The Farmwave Imaging System

Artificial Intelligence in Agriculture

February 1, 2021

Atlanta, GA

Executive Summary

AI, as part of the precision agriculture technology boom, intrigues the agricultural industry. Futurists and growers alike dream of an AI-powered fleet that can make timely adjustments for maximum efficiency. More than one marketing video has depicted sprayers that can identify, target, and spray only diseased crops to reduce the gallons of pesticides used, or combines that can automatically adjust their dynamics to reduce head loss based on grain variety.

In reality, these lofty goals still seem a long way off. AI for agriculture is still in its infancy, as researchers try to contend with the complex, multi-faceted decisions that must be made from planting all the way to harvest season.

Meanwhile, time, as usual, is not on the side of growers. Political turmoil, extreme weather conditions, and a shrinking farming population have contributed to ever thinner margins, with growers needing every advantage they can get just to survive until the next harvest.

Against this backdrop, skepticism about new technologies making big promises is understandable. No one wants to be an early adopter, especially when so many other technologies have failed to live up to their own marketing.

However, in this case, adoption of the time and cost-saving discoveries we have made cannot come soon enough.

Farmwave will unveil a series of breakthroughs achieved through painstaking, time-consuming research that ultimately resulted in bringing AI out of the lab and into the field. Our discussion will dispel myths about AI in agriculture, revealing what works and what doesn’t. Readers will come away with more realistic expectations concerning the limitations and capabilities of current AI technology and see for themselves how this particular technology focused on harvest can immediately start saving farmers money during harvest.

Table of Contents

Introduction

Agriculture is a very visual industry.

At every stage of the food value chain, industry workers with varying education levels, experience, and age are performing thousands of tasks and making decisions- primarily based on visual inspection. The quality and accuracy of these actions and decisions can vary greatly and have real economic consequences.

Industry innovators have responded with a “more is better” approach: add more sensors, add more automation, add more inputs like drones or satellite imagery. The result is the availability of more agricultural data than ever before, for growers who are less equipped to make sense of it all. In short, growers are often data rich, but information poor.

This paper will discuss how actionable insights derived from visual information can provide support for growing decisions through the use of purpose-built, practical artificial intelligence (AI) models. The development of foundational data sets can be leveraged to build sophisticated capabilities, such as the creation of yield prediction tools that can benefit all growers with a more profitable and efficient harvest- whether they are a smallholder or large co-op.

Learning about Learning

While a comprehensive discussion of the relative merits of AI, machine learning, and deep learning is outside of the scope of this paper, what follows is a brief review that will provide readers with a basic, working knowledge of these concepts.

Artificial Intelligence

Artificial intelligence encompasses several techniques and fields of study which are rapidly evolving, including machine learning and deep learning. We will explore each in turn.

John McCarthy coined the term “Artificial Intelligence” at the 1956 Dartmouth Artificial Intelligence (AI) conference, widely regarded as the birth of AI as a formalized field of study.[1] Today, artificial intelligence refers to the “theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.”[2]

Specifically for agriculture, AI is built to make decisions on when and how to apply capabilities exceeding that of humans in processing speed, accuracy, analysis, or work orchestration. These decisions must be made considering multiple data points received from sensors as well as human experience and input.

Introduction

AI, Machine Learning, and Deep Learning

Machine learning is a subset of AI that encompasses various algorithms designed to mimic human decisions. Deep learning is one of many modeling techniques used within machine learning algorithms. The goal of modern deep learning algorithms is to teach AIs how to identify objects such that they can approximate human learning behaviors.

A deep learning model must first learn about its environment before it can begin to process data. After learning about the data, the model could then start making decisions and correct its behavior based on feedback about the decisions it makes.

For example, if deep learning were used to detect an apple, the algorithm would first need to learn how to “look” at an image. Once the AI was able to process an image, it would then need to understand shapes to determine what an apple looked like. The AI would need feedback to understand that it has successfully understood vision and shape.

What would happen if an AI trained only on green Granny Smith apples encountered a red Honeycrisp apple?

The results would vary based on how the deep learning algorithm was trained during the feedback phase. If a person were taught to identify apples by repeatedly being shown green Granny Smith apples, but then saw a red Honeycrisp apple, it would be reasonable for that person to either guess that it was also an apple (based on the shape) or respond that they didn’t know. Deep learning models show the same behavior and would respond in a very similar manner.

On the other hand, if the purpose was to identify the type of apple, the person (or the AI), would not be able to specifically identify that the apple was a red Honeycrisp apple. The AI would need to be given feedback that it needed to now also understand the color of the apple to differentiate between a green Granny Smith apple and a red Honeycrisp apple. Similar to the first example requiring the AI to first learn vision before shape, the AI would need feedback to determine how to understand color before associating the color of the apple to the type of apple.

Data Sets and AI

Now that we have discussed AI, machine learning, and deep learning, we need to discuss the "food" an AI needs to ingest to learn: data sets. Data sets are used to train an AI, providing the interpretive instructions, permutative subroutines, and other information on a specific subject. Developing AI for the agricultural industry is especially challenging, since the consequences of misapplying it are so drastic- faulty decisions made on incomplete data sets could ruin an entire year's harvest. Yet, many commercial AI solutions base their technology on open source data sets that are publicly available.

There are many reasons an organization may be willing to incorporate these types of data sets into their technology: achieving first mover advantage, quickly gaining market share, or simply differentiating themselves from the competition. However, these sources are a double-edged sword. Due to their inherently open nature, they are susceptible to bias. Sources without stringent vetting can also produce highly inaccurate results. By contrast, developing data models that are custom-built for specific yield estimation activities like kernel counting for corn or harvest loss analysis for soybeans can produce highly accurate results.

ImageNet

One popular example of an open source data set is ImageNet. ImageNet has a vast catalog of thousands of categorized images, with descriptions and information that make it an attractive starting place for image recognition applications.

There are two important limitations, however. First, ImageNet's license specifies that it is intended to be used exlusively for research.[3] Our efforts to commercialize practical AI tools for agriculture meant that we needed to have a solid foundation that we could built on with both control over the image data set. Second, pre-trained models (using ImageNet data); Use these pre-trained models to train their own neural network (a process called transfer learning). However, this results in inherent biases in ImageNet does not claim to hold the copyright for these images. Much like Google Image Search, ImageNet retrieves search results, but that does not give a user the rights to use the images retrieved.

Another drawback of using a source like ImageNet is a vulnerability to inaccurate information. Since ImageNet draws from submitted images, there is no gatekeeping mechanism that ensures the image description is accurate. Whether malicious or innocuous in nature, a user could potentially submit photos that are inaccurately labeled, compromising the accuracy of the entire system.

Lastly, ImageNet and other sources like it are beholden to those submitting the images, relying upon them to take high quality imagery that is usable for research- an assumption that cannot be relied upon for systematic AI model building due to inconsistencies and inaccuracies that accompany such sources. Put another way, publicly available data sets simply cannot yield the same accuracy, quality, and quantity of images required to build the next generation of AI for agriculture.

Now that we've laid the foundation for understanding AI, we will put it in the context of the research journey undertaken to build Farmwave's yield estimation tools.

Research Questions and Goals

1. Can image recognition be used to identify and diagnose pathogen and pest infestation in crops?

2. Can a decision model be applied to artificial intelligence, automating a response providing decision support for data-based actions?

3. Can data points be applied from other sources such as location, weather, drone, satellite, soil moisture sensors, to add additional classifiers supporting this decision model?

4. Can capabilities developed for pest, pathogen, or weed recognition be enhanced to utilize additional data models to enable yield prediction and forecasting via measurement or counting?

The above research questions can be answered with a resounding “yes”. It is possible- but only with a sufficient quantity and quality of data. Only clean data can provide the ability to achieve and maintain accurate results.

Farmwave began utilizing open data sources to test initial algorithms being applied to image recognition. One of the initial yield estimation tools utilized vision computing to count kernels on an ear of corn. OpenCV (Open Source Computer Vision Library) an open source computer vision and machine learning software library[4], was used in early versions of the Farmwave kernel counter for counting kernels on ears of corn. The counter worked by enabling users to simply take a picture, rotate the ear 180 degrees, and take another photo, returning a result in mere seconds. The industry standard for accuracy accepted by researchers and growers for yield prediction in corn using hand-count methods was 72-80%. Today, Farmwave’s own internally-developed, proprietary tools are exceeding this accuracy rate using AI instead of open source packages.

This experience reinforced a bedrock principle for our AI development: in order for Farmwave to own data, control biases, and maintain a validated data source, it was necessary to begin building our own data set.

The Road to Data Collection

"Colleagues across academia and the AI industry are hammering away at the same concept: a better algorithm would make better decisions, regardless of the data. But I realized a limitation to this approach— the best algorithm wouldn't work well if the data it learned from didn't reflect the real world.

My solution: **build a better data set.**"

Dr. Fei-Fei Li
Co-Director of AI Institute, Stanford University
Former Head of AI, Google

Dr. Fei-Fei Li's statement succinctly captures the problem with AI research: the effort to collect clean data and achieve a high degree of accuracy is time-consuming, tedious, and manual, but absolutely necessary. Her conclusion was the same as ours: in order to achieve the goals we set, we needed to build a better data set from the ground up.

Our Experience: Early Data Set Evaluations

Farmwave began working with various universities throughout the United States as well as research groups located in the United States and France. The University of Georgia College of Agriculture and Environmental Sciences has maintained a database called the Consortium Internet and Imagery Database System (CIIDS) for over 18 years in collaboration with 16 other universities. Importantly, it is one of the more extensive agricultural and horticultural databases in existence.

However, further inspection of the CIIDS images revealed a concerning limitation: the low quality images captured from the flip phones of 18 years ago did not contain any usable information. In addition, much of this imagery was not clean imagery with sufficient lighting, focus, and clarity of the main subject. Several times, in addition to the focus of the disease or pest issue among the crop in the picture, irrelevant elements such as a truck or farm vehicle would figure prominently in the frame, which would prevent cropping or additional manual processes and render the photo of no value. Ultimately, only one third of the 18 years' worth of data was deemed usable.

Farmwave additionally ran trials to collect data with Iowa State and Michigan State Universities. These custom trials proved more valuable, as image quality could be specified and controlled for addition to the Farmwave image database.

The Road to Data Collection

Digital Phenotyping: Early Disease and Pest Identification

Research on digital phenotyping in partnership with the Institute of National Research in Agriculture (INRA), in Dijon, France took our image recognition into the shoots and roots of plants. This effort is a continuing project with INRA as we work through multiple hardware challenges around capturing imagery in the field with the right tools to reproduce the lab environment. This level of image recognition, combined with digital phenotyping, adds value to the grower level because it can identify and diagnose a crop problem far earlier, meaning a less invasive and more economical spraying treatment can be applied.

Image quality is paramount. To achieve desired results, researchers must prescribe criteria for capturing photos of sufficient resolution, composition, and quantity.

We also identified the need for a data processing system that could keep pace with our data capture methods, as data collection in large volumes is a challenge to sort through manually. This is known as Farmwave's Cloud Optimized Recognition Engine ("CORE"), an internal system and process for the intake of large quantities of imagery. Our proprietary system ingests imagery at an astonishing rate for building models while simultaneously tagging and augmenting the imagery very specific ways. This methodology allows Farmwave to load data sets faster, build models in a more streamlined fashion, and produce results in a faster turn-around time than anything else on the market today.

The Road to Data Collection

The image below represents soybean header loss, showing our CORE processing, identifying, and displaying the kernels. Typically, many models require hundreds, if not thousands, or high resolution images to produce accurate results. In contrast, our sophisticated CORE achieved the results below with only 36 images. By Model 3, no human intervention was required; the CORE was teaching itself.


Original Photo


Model 1

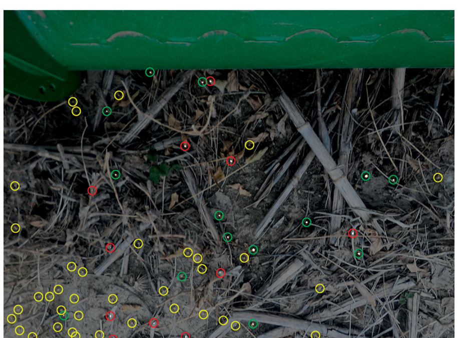
Model 2

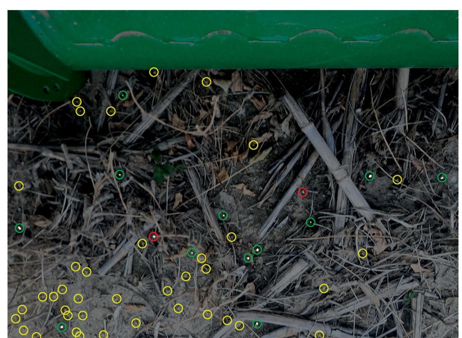
Model 3


Model 4


Model 5

The Road to Data Collection

Farmwave's Current Infrastructure

The Two Types of Data Required for AI in Agriculture

To operate as needed, both the data science behind image recognition and the raw data from third-party platforms are essential to address the daily decision-making that takes place in agriculture.

While many companies in the agriculture industry have closed, proprietary systems, our team saw the potential for maintaining true artificial intelligence in agriculture by encouraging interoperability through channels that could ingest third-party data as part of the overall ecosystem.

The Application Programming Interface (API) between these third-party sources and Farmwave is part of a longer-term solution to enable true automation in agriculture.

However, it is equally important that this external data meet the same stringent requirements as our imagery data in order to maintain the level of accuracy required by Farmwave or that should be reasonably expected by any other SaaS platform.

The Road to Data Collection

Results-Driven Data Science; Building AI Models for Commercialization

Within our offices, we often discuss what it means to build results-driven data science, that is, data science that goes beyond theory and research and development. Artificial Intelligence itself is not very new. What is new is the production or commercialization of AI technologies in real-world business environments that can be deployed to add value and produce results.

Since approximately 2013, the market introduction of tools like Google Glass or Microsoft's HoloLens brought image recognition to the forefront of both the consumer and industrial market's attention. These devices touted the ability to capture pictures and use voice recognition to move and transmit data in a whole new way: hands-free. This escalated quickly from 2013 to today, ushering in an entire cottage industry dedicated to consulting on AI and how it can be implemented into an organization's strategy.

Farmwave's early research and development efforts quickly revealed that in order to resolve challenging problems with real working solutions, the specified AI would need to not only address an organization's technical issues, but also align with its business strategy. Business models that provide value, scalability, revenue, and eventually profit need to be presented before budgets will be established. If these elements are not addressed, the project is closer to a research project with no mandate and an open-ended, indeterminate timeline, meaning for-profit companies have less of an incentive to seek an outcome.

With these objectives in mind, the Farmwave team was keenly aware of what assets, research, and technology we could bring to bear to address real-world challenges. While building our framework and platform, we constantly focused on results that could be measured, show value, and meshed with a viable business model. Our experience has proven that this relentless focus is the path forward not only for commercializing a product or SaaS model utilizing AI, but also for adding value and igniting competition in the marketplace.

Field Testing: The Proving Ground

After developing a proprietary framework with capabilities that could be tailored for related, yet specific applications, it was clear that there was an opportunity to help growers significantly increase the efficiency and profitability of their operations. However, this opportunity came with a heavy barrier to entry: solving the riddle of how to conduct real-time measurements of harvest loss, one of the most challenging problems facing agriculture today. To do this with clarity and efficiency, we recognized that the best way to get feedback, gather data, create a product, and test that product was to get out of the lab and into the field. This was the best channel to get an authentic taste of what our future customers experience day in and day out.

In 2019, we set out to understand and visually build out the ecosystem of harvest loss during the field portion of the harvest season. During our journeys, we had the pleasure of working with some of the greatest producers in the Midwestern United States. We were humbled and encouraged by the opportunity to spend time with these hardworking men and women who are truly what make farming great- in their communities, their states, and this country.

Across five different Midwestern states, we tested in both very flat land and extreme hills. We also experienced weather from extreme heat to snow. All of these environments were incredibly important to understand how our prototype hardware and software would respond when subjected to everything Mother Nature could muster. Our experiences highlighted the necessity for having a robust solution that would weather the spectrum of climates, and illustrated how the slightest changes could have the most profound impact on harvest loss.

Deployment

Before we embarked on our field work journey, there was one thing that all the growers we contacted made very clear: the deployment of our prototypes on their combines needed to be able to be completed in less than 30 minutes. Since the prototypes we built were self-contained units, we knew we could be fast and were focused on a deployment of four units in less than 10 minutes. What we found on day one of our trials was the opposite of our expectations: we were not as quick as we thought. The first deployment took a couple hours. Fortunately, we had arrived early, and heavy dew on the crop meant we could take the necessary time without slowing down the growers' harvest.

Our time improved: two weeks after our initial slow deployment, we had achieved our deployment target of 4 units in less than 10 minutes. The third week, we had deployment down to three cameras in seven minutes.

The next question we received from growers was, "How long does deployment take with one person?" Our real-world testing showed that three cameras could be deployed by one person in about 20 minutes.

Field Testing: The Proving Ground

Locations

Once we could consistently deploy within the growers' requirements, we set out to understand what growers needed to see as they operated their machines. This led to an exploration of what locations were feasible for positioning our camera units to obtain usable results.

What followed was a trial-and-error process of experimenting with different heights, locations, lighting conditions, and cameras to ensure the software was able to return high quality results, taking into account as many scenarios as a grower might experience in one season. In total, thirteen unique positions inside and out of a combine yielded the necessary data, including incredibly clear, useful, and highly relevant imagery. One of the most encouraging and impactful conclusions was that each of the locations we tried provided a unique look into how a combine processes, retains, and even loses grain.

Day	Location	Weather	Crop	Machinery
Day 1	Wheatfield, IN	Mild; 60 degrees each start of morning	Soybeans	John Deere 8510 Combine Harvester/ John Deere 925 26; Soybean Harvester Head
Day 2	Wanatah, IN	Damp ground with morning dew; beans dry enough by 1100 to start harvesting	Soybeans	2015 Case 7240 / MacDon FD75-S Head
Day 3	Wanatah, IN and Goodland, IN	56 degrees, Conditions: soft ground, high morning dew. Relatively dry and sun out will dry things out well	Soybeans (morning) Corn (afternoon)	2015 Case 7240 / MacDon FD75-S Head <---> 2016 Case 8240 / Geringhoff NorthStar
Day 4	Hamburg, IA and part of Nebraska	Sunny, mild, 65-70 degrees, terrain was very hilly with varied terraces	Corn, Enogen for Ethanol	John Deere 9670 STS with HillCo Tech / Head 608C
Day 5	Watson, MS	Sunny, colder, very windy, 58-65 degrees	Corn	John Deere S780 / Head 612C
Day 6	Hamburg, IA	Sunny and warmer, about 70 degrees	Soybeans	John Deere S660 with HillCo Tech/ Head John Deere 630FD Draper Head
Combined Trials	Creston, IL	Varying temperatures and weather conditions	Corn	Case 7120/Head Geringhoff Rota Elite XL

Lessons from the Field

Lessons from the Field

During the 2019 harvest seasons, we had the opportunity to experience just how our technology could help growers. Each day, a new insight. Each week, new refinements to our approach, technology, and assumptions. While we learned invaluable lessons from our time in the field, the main takeaway was this: we can identify, measure, and deliver loss information in near real-time with the extremely powerful technology we have developed. From a process perspective, we found that if we can show growers how much of their harvest is being lost, and where that loss is coming from, they can make the necessary adjustments to increase the grain retention in near real-time. We provided some of our initial raw imagery to a grower in Northern Illinois with tremendous results.

> **“With the imagery provided by Farmwave and an outside equipment consultant, we made two major changes to the combine.**
>
> **Adjusting sieve levels and changing the rotor settings together helped us to retain more grain.”**
>
> **Ben Hill**, Hill Farms

> **“At harvest there are two places for loss: threshing and front end equipment. Threshing loss sensors exist, but could use substantial improvement to give insight in various crop conditions.**
>
> **Front-end loss has no loss quantification sensors that I am aware of. A huge opportunity exists in this area.”**
>
> **Steve Pitstick**, Pitstick Farms

Lessons from the Field

The charts below illustrate the time-savings, quantity of counts conducted, and counts per square footage.

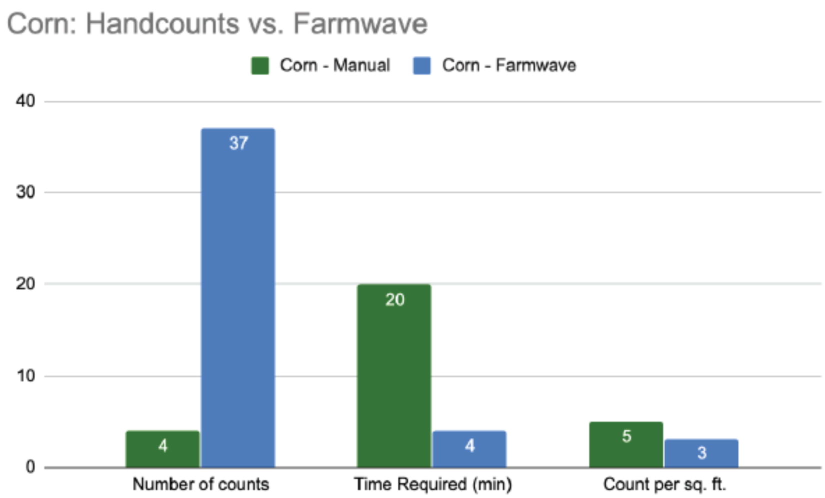

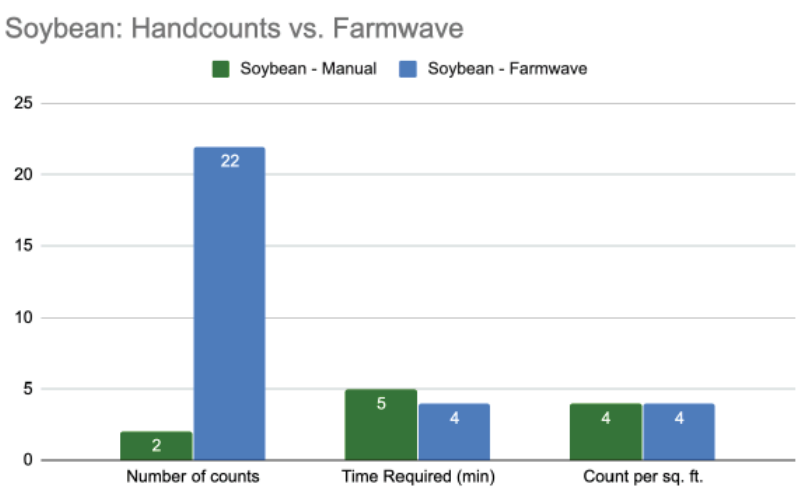

The ability to conduct multiple scans in the same amount of time as an individual grower could perform one hand-count is invaluable as both a time-saving measure and, as a result, a way to effectively gather more data on harvest equipment performance metrics, for overall harvest loss statistics and individual equipment calibration and performance.

Shifting the Paradigm

Harvest Loss Measurement and Benchmarking

Growers have been measuring harvest loss using the same manual techniques for many years, such as those outlined in the University of Georgia's Bulletin 973.[4] While these techniques are designed to help growers, the process is so time-consuming that they are very rarely utilized.

During our 2019 harvest trials, the growers we worked with expressed a willingness to receive, learn, and study their loss differently than they have in the past. They were eager to reap the benefits of tools that provided visual analysis of grain loss, which our trials proved was possible and could be performed in near real-time, providing growers with thousands of data points of loss in a single harvest season.

The main goal during harvest season for any grower- whether independent grower or a major Original Equipment Manufacturer (OEM)- is to get as much grain as possible out of the fields and into elevators to be sold at the best price possible. As an individual grower, the harvest loss calculation consists of an equation based on the crop, size of the head, and kernels counted before and after harvest has occurred. The test areas can range anywhere from one square foot to ten square feet depending on each grower's specified thoroughness and accuracy. Then, with all measurements taken into account, growers would average and obtain an estimated loss amount.

There are also other tools out there that are deployed by OEMs to collect all the post-harvest material that passes through a combine, re-sort the material, and remove any grain that passes through the harvester. Then the extra machine weighs that additional grain to understand what is being lost. These machines are roughly $1,500,000, putting them out of reach for the general producer. Over the years, combines have become much more efficient, especially with the introduction of automation to optimize the combine and retain grain. However, one grower said it best: "The machines are automating using a series of sensors and technologies that in many cases are 20 or more years old. How accurate can they be?"

Estimating yield loss and making the correct adjustments to a combine is a time-consuming process. During our trials, we experienced only two growers that stopped and checked their grain loss. Both of these farmers stopped only once in their respective fields. So, in the best case scenario a grower task someone with monitoring loss in every new field a combine enters. The other growers we worked and spoke with said they very rarely stopped to check loss due to time constraints. If they did check grain loss, it was done at the beginning and sometimes the middle of the season as the weather changed. Other than those two times, loss was not checked again- thus making the understanding of what they are losing highly subjective and highly variable.

Shifting the Paradigm

During the trial period to prove out our technology, we started with a processing time of 14 minutes in the office, which was lowered to 4-5 minutes for our field work, and has since been dramatically reduced to 4-6 seconds based on power availability. These are results from our prototype standalone unit; full integration will have the capabilities to run in real-time drawing power from a combine and allowing for potentially even faster machine automation. With processing speeds that fast, the technology is processing about 100 plus image counts per acre, giving the grower a far more complete depiction of what is being left in the field.

Based on the feedback we have received, if Farmwave simply provided an average count every 60 seconds based on the images collected during the previous 60 seconds, we would be pushing 12 data points per acre. This type of data would provide growers with directional analysis of real-time loss and provide a basis for whether or not to make adjustments.

Camera Units	Analysis Time	Photos per Acre
Pre-Prototype	13 minutes	1
Field Deployed	4 minutes, 30 seconds	3
Minimum Viable Product	4 seconds from direct power source (combine/tractor power)	180
	6 seconds (battery powered)	120

With the proper integrations, processing time would yield real-time analysis, with these data points integrated into the machines' understanding of the fields, crops, and unique topographies- all orchestrated to dictate how the machine responded. Over time, this data would build out a decision model for the machine- based on each field and all its variables, creating a machine that could not only make adjustments automatically but also anticipate changes and make the proper adjustments earlier, helping to retain more grain for the grower.

The chart below illustrates the different in kernel identification, recognition, and display for both corn and soybeans.

Actual vs. AI Comparison	Control Hand-Count	AI
Corn	3-4 kernels per sq ft (4 test plots)	4-6 kernels per 3 sq ft (12 AI analyzed photos)
Soybeans	4-5 beans per sq ft (4 test plots)	12-16 beans per 3 sq ft (7 AI analyzed photos)

The idea of artificial intelligence working on a piece of equipment is amazing to me. A piece of hardware able to watch my actual loss out of my combine will be a game-changer. Conditions change constantly, and each field is incredibly different.

A real-time, accurate assistant watching what I am doing is huge. I can make adjustments on the fly and I know right away if that was the correct thing to do.

The approach today is less than scientific. Run the combine, have a guy look at the loss, guess if it is acceptable, repeat. Farmwave would allow me to quantify loss, let me know as changes happen, and alert me.

In today's low-margin environment, every bushel counts. With Farmwave, I feel more confident knowing that I got every bushel out there.

Jake Smoker, G&D Smoker Farms

Proving the Value

The Value of Farmwave - Harvest

Technology for the sake of technology means nothing. If there is no impact in the technology being built, what value is it to anyone? Farmwave set out to not only prove its impact to farmers, but also demonstrate its value to the commodities market. In artificial intelligence, this impact is a compounding value not typically seen in traditional software-as-a-service (SaaS) companies. By identifying harvest loss in real-time and providing this information to farmers, they can make decisions that translate into real dollars. The Farmwave Harvest Vision System (FHVS) also effectively extends the value of a farm's operational capabilities at a fraction of its cost. Therefore, the value of technology like Farmwave on a global scale and in various crops is exponential- benefiting the farmer, commodities markets, and machinery companies through the power of automation integration.

Cumulative Value Addition by Farmwave - Corn and Soybeans (12 Year Period)

Harvest data obtained from the USDA and Farmwave field trials project an increase of 5 bushels per acre for corn and 3 bpa for soybean. The table and charts below represent the added value Farmwave technology would bring to the market with conservative estimates on decreased loss.

Historical Economic Harvest Data (U.S.)				
Corn				
Year	Actual Harvest	Optimized Harvest	Farmwave Value Add	Cumulative Value Add
2020	$51,923,000,000	$54,672,645,000	$2,749,645,000	$24,437,060,400
2019	$50,416,000,000	$51,886,160,000	$1,470,160,000	$21,687,415,400
2018	$51,696,000,000	$53,205,164,200	$1,509,164,200	$20,217,255,400
2017	$52,414,000,000	$53,915,768,800	$1,501,768,800	$18,708,091,200
2016	$52,397,000,000	$54,032,480,400	$1,635,480,400	$17,206,322,400
2015	$51,952,000,000	$53,454,711,600	$1,502,711,600	$15,570,842,000
2014	$54,812,000,000	$56,454,816,000	$1,642,816,000	$14,068,130,400
2013	$58,797,000,000	$60,765,049,800	$1,968,049,800	$12,425,314,400
2012	$81,216,000,000	$84,390,643,200	$3,174,643,200	$10,457,264,600
2011	$73,780,000,000	$76,069,560,000	$2,289,560,000	$7,282,621,400
2010	$76,916,000,000	$79,430,120,000	$2,514,120,000	$4,993,061,400
2009	$51,348,000,000	$53,826,941,400	$2,478,941,400	$2,478,941,400

Proving the Value

Historical Economic Harvest Data (U.S.)				
Soybean				
Year	Actual Harvest	Optimized Harvest	Farmwave Value Add	Cumulative Value Add
2020	$49,059,000,000	$52,210,123,560	$3,151,123,560	$30,988,832,460
2019	$31,755,200,000	$33,717,600,000	$1,962,400,000	$27,837,708,900
2018	$40,633,000,000	$43,051,827,000	$2,418,827,000	$25,875,308,900
2017	$43,636,600,000	$46,349,723,000	$2,713,123,000	$23,456,481,900
2016	$45,082,600,000	$47,663,814,200	$2,581,214,200	$20,743,358,900
2015	$35,968,200,000	$37,648,286,400	$1,680,086,400	$18,162,144,700
2014	$39,928,800,000	$43,059,096,000	$3,130,296,000	$16,482,058,300
2013	$43,987,300,000	$46,984,452,900	$2,997,152,900	$13,351,762,300
2012	$42,711,900,000	$46,001,993,400	$3,290,093,400	$10,354,609,400
2011	$34,761,600,000	$37,206,272,000	$2,444,672,000	$7,064,516,000
2010	$44,590,000,000	$46,802,600,000	$2,212,600,000	$4,619,844,000
2009	$35,044,800,000	$37,452,044,000	$2,407,244,000	$2,407,244,000

The Farmwave Vision System provides a significant amount of value in a single country in a single crop during a single season: harvest. This data does not even begin to contemplate the value that can be gained during the planting and spraying phases of the growing season as well. The value Farmwave represents to farming on a global scale exponentially increases when extrapolating these figures to multiple crops in multiple countries.

Grower ROI on Average 2,000 to 50,000 Acre Operations (1 Year)

The table below represents the added value Farmwave brings to growers who license the Farmwave Harvest Vision System.

Grower Annual Economic Impact - Harvest						
			Farm Operation Size (in acres)			
Crop	Price	Farmwave Value Add Per Acre	2,000	5,000	10,000	50,000
Corn	$5.48	$27.40	$54,800	$137,000	$274,000	$1,370,000
Soybean	$13.66	$40.98	$81,960	$204,900	$409,800	$2,049,000

(Numbers are based on the above averages of decreased loss of 5 bushels per acre in corn and 3 bushels per acre in soy at current market pricing from the commodities market).

Proving the Value

Over a single season, in less than one year, the grower reaped increased value to their operation and obtained increased profits. It is important to emphasize that the value goes beyond commodity pricing but to the end user as well. In countries such as Brazil, where it is common to have individual growers with operations with 600,000+ acres, the value returned is estimated to be in the multi-millions annually.

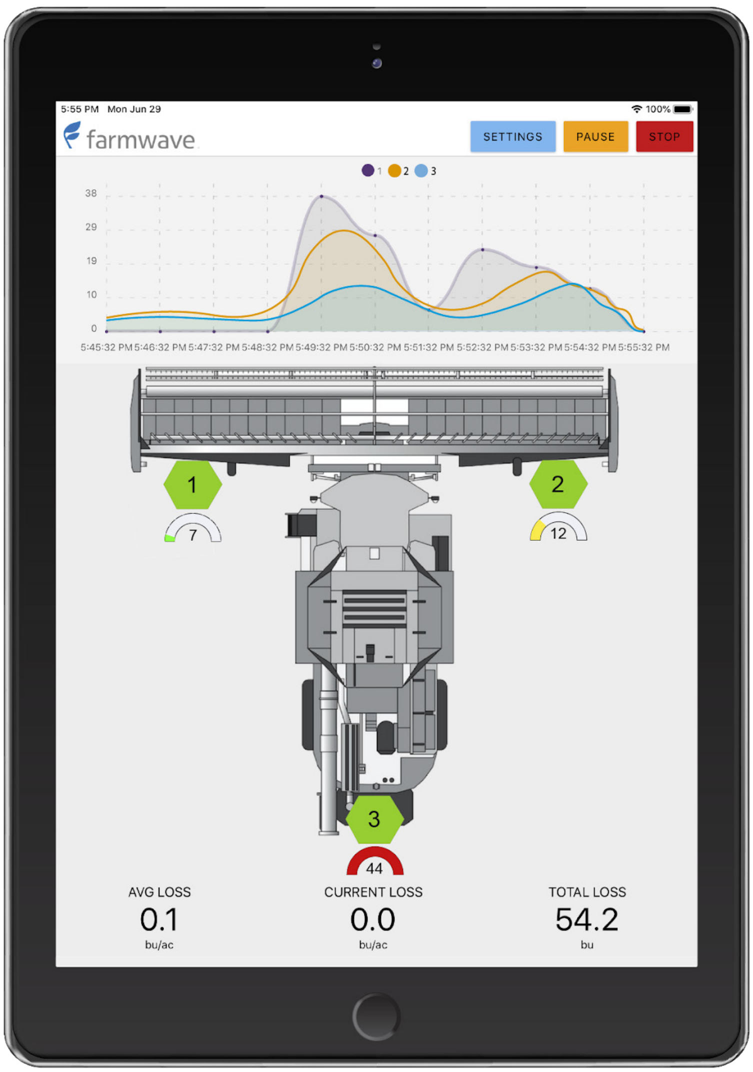

This screenshot from the Farmwave Harvest Vision System iOS app shows the advantage farmers have with real-time decision making.

As mentioned earlier in this whitepaper, the default hand count method is inadequate, inefficient, and inconsistent at best. It typically requires 30 to 40 minutes and is often not even performed by growers in the field during harvest season.

In contrast, Farmwave's Harvest Vision System can count 144 times per acre, dramatically improving the accuracy of measuring harvest loss.

Proving the Value

The Farmwave Harvest Vision System was designed to integrate with equipment produced by major machinery companies, but has also proven its value in field trials through externally mounting to existing machinery of various makes, model, and manufacturer.

This screenshot shows how Farmwave detects, analyzes, and displays loss trends and where loss is occuring on the machine in real-time.

Taken in aggregate from all farms equipped with this technology, it is easy to imagine the compounding value Farmwave brings to farmers, with savings increasing directly proportional to increasing acreage.

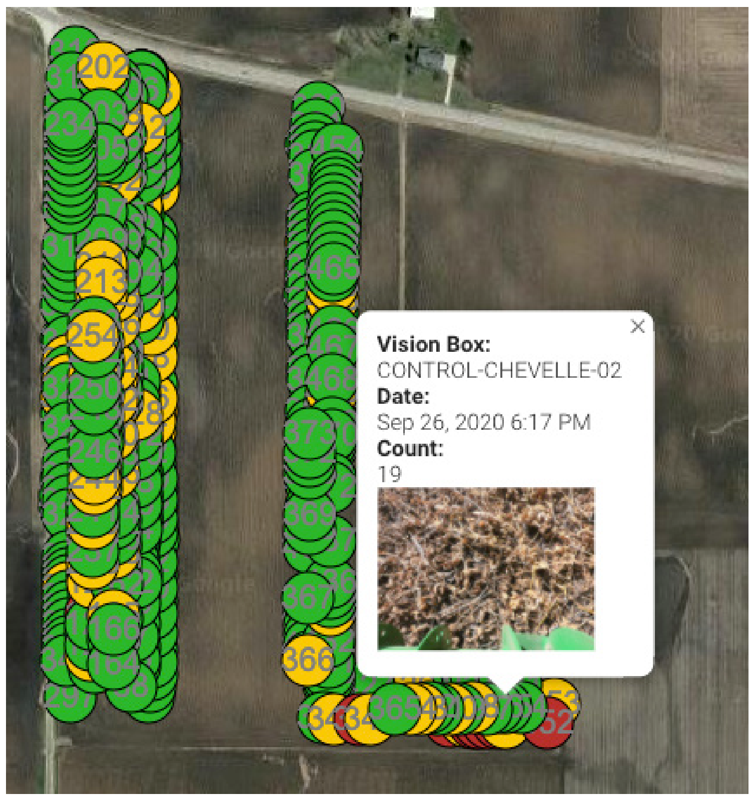

Proving the Value

The Value of Farmwave - Sprayers

But harvesters account for only one part of the growing season. What if we could repurpose the Farmwave handheld technology for outfitting sprayers? How many tasks could Farmwave perform in a single pass across the field? Feedback from growers is that they would see at least a 15% reduction in sprayer cost. Below, you can see the value this returns. Once again, it shows a compounding value add, with savings increasing directly proportional to increasing acreage.

Average Operational Cost and Value Savings - Sprayer				
	Farm Operation Size (in acres)			
	2,000	5,000	10,000	50,000
Annual Sprayer Investment	$54,706	$136,765	$273,529	$1,367,647
15% Savings	$8,206	$20,515	$41,029	$205,147

(Conservative estimates, assuming input costs of $25 per acre, a $200 per hour sprayer operational cost, and average rate coverage of 85 acres per hour).

Farmers are no strangers to multi-tasking, and their technology should perform the same way. In a single pass, the Farmwave Sprayer Vision System performs the following tasks:

- Nozzle performance,
- Sprayer plant count,
- Pest and disease diagnosis,
- Application coverage and droplet size analysis,
- Growth stage measurement, and
- Weed pressure measurement.

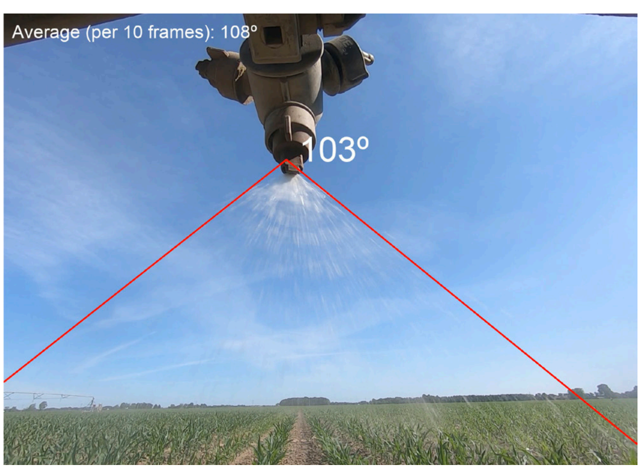

The photo above is a real-time display of sprayer metrics.

The photo above shows plant count in real-time.

Proving the Value

The Value of Farmwave - Web App

In addition to machinery, the handheld mobile web app of Farmwave continues to see traction and adoption across the globe. Multiple use cases, as seen below, among cooperatives in France for counting wheat seedlings proved to reduce human error by 24% while maintaining an 87% accuracy rate. This simple, yet impactful, technology continues to add value with small-holder growers in developing countries as well as empowering automated decision making in the field from the palm of your hand.

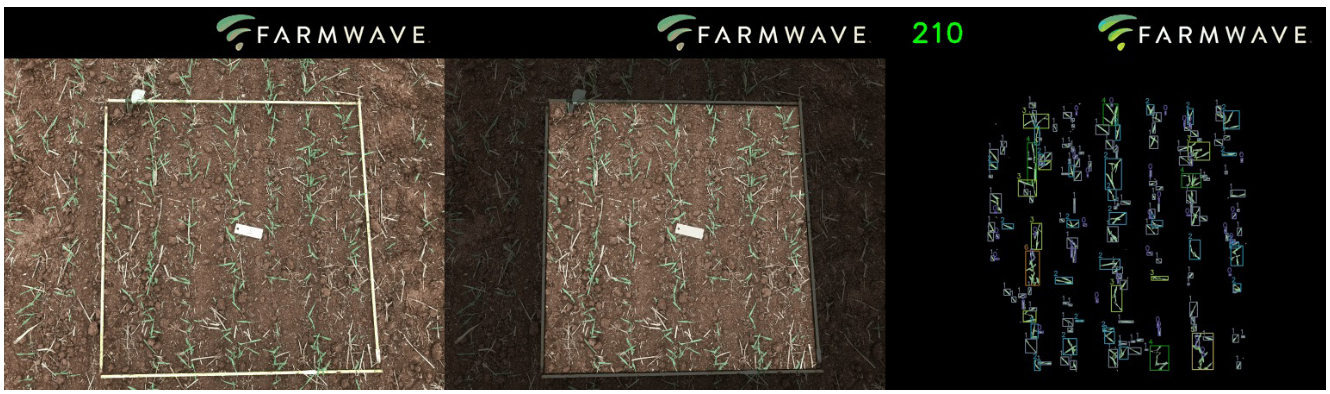

The photo above is a sequence of wheat plant counts taken with a handheld mobile device.

Our Path Forward

We have heard many times that what we are doing seems almost too good to be true. We attribute this skepticism to the many unsubstantiated claims made by many companies in the agtech space about what their technology is purportedly capable of. Farmwave set out to be different. Over the past months, Farmwave has demonstrated and proven to the market a number of things, many which were thought to be impossible, and some that others have only mocked up for marketing purposes.

The difference between what Farmwave has built and the rest of the industry? Farmwave actually works and produces real results, and adds real value that a producer can use as a baseline for real-time decision-making. Faster real-time decisions to retain more grain help to improve the profitability for independent growers.

Integration

We are moving towards building integrations with OEMs that would bring full-scale automation to agricultural machinery that does not require human intervention. Farmwave integrations into machinery would mean adjustments could be systematically made once specific thresholds were met. This level of automation preserves the future of trained combine operators, as opposed to drivers who often cannot fully appreciate the importance of operating a combine versus driving one.

Product Launch

We believe Farmwave's future will be an exciting one, starting with the commercialization of our app in June 2019.

Since that time, we have received tremendous feedback on the capabilities we have developed. The results from our real-world field tests are proving that our technology provides the answers growers need and want. Farmwave brought the Farmwave Harvest Loss Device to market in 2020, which provided incredible results to growers and proved the value it could deliver to markets across the globe.

During 2020, Farmwave recovered a loss value of $874,404.20 in only 10 weeks on 15 unique farms. We spanned almost 40,000 acres and collected just over 1.1 million images which have already begun to improve the core visions system of Farmwave's AI.

This device allowed farmers to view, analyze, and make decisions based on what they were losing in near real-time, assisting them in retaining more grain and improving their profitability. The demonstrated results of our technology caught the attention of dealers around the world and there continues to be increased interest from North America, South America, and Western European farmers who are looking for an aftermarket option to use Farmwave as soon as they can. While integration is the optimal desired outcome for all, farmers do not want to wait to take advantage of the value Farmwave brings.

Our Path Forward

Additional Equipment Research

In addition to field testing on combine harvesters in 2020, ongoing efforts are underway to evalute how Farmwave technology and hardware can synchronize with planters, sprayers, and tillage to solve additional problems via visual inspection.

Weed pressure, spray coverage percentages, disease and pest infestation in-motion, plant counts in-motion, growth stages in-motion, furrow seed orientation, furrow quality, and other anticipated capabilities will be tested in 2021. The spring planting season of 2021 will initiate the minimum viable product launch to early adopters of the Farmwave Sprayer Vision System. Pre-ordering has already begun.

But Farmwave won't stop there. The research listed below is already underway to continually improve our AI's capabilities and increase the value returned to farmers and markets for seasons to come.

Farmwave Research Efforts	
Additional Value Add Harvester Technology	**Additional Value Add Sprayer Technology**
Loss Monitoring (Preharvest, Header loss, Machine Loss)	Pest & Disease ID
Header automation Pest & Disease ID	Visual nozzle monitoring
Sieve area visual automation	Growth-stage/Crop height
Grain quality image recognition	Field overall plant health
Cleaning area automation(s)	Stand counts
Rear-of-machine coverage mapping	Dry fertilizer coverage during application
Residue Size/Shape	Percentage of application coverage
Residue Coverage	Weed detection/see-'n-spray
Tailings Analysis	Aftermarket Product (Sprayer Vision)
Machine Health/Diagnosis/Monitoring	
Aftermarket Product (Harvest Vision)	

Takeaways

This paper discussed how our team discovered, developed, and demonstrated real-world AI solutions that can help growers increase their yield by reducing harvest loss. Here are the main takeaways from our body of work:

1. **AI is a supplement, not a replacement for agriculture.** No AI can completely replace human experience, but can augment grower capabilities if developed correctly.
2. **One size does not fit all.** AI is an art and a science. The solutions to solve real-world problems cannot be found in a ready-made kit, but must be carefully cultivated with sufficient amounts of quality data and human experience.
3. **Customer discovery is key.** Getting out into the field to gather grower feedback is essential to understand the scope and nature of the problems to be solved- whether technical, cultural, or financial.
4. **It's time for a new way to understand and monitor harvest loss.** "Because its always been done this way" is not an acceptable answer. Growers demand more, and Farmwave's proprietary tools make it possible to understand real-time loss at the field level.
5. **The future is bright.** While AI for agriculture is in its infancy, the solutions we have developed are just the beginning. Machine integrations, additional capabilities, and growers eager for more efficient, purpose-built tools will drive and sustain this wave of innovation.

About Farmwave

Farmwave's mission is transforming the world's agricultural information into AI data models that power decision-making and preserve the future of farming- the legacy and experience of knowledgeable farmers. Our vision is to build the new decision support ecosystem for agriculture- transforming information from technology, people, and data into decisions to reduce crop destruction and to increase yields.

Farmwave set out to build a decentralized neural network of automated intelligence coming from various data points to aid in the speed, accuracy, and automation of these visual-based decisions.

Today, Farmwave is the only company to build a fully dedicated data set of imagery for various industry tasks to create a standardized set of agricultural data models. To fully solve for the complexities inherent to agriculture, we recognized that we needed capabilities that were lacking in readily available, open source solutions. Thus, we developed home-grown artificially intelligent models from our own algorithms and methodologies.

Endnotes

1. “Artificial Intelligence: Past, Present, and Future” Vox of Dartmouth, 24 July 2006.
2. Oxford University Press (OUP), Artificial Intelligence, https://www.lexico.com/definition/artificial_intelligence, Lexico.com, 2019.
3. “Overview,” ImageNet, Stanford Vision Lab, Stanford University, Princeton University. http://image-net.org/about-overview.
4. Paul E. Sumner, E. Jay Williams. “Measuring Field Losses from Grain Combines”, Bulletin 973, University of Georgia Cooperative Extension, May 2012.

Contributors

Chuck Adams, Charles Bassham, Frank Benoit, Brian Bull, Chris Chan, Craig Ganssle, Tom Hyatt, Steve Hyland, Noriaki Kono, Krishna Padmanabha, and Nick Palczynski.