and remove subjectivity and bias. To this end, a 2-stage method is presented here that uses computer vision and deep learning to identify the location of flies in digital images. The first stage segments the salient cow from all other parts of the image to remove flies on neighboring cattle from consideration. The second stage processes full-resolution patches of the original image and produces a heat map at the location of flies in the images. The method was trained on a set of 375 human-annotated images and tested on 120 images, where significant variation was observed amongst the human scorers. Counting results are compared to four separate human scorers and demonstrate that the neural network produces consistent results and that the method is reliable. Thus, the developed method can be used for monitoring changes in horn fly populations over time by anyone and allows for increased rigor and repeatability. An examination of individual images where the method was closest to and farthest from the human counts provides valuable insights regarding photographic processes that lead to success and failure.

1. Introduction

Horn flies (*Haematobia irritans* L.) are obligate ectoparasites of cattle remaining on their hosts for their entire adult life, blood feeding day and night, and even mating on the host (Bruce, 1964; Cupp et al., 1998). They can occur in large numbers during the summer and fall months, often reaching > 1000 flies per animal. Horn flies have negative physiological and economic consequences on cattle. Horn flies blood feed on cattle > 30 times a day, which alters the behaviors of animals, decreases milk production and weight gain in pastured cattle (Cupp et al., 1998), and allows for mechanical transmission of pathogens such as *Staphylococcus aureus*---causing mastitis and *Salmonella* (Gillespie et al., 1999; Edwards et al., 2000; Oliver et al., 2005; Anderson et al., 2012; Olafson et al., 2014), which can lead to zoonotic illness in humans who consume

beef and dairy products (Omer et al., 2017). Treating heifers with insecticidal ear tags for horn fly control resulted in 14% higher weight gain than untreated heifers allowing \$5.74 to \$8.38 return to the producer for each \$1 spent on horn fly control (DeRouen et al., 2003). The economic threshold at which horn flies cause economic damage to cattle can be as low as ten flies per animal if pathogens are being transmitted (Gordon et al., 1984); however, it is traditionally accepted as 200 flies per the side of an animal (Haufe, 1979; Kunz et al., 1991). The most recent research examining the cost of horn flies to the cattle industry was conducted in 1991 and estimated their impact to be in the range of \$700 million (Arther, 1991) to \$876 million annually (Kunz et al., 1991; \$1.5 billion and \$1.6 billion in 2018 dollars respectively). While staggering, these values may still underestimate the economic impact on producers as evidenced by a recent review of economic expenses and losses caused

https://doi.org/10.1016/j.compag.2020.105927

Received 26 August 2020; Received in revised form 23 November 2020; Accepted 24 November 2020 Available online 7 December 2020 0168-1699/© 2020 Elsevier B.V. All rights reserved.







^{*} Corresponding author. *E-mail address*: epsota@unl.edu (E.T. Psota).



Fig. 1. An example of the presentation of horn flies in cameras with various resolutions. These results assume a 2×3 aspect ratio with the cow's body captured entirely from head to tail.

by stable flies (Taylor et al., 2012).

Current horn fly management strategies are designed to take advantage of their obligatory ectoparasitic behavior (Geden and Hogsette, 2001). Treatments to control adult flies are often applied directly to the host, as horn flies tend to be found on the shoulders, midline, and underbelly, or targeted to the immature stage using feed-through insect growth regulators. The use of insecticides such as pyrethroids, organophosphates, and cyclodienes is the primary strategy for horn fly prevention and control. Importantly, in order to evaluate different methods of horn fly management and health and welfare impacts resulting from horn fly feeding, surveillance of these flies needs to be easy, accurate, and precise. Surveillance methods for horn flies have slowly developed with technology, but adoption by producers is rare and often limited to researchers evaluating control products. Current methods for horn fly surveillance are limited to (1) training a person to estimate fly populations on an animal and (2) taking a digital image of that animal and counting flies from the image of that animal (Mochi et al., 2009; Mullens et al., 2016; Smythe et al., 2017). Both procedures have advantages and disadvantages, but the primary and concerning disadvantages are reproducibility, accuracy, and adoption by producers. Since the first step of all integrated pest management methods is monitoring and establishing a surveillance program for pests (Geden and Hogsette, 2001), there is a need to develop a quicker, accurate, and user-friendly method for assessing horn fly populations (and other pests) on cattle.

Computer vision and machine learning have been successful with a number of pests and pathogens (Ding and Taylor, 2016; Cheng et al., 2017; Yang et al., 2019), and is already used by beef and dairy producers to monitor for health and welfare scores (Spoliansky et al., 2016). Because horn fly infestation levels can vary significantly among animals within a herd (Pruett et al., 2003), the development of an automated, accurate, and precise method would aid producers and other stakeholders in determining horn fly densities and monitoring for animal health and welfare concerns as well as managing insecticide resistance and treatment effects. Furthermore, computer vision could facilitate selection for increased cattle horn fly resistance through a national sire evaluation for horn fly carrying capacity. Digital images captured on registered cattle with pedigree and genotypes previously recorded could

be exploited with computer vision to make a national sire evaluation feasible. This development can aid in maintaining sustainable agriculture by improving the natural resistance of cattle, reducing the use of pesticides, lowering disease transfer, and improving efficiency.

Here we developed and validated a convolutional neural networkbased approach to fly counting by incorporating a diverse set of annotated images (e.g., geography, breed, lighting, fly species, environmental conditions, etc.) with digital count data. We expect that this neural network can be used for monitoring horn flies and that it will be more efficient than, and equally (if not more) reliable than, traditional phenotyping methods at quantifying horn fly infestation levels on cattle.

2. Materials and methods

2.1. Image capture

The goal when assessing horn fly burden is to count the number of horn flies on the body surface of a cow. It is impossible to count all the flies on all aspects of the entire animal. Here, the two traditional views for assessing flies on the animal, side and underbelly, are used in this study. It is assumed that the overall fly burden can be approximated from the single capture or that multiple images from different angles could be combined for a more accurate count. However, even when considering just a single image, capturing sufficiently detailed images of horn flies is challenging. A massive size mismatch obviously occurs between cows and horn flies, where the average head to tail length of a cow is 2.5 m and the average length of a horn fly is only 4 mm. The implications for image capture are that if, for example, a side view of a cow is perfectly captured in a 24-MegaPixel (MP) image (4000 \times 6000 pixels), each horn fly on that cow would be contained within a 10×10 pixel area. A second challenge exists with the body site preference of flies which varies between the back, side, and belly areas. Flies on the back and belly are particularly difficult to capture in side-view images, due to their tendency to overlap and occlude one another.

Fig. 1 illustrates the visual presentation of horn flies assuming a head-to-tail image capture of a cow. It should be noted that these images are subsampled from a sharp, high-resolution photo, so they represent

Computers and Electronics in Agriculture 180 (2021) 105927



Fig. 2. A flowchart illustrating the processing stages involved in the proposed horn fly detection method. The original image is processed by a Salient Cow Segmentation Network in order to isolate the center cow from the background. The original image is also broken up into 1024×1024 overlapping tiles and each tile is processed by the Fly Detection Network to detect the location of horn flies. The detection tiles are then combined into an image and masked by the salient cow mask to eliminate any fly detections outside the area of the salient cow. The image at the bottom illustrates the combination of segmentation (where flies are colored red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the best-case capture at their respective resolutions. Furthermore, resolutions>100MP may introduce considerable challenges due to limitations on lens sharpness (resolving capabilities) and the ability of photoreceptors on the image sensor to gather enough light to overcome noise. On the other side of the spectrum, this figure illustrates that 8.64-megapixel images do not provide enough detail to reliably differentiate flies from other artifacts.

At this time, the vast majority of consumer cameras are available within the 54–13.5-megapixel range. In our experiments, we used Nikon Coolpix P1000 16.0- Megapixel Digital Camera (Nikon, Tokyo, Japan) to capture images in the field. Distances from the cow ranged from 0.3 m to 0.9 m for underbelly pictures and 3.1 m to 4.6 m for side pictures.

2.2. Detection and counting method

The proposed fly detection and counting method requires two parallel stages of processing. The stages, illustrated in Fig. 2, are used to (1) isolate the salient cow from both the background and other cows in the image and (2) identify horn flies. Here, we define the salient cow as the subject animal or the cow we are attempting to analyze.

To detect the salient cow, a semantic segmentation network classifies pixels as either salient cow or background. The network used for salient cow segmentation is based on the DeepLabV3 + architecture (Chen et al., 2018) with a ResNet18 pre-trained front end (He et al., 2016). This particular network was chosen due to its exceptional performance in terms of segmenting objects in images, particularly when the scale of the object varies considerably, as is the case for the problem of salient cow segmentation when both close-up and far-away images are used in the

same dataset. The cow images are down-sampled from their original resolution down to one that is 448x448 where, prior to down-sampling, the image is made square by padding either the top and bottom or the left and right side with white pixels. The network is trained to output pixel-wise classification results that label each pixel as either salient cow or background. It is assumed that the salient cow occupies the center pixel in the image.

In parallel with salient cow segmentation, the image is broken up into a collection of overlapping tiles, as shown below the original image in Fig. 2. This is required because the original image is too large to process via a single pass on GPU hardware due to limitations on RAM. The network instead processes tiles with size 1024×1024 in small batches and produces tiled fly detection maps. The network used for fly detection is also based on the DeepLabV3 + architecture (Chen et al., 2018) with a ResNet18 pre-trained front end (He et al., 2016). Finally, these tiles are recombined (untiled) to reconstruct an image with the same dimensions as the original image and the segmentation output is used to mask out areas of the image that do not pertain to the salient cow. The output image, shown at the bottom of Fig. 2, illustrates the effects of masking using black pixels and fly detections are illustrated by red dots.

2.3. Datasets

To train the segmentation network, a combination of images from the COCO dataset (Lin et al., 2014) and our own dataset were used. In either case, a single target was chosen from the list of instances in each image, the centroid of its mask was identified, and the image was shifted



Fig. 3. A set of 20 samples from the training set used to train the salient cow segmentation network. The entire training set consists of 4,687 images.

to place the centroid in the middle of the image. Here, the centroid is defined as the center of mass for all pixels belonging to the target of interest, i.e., the average row and column position of the target's spatial location in the image.

From the COCO dataset, images of cows, zebras, sheep, goats, elephants, and giraffes were used. This effectively expands the capabilities of the network to handle a wider range of 4-legged animals and encourages the network to generalize to unexpected appearances. Fig. 3 shows 20 samples of the original images and their segmentation targets. The overall set of images used for training consists of 4,687 images.

Pictures (n = 375), taken of both the back and ventral midline area of Holstein and Jersey dairy cattle and beef cattle with color patterns representative of common industry breeds were used to train the fly detection network. Samples of images from the dataset are shown in Fig. 4, illustrating the body presentation of cattle in the images. In some views, the salient cow is fully captured within the bounds of the image space whereas, in others, only a small section of the cow's body is provided.

Images were annotated by personnel using Microsoft Paint (Microsoft, Redmond, WA) to train the neural network. Pure red (255 Red, 0 Green, and 0 Blue) was used to mark the flies seen on the center cow in each image, and images were saved using the lossless compression Portable Network Graphics (png) format. Pure colors are almost never seen in natural images and lossless compression ensures that each pure red annotated pixel can be reliably extracted prior to creating target images for neural network training.

Human annotation styles vary from user to user and the precision of the annotation locations is not exact. Therefore, in this work each of the annotations was converted to a standardized form, as shown in Fig. 5. In each image triplet, the original image is on the left, the human annotation is in the middle, and the standardized form is shown on the right. Notice that humans used either circles, lines, or full-coverage blobs to annotate each fly. Regardless of the human style, the standardized form is created by extracting the centroid of each connected component (red mark) and replacing the human annotation with a circular marker at that position. The circles' diameters were manually adjusted to a single value for each image in order to closely match half of the head-to-tail length of flies in the images.

To account for the imprecision of the human annotator, a buffer region (shown in gray in Fig. 5) was placed around each standardized annotation. Within this region, the network is allowed to label the pixels as either "fly" or "no fly" without penalty. The rationale behind this methodology is to make sure the network is not training itself to mimic the human annotation errors by outputting a circle with the exact same diameter and centroid location that the human chose, knowing that the results of human annotators are variable.

2.4. Human-based predictions and sensitivity analysis

Pictures (n = 120) were annotated and counted by four individuals, along with the neural network, in order to determine interrater and intrarater reliability. Both the individuals and the neural network repeated the process three times with the same set of pictures. The total number of flies counted by an individual rater and the neural network for one picture were averaged amongst the repetitions. The coefficient variation (CV) was then calculated by dividing the standard deviation by



Fig. 4. A set of 16 samples from the training set used to train the horn fly detection network. The entire training set consists of 375 images.

the mean for every rater's picture. The CV was calculated to determine a rater's dispersion around the mean for each picture counted. A lower CV means the rater had a more precise estimate every repetition. Mixed model analysis of variance in SAS 9.4 (PROC GLIMMIX; Cary, NC) was used to assess interrater reliability, with the random effect of image and rater*image, while blocking on the image. To assess intrarater reliability, a no variable model in SAS 9.4, (Cary, NC) was used with random effect of intercept. An interclass correlation (ICC) was then calculated for each rater by dividing the estimate by the residual. ICC was used to determine the correlation within a rater. An ICC > 0.75 is considered excellent, while an ICC < 0.40 is considered to be poor (Portney and Watkins, 2009).

3. Results

3.1. Description of network training and post-processing

The cow segmentation network was implemented and trained in MATLAB using the Deep Learning Toolbox (Mathworks, 2020). Segmentation results are produced via a 2-channel output with softmax classification and cross-entropy loss at each pixel location. Training was performed using Stochastic Gradient Decent with Momentum (SGDM) with minibatches of eight images and an initial learning rate of 10^{-3} . Augmentations included random rotations ($\pm 25^{\circ}$), random horizontal reflections, random horizontal and vertical scales (0.5–1.5), random shearing (± 0.15), and random horizontal and vertical translation (± 10 pixels). These augmentations are meant to maintain the upright presentation of animals in the center of the image while spatially varying their presentation as much as possible within reason. The trained network used in this work was obtained after processing 100 epochs of the dataset.

MATLAB was also used to train and evaluate the horn fly detection network. Both network inputs and outputs are trained with 1024 × 1024 overlapping tiles extracted from the full-resolution original image. Overlaps are chosen so that only the middle 512 × 512 section of each output is used to reconstruct the full-resolution output. Thus, the 256pixel wide border region is used only to give the center portion of the tiles more spatial context to detect flies. So, for example, a 2048 × 3072 input image would be constructed using 24 non-overlapping 512 × 512 tiles, and each of these is extracted from the center of 24 overlapping 1024 × 1024 tiles. Fig. 2 demonstrates the image tiling process on a sample image.

For horn fly detection, augmentations of tiles included random rotations ($\pm 45^{\circ}$), random horizontal reflections, random horizontal and



Fig. 5. A set of 18 sample crops from the training set used to train the horn fly detection network. The left sides illustrate the original image crops and the middle contains user annotations. Because the annotation marks vary greatly between annotators, they were converted to a standard form shown in the right sides. The white regions are marked as flies, the black regions are marked as background, and the gray regions illustrate regions of ambiguity. During training, the gray regions do not affect the cost and the network is free to assign this as fly or background without penalty.

vertical scales (0.5–1.5), random shearing (±0.2), and random horizontal and vertical translation (±100 pixels). The reason rotations were limited to \pm 45° was to preserve the orientation of horn flies, which typically orient downward towards the ground when feeding on the sides of cattle (Geden and Hogsette, 2001).

To convert a pixel-wise segmentation of horn flies into fly counts, regional maximums are extracted in post processing. This assumes that the confidence of the output classification is strongest near the center of each fly and that confidence decreases for pixels that are farther from the center until it eventually crosses 0.5 and classifies pixels far from the center as background. In this case, the pixel that is equal to the maximum value within its own region is likely to correspond to the center point of the fly. To detect these region maximum locations, during forward inference a 9×9 max-pooling layer is added to the end of the network and its output is concatenated with the original output. By comparing these two outputs pixel-by-pixel, regional maximums can be detected by only considering pixel locations where the two outputs are equal. Because the radius of the max-pooling region is equal to four pixels, the method cannot detect flies that are less than five pixels apart from one another. In practice, this is acceptable since flies that are<5 pixels wide would be nearly impossible to differentiate from artifacts.

3.2. Performance evaluation

Salient cow detection was trained on 4,270 images that were extracted from the COCO dataset and 417 images that were labelled from our fly detection dataset. Fig. 6 illustrates the results of the trained network on 14 randomly chosen samples from the dataset. The

segmentation results successfully identify and mask the salient animal in each of the examples, even in the presence of nearby animals and partial obstruction.

The results in Fig. 7 demonstrate the network's performance on 10 images from the 103-image test set. The performance on the test set allows us to evaluate the level of overfitting that might occur during training. While there are some unwanted artifacts and errors in the results, the level of success on the test set indicates that the network was able to generalize its understanding of the salient animal in images.

Numerically, the salient cow segmentation method produces the precision and recall presented in Fig. 8 over a wide range of thresholds from 0.01 to 0.99. Fig. 8 also presents the Intersection over Union (IoU) of the salient object segmentation, which is perhaps a better metric for evaluating segmentation performance when the sizes of targets varies. These results, obtained on the test set, indicate that a peak IoU of 0.83 can be obtained when using a 0.5 threshold to differentiate between salient cow and background. Using this same threshold, the precision is approximately 0.9 and the recall is 0.92.

3.3. Human-based predictions and sensitivity analysis

When ICC was calculated for each individual rater, all raters had an excellent ICC (Table 1); meaning, for every picture, a rater counted a similar number of horn flies every repetition. Combined, all raters counted a similar number of horn flies for each picture, with a combined ICC of 0.95. When the CV was compared amongst the raters, raters were not different from one another and counted pictures very similarly (P = 0.9; Table 1). The neural network always had a CV of 0, regardless of the



Fig. 6. The results for 30 samples from the salient cow segmentation training set. Original images, ground truth segmentations, and network outputs are illustrated in each sample from left to right. The network outputs are practically indistinguishable from the ground truth in most cases.



Fig. 7. The results for 10 samples from the salient cow segmentation validation set. Original images, ground truth segmentations, and network outputs are illustrated in each sample from left to right. While most of the network outputs resemble the ground truth, there are some errors in situations where nearby cows overlap with the salient cow (4th row, 1st column and 5th row, 2nd column). Overall, results illustrate strong generalizability from training to validation sets.



Fig. 8. Precision vs. recall averaged over the 103-image test set (left). Threshold values 0.1, 0.2, ..., 0.9 are labeled on the curve. Intersection over Union (IoU) is also plotted (right) as a function of threshold for values ranging from 0.01, 0.02, ..., 0.99.

Table 1
Comparison of human raters and the neural network mean fly count outputs for
assessing flies on cattle. Raters were consistent among themselves (high ICC
scores), but differed between one another and the neural network (P < 0.0001).

Rater	Mean	Standard Error	Letter Group	ICC
Rater 1	53.0	6.8	С	0.99
Rater 2	55.4	6.8	BC	0.99
Rater 3	46.8	6.8	D	0.99
Rater 4	69.5	6.8	А	0.99
Neural Network	50.0	6.8	В	0.99

*Raters that share similar letter groups do not differ at alpha = 0.05.

picture, meaning it counted the exact same number of horn flies for each repetition, which is expected from a deterministic algorithm. When assessing interrater reliability, raters and the neural network were significantly different to one another (P < 0.0001). Rater 2 counted similar to the neural network and with rater 1. Rater 4 routinely counted higher than the other raters, while rater 3 routinely counted lower (Table 1). The neural network has a strong reliability with consistent counts for each repetition, while human raters are more variable in their counts.

Examples of network failures and success shed light on how the performance is affected by the cow, the environment, and the image capture conditions. The nine images where the network was closest to the human annotation average are shown in Fig. 9. In all of these images, the cow is mostly white or light brown. This is unsurprising, because the dark flies stand out mostly on light fur. Seven of the nine images are of the belly region of the cow and the other two depict the cow in its entirety from a side view. When only the belly region is photographed, the flies appear larger in the image, thus they are easier for both the rater and the network to identify. In terms of image capture properties, each photo is exposed well (no blown-out highlights or noisy shadows) and focused properly on the subject. Again, this detail makes it relatively easy for both the rater and the network to identify individual flies and differentiate them from other artifacts.

The nine images where the network produced the worst results (relative to the average human annotator) are given in Fig. 10. The cows in the failure set are mostly darker than the cows in the success set. The

two exceptions are lighter cows that have dark patterns with small spots. Six of the nine failure images depict the side-view and the images are generally captured from farther away than successful images. This causes the flies to appear as little more than black dots on white fur or slightly lighter dots on black fur; however, it appears that the main contributor to network error is small patterns in the fur and dirt/debris.

A closer inspection of the successes and failures is provided in Figs. 11–13, which contains samples of true positives, false positives, and false negatives. The true positive examples in Fig. 11 include flies of various sizes and presentations. Some are large, crisp images of flies that are easily discernible, while others are merely dark blobs on a lighter background or vice versa.

The false positive set in Fig. 12 sheds some light on the most common type of error that the network tends to make. There are clearly some flies that the human annotators missed and the network detected, which is an encouraging sign. It also appears that the majority of the false positives that are clearly not flies come from small artifacts that stand out from the background. In some cases, the natural texture of the fur is mistaken for a fly. In other cases, it appears to be a speck of dirt.

Finally, the false negatives in Fig. 13 illustrate situations where the network fails to identify the fly. In nearly half of these samples, it is difficult to validate if the human annotation is correct. It may be easier to discern from the larger context of the image, but many of these small crops are blurry and/or depict very small flies. It is worth noting that, in a couple of samples, the distance between flies may be making it difficult for the regional maximum to separate flies from one another.

4. Discussion

To the best of our knowledge, this is the first attempt to use a deep neural network to count flies on livestock. The results demonstrate the potential of the method and an inspection of successes and failures point towards some generalizable lessons for further development. Under ideal conditions with large flies in sharply focused and properly exposed images, the proposed method performs as well as human annotators. Future improvements in digital imaging technology, such as higher resolution sensors, better noise handling, and improved autofocus will directly result in improved performance even if the processing framework is unchanged.



Fig. 9. The nine images from the testing set with the smallest relative error between the average human annotator and the network-based approach.

An inspection of the results shows that they are likely to improve with a more tightly constrained image capture process. For example, images of the side-view should be taken at an angle perpendicular to the cow. If the cow's shoulders are closer to the camera than the tail, it is likely that depth of field will become an issue and some region of the cow's body will be out of focus. Images captured in low light (indoor) conditions make it difficult to achieve sharp, low noise results. This is because the camera's only two mechanisms for handling low light are to open up the aperture or increase the gain, resulting in a narrow depth-offield or high-amplitude additive noise, respectively.

Current technologies exist that might improve the results, if integrated into both the image capture routine and processing framework. For example, multi-image super-resolution has shown promising performance on satellite images (Deudon et al., 2020) and similar techniques could likely be developed to merge multiple images of slowmoving targets like cattle into a 100MP or greater resolution image. Another technique that might improve performance is to encode the cow body location into each high-resolution tile. In this work, each tile was processed independent of every other tile and the detection method had no information regarding where that tile was located on the original cow. By encoding this location into the tile for processing, the detection method may have an improved level of expectation in terms of how flies should present themselves in the tiled image because horn flies are not uniformly distributed over the surface of the cow's body. Horn flies typically occur in higher densities on the animal's back and belly; this information could potentially improve the accuracy of the neural network in the face of ambiguous horn fly images.

Our computer vision system could be employed to screen large numbers of cattle to identify animals with extreme horn fly carrying capacity (both high and low) from multiple breeds and family lineages. Extreme horn fly carrying capacity animals subsequently could be tested for volatile semiochemicals to identify a diverse repertoire of attractants and repellants which could be synthesized and employed as part of a push–pull horn fly control program (Oyarzún et al, 2008). Alternatively, small groups of naturally attractant cattle could be employed to attract flies away from the main herd or conversely repellant cattle to dissuade flies. Horn flies use complex volatile chemicals to choose their targets. Exploiting this horn fly sensing and signaling mechanism may lead to cost effective horn fly control measures that are sustainable with little harmful impacts to the environment or beneficial insects.

5. Conclusion

Here we present a neural network for counting horn flies on dairy and beef cattle, which is an efficient and equally (if not more) reliable method as traditional quantifying methods. We were surprised to see that the annotation of digital images was relatively uniform by each rater, but intrarater reliability was low. Importantly, the neural network was significantly more repeatable and reliable compared to human annotators. Logical next steps are to validate the network with additional host species (e.g., beef cattle), and develop the network to identify and differentiate additional pest species (e.g., stable fly) and horn fly damage (e.g., mastitis). Using the network, we envision the creation of a relational and accumulating database which incorporates phenological data

Computers and Electronics in Agriculture 180 (2021) 105927



Fig. 10. The nine images from the testing set with the largest relative error between the average human annotator and the network-based approach.



Fig. 11. Random samples of true positives from the testing set.



Fig. 12. Random samples of false positives from the testing set.



Fig. 13. Random samples of false negatives from the testing set.

of horn flies from different farms using different management techniques. This resulting database will be essential in the future monitoring and modeling of horn flies and insecticide resistance and planning sustainable management options.

Declaration of Competing Interest

None.

Acknowledgements

This study was partly supported by OREI grant no. 2015-51300-24140 from the USDA National Institute of Food and Agriculture. Additional support came from the S1076 Regional Hatch Project (Fly management in animal agriculture systems and impacts on animal health and food safety). We appreciate the dairy producers for agreeing to participate in the collection of pictures. Thank you to all the undergraduate students (Shanon Park, Adam Allen, Diondria Richards, Brett Anderson, Amanda Young, Caitlin Zaring, Pari Baker, Sarah Fiedler, Calyn Fulton, Emma Graf, Jalyn Westmoreland, Kaitlin Cracchiola, Malone Young, Tanner Thornton), graduate students (Hannah Malcomson), and research specialists (Dave Paulsen, Wyatt Simmerman, Collin McCorkel, Natasha Bales) that assisted with the data collection and annotations.

References

- Anderson, K.L., Lyman, R., Moury, K., Ray, D., Watson, D.W., Correa, M.T., 2012. Molecular epidemiology of *Staphylococcus aureus* mastitis in dairy heifers. J. Dairy Sci. 95, 4921–4930.
- Arther, R.G., 1991. Management of horn fly resistance, p. 7pages. In: 40th Annu. Florida Beef Cattle Short Course Proc.
- Bruce, W.G., 1964. The history and biology of the horn fly, *Haematobia irritans* (L.): with comments on control. United States Department of Agriculture No. 157.
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H., 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In Eur, Conf. Comput. Vis, pp. 801–818.
- Cheng, X., Zhang, Y., Chen, Y., Wu, Y., Yue, Y., 2017. Pest identification via deep residual learning in complex background. Comput. Electron. Agric. 141, 351–356.
- Cupp, E.W., Cupp, M.S., Ribeiro, J.M.C., Kunz, S.E., 1998. Blood-feeding strategy of Haematobia irritans (Diptera: Muscidae). J. Med. Entomol. 35, 591–595.
- DeRouen, S.M., Foil, L.D., MacKay, A.J., Franke, D.E., Sanson, D.W., Wyatt, W.E., 2003. Effect of Horn Fly (*Haematobia irritans*) Control on Growth and Reproduction of Beef Heifers. J. Econ. Entomol. 96, 1612–1616.
- Deudon, M., Kalaitzis, A., Goytom, I., Arefin, M.R., Lin, Z., Sankaran, K., Michalski, V., Kahou, S.E., Cornebise, J., Bengio, Y., 2020. HighRes-net: recursive fusion for multiframe super-resolution of satellite imagery. arXiv Prepr. 2002.06460.

Ding, W., Taylor, G., 2016. Automatic moth detection from trap images for pest management. Comput. Electron. Agric. 123, 17–28.

- Edwards, J.F., Wikse, S.E., Field, R.W., Hoelscher, C.C., Herd, D.B., 2000. Bovine teat atresia associated with horn fly (*Haematobia irritans irritans* (L.))-induced dermatitis. Vet. Pathol. 37, 360–364.
- Geden, C.J., Hogsette, J., 2001. Research and extension needs for integrated pest management for arthropods of veterinary importance. In: Proc. a Work. Lincoln, Nebraska, April 12-14, 1991.
- Gillespie, B.E., Owens, W.E., Nickerson, S.C., Oliver, S.P., 1999. Deoxyribonucleic acid fingerprinting of *Staphylococcus aureus* from heifer mammary secretions and from horn flies. J. Dairy Sci. 82, 1581–1585.
- Gordon, D.V., Haufe, W.D., Klein, K.K., 1984. Determination of economic thresholds for horn fly control in western Canada: a farm level simulation approach. Candian J. Agric. Econ. 32, 399–421.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Eur. Conf. Comput. Vis., pp. 770–778.

- Haufe, E.O., 1979. Reduced productivity of beef cattle infested with horn flies En: Croome G.C.R., Holmes N.D., Research highlights, Agriculture Canadian Research Station, Lethbridges, Alberta, p. 61-63.
- Kunz, S.E., Murrell, K.D., Lambert, G., James, L.F., Terrill, C.E., 1991. Estimated losses of livestock to pests. In: Pimental, D. (ed.), CRC Handb. Pest Management Agric. Vol 1. CDC, Boca Raton, Florida, pp. 68–69.
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C. L., 2014. Microsoft coco: Common objects in context. In: European conference on computer vision. Springer, Cham, pp. 740–755.
- Mathworks. 2020. Deep learning toolbox TM user's guide how to contact MathWorks. (https://www.mathworks.com/help/pdf_doc/deeplearning/nnet_ug.pdf).
- Mochi, D.A., Monteiro, A.C., Simi, L.D., Sampaio, A.A.M., 2009. Susceptibility of adult and larval stages of the horn fly, *Haematobia irritans*, to the entomopathogenic fungus *Metarhizium anisopliae* under field conditions. Vet. Parasitol. 166, 136–143.
- Mullens, B.A., Soto, D., Gerry, A.C., 2016. Estimating field densities of *Haematobia irritans* (Diptera: Muscidae) using direct visual field counts versus photographic assessments. J. Med. Entomol. 53, 703–706.
- Olafson, P.U., Lohmeyer, K.H., Edrington, T.S., Loneragan, G.H., 2014. Survival and fate of *Salmonella enterica serovar* Montevideo in adult horn flies (Diptera: Muscidae). J. Med. Entomol. 51, 993–1001.
- Oliver, S.P., Gillespie, B.E., Headrick, S.J., Lewis, M.J., Dowlen, H.H., 2005. Prevalence, risk factors, and strategies for controlling mastitis in heifers during the periparturient period. Int. J. Appl. Res. Vet. Med. 3, 150–162.
- Omer, M.K., Álvarez-Ordoñez, A., Prieto, M., Skjerve, E., Asehun, T., Alvseike, O.A., 2017. A systematic review of bacterial foodborne outbreaks related to red meat and meat products. Foodborne Pathog Dis. 2018 Oct, 15(10), 598–611. doi: 10.1089/ fpd.2017.2393. Epub 2018 Jun 29. PMID: 29957085.
- Oyarzún, M.P., Quiroz, A., Birkett, M.A., 2008. Insecticide resistance in the horn fly: alternative control strategies. Med. Vet. Entomol. 22, 188–202.
- Portney, L., Watkins, M., 2009. Foundations of clinical research: applications to practice, 3rd edition. Pearson Prentice Hall, Upper Saddle River. USA, p. 892.
- Pruett, J.H., Steelman, C.D., Miller, J.A., Pound, J.M., George, J.E., 2003. Distribution of horn flies on individual cows as a percentage of the total horn fly population. Vet. Parasitol. 116, 251–258.
- Smythe, B.G., Urias, S., Wise, M.E., Scholljegerdes, E.J., Summers, A.F., Bailey, D.W., 2017. Comparing visual and digital counting methods to estimate horn fly (Diptera: Muscidae) populations on cattle. J. Med. Entomol. 54, 980–984.
- Spoliansky, R., Edan, Y., Parmet, Y., Halachmi, I., 2016. Development of automatic body condition scoring using a low-cost 3-dimensional Kinect camera. J. Dairy Sci. 99, 7714–7725.
- Taylor, D.B., Moon, R.D., Mark, D.R., 2012. Economic impact of stable flies (Diptera: Muscidae) on dairy and beef cattle production. J. Med. Entomol. 49, 198–209.
- Yang, F., Poostchi, M., Yu, H., Zhou, Z., Silamut, K., Yu, J., Maude, R.J., Jaeger, S., Antani, S., 2019. Deep learning for smartphone-based malaria parasite detection in thick blood smears. IEEE J. Biomed. Heal. Informatics. 24, 1427–1438.