




Harnessing Unstructured Big Data Through Computer Vision & Machine Learning

Presented by Hive.AI

August 2017



CONTENTS

Contents

- ABOUT HIVE AI3**
- INTRODUCTION4**
- COMPUTER VISION AS A SOLUTION5**
 - Sample projects & partners6
- HIVE’S ADVANTAGE6**
 - Case study 1 – Celebrity7
 - Comparison with Amazon Rekognition7
 - Business case for celebrity modeling – media10
 - Case study 2 – Biometrics.....11
 - Business case for biometric modeling – Melon13
 - Case study 3 – Violence14
 - Business case for violence – protecting communities.....14
- CONCLUSION15**

About Hive AI

Hive AI is a deep learning company focused on building practical applications leveraging visual intelligence. Founded in 2013, we have taken a full stack approach in building an independent artificial intelligence business understanding all facets of the machine learning workflow. Today, our products include:

1. Hive Work - A data labeling solution built for machine learning engineers by machine learning engineers. Hive work has 50,000+ workers ready to label data for any task, and we currently have capacity to label 10M items a day.
2. Hive Pre-Trained Models - An API as a service solution to solve your visual recognition problems. We have a variety of models including NSFW, Gender, Brand Detection, Celebrity Detection that we have built for clients.
3. Hive Enterprise Solutions - Our enterprise solutions team works with large and small companies on custom engagements where we build custom deep learning solutions to their problems. We have several clients in the media sector already.

Through our close partnerships with each of our customers, we have built productive symbiotic relationships that give them the insights to compete in their respective industries and grant us access to a steady stream of data that allow us to continuously improve our models. In conjunction with utilizing state-of-the-art machine learning techniques, this allows us to build machine learning models that are among the best in the industry. See thehive.ai for more information.

Introduction

We live in a world of exponentially growing data, with more data created in the last two years than in the rest of human history. In a digitized world, companies churn out trillions of bytes of data from their consumers as they go about their day – posting, sharing, and communicating. This presents unprecedented opportunity to derive a level of understanding about the customer unthinkable even a few years ago.

While traditionally this data would be analyzed through data analytics (namely multiple forms of regression and statistical analysis), it is now widely agreed that machine learning algorithms can better solve certain forms of problems. However, limitations still that have made applying machine learning to the vast oceans of data collected and stored difficult and prohibitively expensive.

Increasingly this captured data has taken unstructured forms, namely images and video. Rich media data has long been an intractable task by traditional machine learning approaches. In the past, utilizing this information required a person to act as a scribe between the visual data and the numeric data that a computer is able recognize. It is only then that traditional machine learning algorithms would be able to be applied. This solution does not provide the scale or the cost-effectiveness to keep up with the rapidly growing volumes of data and demand for new machine learning classifiers.

With little means to process or analyze this data, petabytes of image and videos sit in archives and data centers neither used nor useful. From security camera footage in stores to media broadcasts, organizations globally are not able to realize the potential of their vast data stores. And with the lost ability to harness this data, comes the loss of millions to billions of dollars of lost revenues and wasted analysis.

It is only with the extraordinary technological advances in computer vision in the past few years that trailblazers in the machine learning industry have begun to harness this data to provide valuable insights. Hive is one such trailblazer.

Through our position as a full stack machine learning company, we are able to iterate and improve our classifiers faster than our competition to bring impactful computer vision solutions to your company. We partner with leaders and disrupters in their respective industries, giving them the ability to thrive in the business of tomorrow.

Computer vision as a solution

Computer vision began in the 1970s with the goal of allowing machines to “see” and perceive as humans do. Even with all the technological advances over the past 40+ years, the primary goal remains unchanged.

Computer vision in the most general sense allows a computer to take visual data as an input and convert it into numeric information. For machine learning purposes, this creates the necessary structure from the visual data for machine learning algorithms to be able to form classifiers of the visual data.

After decades of research with few practical results, in recent years, rapid advances in technology have dramatically changed the field. Primarily due to advancements in chip manufacturing led by companies such as NVIDIA, the computational power has increased and cost has decreased to a point where successful commercial use-cases are possible in conjunction with advances in machine learning techniques. New technologies such as Deep Learning have emerged, drastically changing the possibilities of computers in the seeing and perceiving space.

Some of the more recognized applications of computer vision and machine learning are models that can describe an image utilizing free-form text or even to drive a car.

Exhibit 1: Samples of computer vision uses



However, this same technology can be applied at a smaller scale and still produce insights that have a significant impact on your company and change how you can compete in the rapidly evolving global economy. Hive already partnered with both entrenched leaders and innovative players across a variety of industries to employ this new technology in commercially sound contexts.

Sample projects & partners

Large Media Company	Hive partnered with a large media company to analyze a critical sporting event, providing a frame-by-frame analysis to allow them to optimize their coverage accordingly
	Yellow has used Hive's pre-trained NSFW classifier to prevent inappropriate user content from circulating
	Melon rose to one of the top apps on the app store by carefully curating users and immediately removing explicit users via the Hive NSFW pre-trained classifier
	Plaza utilized Hive both to monitor for contraband items through its NSFW filter and to better tag items to improve suggested items and reduce user friction

Hive's advantage

With the proliferation of computer vision tools and techniques it is now possible for a dedicated group of machine learning engineers to create vision classifiers relatively quickly given a clean set of labelled data. These engineers can then provide these APIs to other industries as a service. These APIs have become a type of commodity crucial for powering applications and devices. They are the new electricity of today.

Hive possesses a unique advantage in producing new models and ensuring these are ready to power your business forward. Our position as a full stack machine learning company enables us to understand both the opportunity and the challenges at each step of the machine learning workflow. We have products built specifically for data labelling, machine learning, and enterprise implementation all under one roof.

Our completely in-house development model allows us to greatly increase the speed at which we improve our existing machine learning classifiers by shortening iteration cycles. Through relationships with our partners, we are able to acquire the raw data behind new people, places, and ideas. We are then able to utilize our internal data labelling platform to immediately structure this data for machine learning use.


We have found great success in utilizing our improved product design model in the computer vision space. There, we have developed state-of-the-art models in violence, logo, biometric, celebrity recognition. See the below case studies for additional details.

Case study 1 – Celebrity

Celebrity face-recognition models aim to identify celebrities and other public persons in images and videos. While this goal has been in place for the past ten years, there have been major strides recently that have increased the accuracy of these classifiers.

Hive has implemented its own celebrity recognition system using a combination of facial recognition technology and deep neural networks. Our classifier identifies each recognized celebrity in the test image and gives a confidence score. See Exhibit 2 below for a sample of our celebrity model.

Exhibit 2: Sample of Hive’s celebrity model



Comparison with Amazon Rekognition

To better understand the capabilities of our celebrity classifier, we pitted our classifier against Amazon’s facial recognition technology. The test consisted a random set of 500 celebrity and politician photos – to ensure that we covered a wide range of ages. The test sets were selected to be difficult where all the images would be new to the classifier to ensure fair analysis.

Exhibit 3 highlights an example between the two classifiers on celebrities, while Exhibit 4 shows a comparison between an image of politicians

Exhibit 3: Hive vs. Amazon - comparison of celebrity test results





Exhibit 4: Hive vs. Amazon - comparison of politician test results





Both Exhibit 3 and Exhibit 4 illustrate our classifier’s increased precision and our ability to detect more difficult faces.

In Exhibit 3, both models correctly recognize Kayne West with approximately the same level of confidence, however the Amazon Facial Classifier fails to recognize his wife Kim Kardashian. Both faces are difficult for most facial recognition detection classifiers due to the angled shots. We have tuned our models to specifically recognize faces from the all angles, including difficult side shots.

In Exhibit 4, both models correctly recognize Chuck Schumer, however Hive’s model correctly labels Bernie Sanders without the false positive on Ernesto Perna. Again, the image of Bernie Sanders is difficult due to the low resolution and blurriness of the face. Our model has also been tuned against lower quality images to better serve live video chatting platforms, which frequently require classification of low resolution images.

Overall, our classifier was able to outperform Amazon across the test set of 500 celebrities and politicians at similar recall levels. Exhibit 5 below highlights the precision of Hive on this test set.

Exhibit 5: Comparison of results Hive AI vs. Amazon Rekognition

	Hive	Amazon
Precision	90.9%	86.4%
Recall	70.6%	70.1%

As can be seen by the example and the test results, in a short amount of time we have managed to achieve a comparable, if not higher, performance compared to one of the largest established players in the industry. This is due to highly tuned deep learning

architectures and a range of engineering challenges we've overcome, which effectively filter noise out of large amounts of data and quickly index millions of images. As we continue to develop and improve these models with an ensemble of advanced algorithms, we expect our performance to continue to improve.

Already, we have built one of the largest databases of celebrities for machine learning classification in the world using our distributed data labelling platform. We believe that it is this focused approach on the practical applications of computer vision that gives us the necessary edge to propel our future success.

Business case for celebrity modeling – media

Celebrity tagging is an especially useful insight for traditional media and advertising campaigns. Currently, media companies cannot confirm exactly what is airing at any given point in time. This creates difficulties in valuing the effectiveness of a given advertising campaign or sponsorship.

Using its celebrity classifier, Hive has created an interface that specifically tackles this problem in Hive Media. Hive media allows our partners to identify which celebrity is on the screen at any given moment on standard network shows. See Exhibit 6 for a screenshot of the Hive Media product in action.

Exhibit 6: Hive Media production example of celebrity tagging – Donald Trump

Query

7/24/2017 - 7/31/2017



CURRENT CLIP	KGO-HD: Donald Trump 3:45:08 PM - 3:45:09 PM 1 seconds
---------------------	--


By juxtaposing this data with viewership numbers, media companies will be able to derive precise, data-backed insights on impact of celebrity sponsorship. Companies will also be able to understand public perception of celebrities in a quantifiable manner that was impossible prior to the application of computer vision.

Case study 2 – Biometrics

Biometric machine learning models aim to provide a range of user information including gender, ethnicity, and age. Here at Hive, we have implemented the ability to recognize all three of these biometric categories using advanced facial detection algorithms.

Of all the biometric machine learning classifiers, age classification remains one the most popular and most difficult problems to solve. Many companies both large and small have been creating age classifiers, such as Microsoft's viral HowOldRobot. We have created our own version of this age classifier that captures the user's age within 5 years. See Exhibit 7 for a sample of this classifier at work.

Exhibit 7: Sample of Hive's age classifier



In addition to age, Hive has created classifiers to distinguish user gender and ethnicity through visual data. Our classifier can distinguish between persons of African, Arabic, Asian, Caucasian, Hispanic, and Indian backgrounds. See Exhibit 8 for samples of real images that have gone through our gender classifier and Exhibit 9 for an example of the ethnicity classifier.

Exhibit 8: Sample of Hive's gender classifier





Exhibit 9: Sample of Hive's ethnicity classifier



We have reached >95% accuracy on our gender classifiers and >90% accuracy on our ethnicity classifiers. These results are comparable with similar classifiers offered by our major competitors. See Exhibit 10 for more details on the accuracy of our ethnicity classifier.

Exhibit 10: Hive’s ethnicity classification confusion matrix



To date, our recognition models have trained on millions of images have and through our relationship with some of our customers, we continue to ingest thousands of images every day. This allows us a steady flow of data to better train our classifiers and in turn improve our customer’s experience.

Business case for biometric modeling – Melon

Companies can utilize this information to better understand their users and derive valuable insights on their business. We partnered with Melon for them to derive those insights.

We have implemented our biometric classifier technology into Melon’s video app to allow them to better enforce their age-requirement and to give them additional data on the hundreds of thousands of video calls they have on their platform every day. They then use this data to identify which types of calls last longer, whether users are being truthful about their profiles, and optimize their matching algorithm to improve user experience. Working with Melon, we fine-tuned our biometric classifiers for the lower-

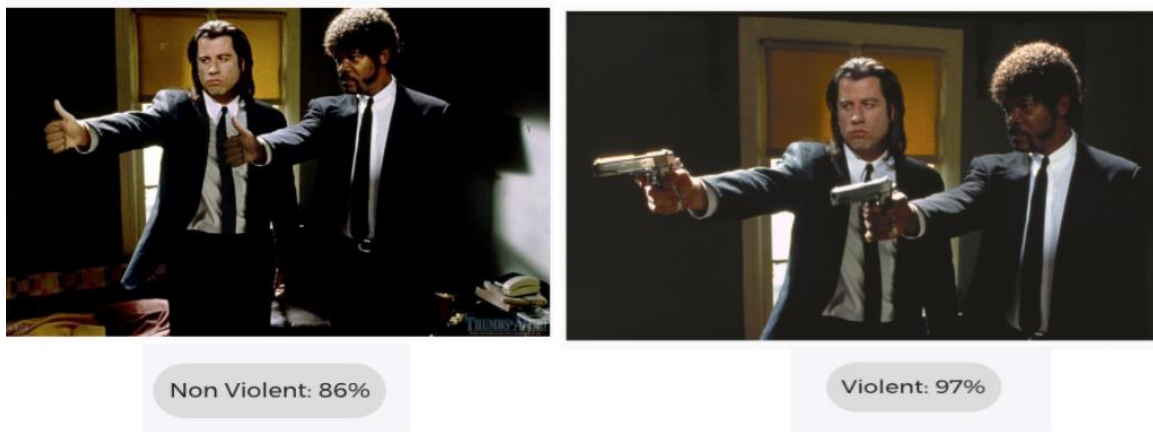
quality videos common on chat apps. From this partnership, we became one of the leading machine learning companies at extracting biometric information from low-quality images.

Case study 3 – Violence

Violence models and their counterpart in not-safe-for-work (NSFW) models are primarily used as screening methods for marketplaces and other platforms to prevent contraband from entering these platforms. However, these violence models can also be used to in traditional media to better understand the effect of violent scenes on viewership numbers and ratings.

Hive has developed a violence model that can detect various classes of violence including guns, blood, explosions and fights. This classifier is again based on deep neural network technology. Our training set primarily leverages images of contraband found on marketplace apps and various scenes from movies. See Exhibit 11 below for an example output from our violence model.

Exhibit 11: Sample of Hive’s violence classifier



Business case for violence – protecting communities

Violence machine learning models are critical to creating safe communities and increasing the overall user experience. You don’t need to go back far to find examples of graphic events occurring in safe communities. Instances of such events occurring on Facebook Live have become more common and more extreme. Traditional means of reporting simply cannot keep up with the pace and scale at which videos are uploaded.

Hive has already implemented protections for many of our community based partners including Plaza (a new take on marketplaces), Melon (a live video chatting app), and Chatous (a live chat app). By running the user created images and video streams through

our pre-trained violence model, our partners can regulate exactly who and what are allowed take part in these communities.

Conclusion

As we move away from text towards an increasingly image and video dominated world, new methodologies are necessary to derive insights from this new treasure trove of consumer data. One such technique that can handles this increased throughput is the combination of computer vision and machine learning.

While advances technology and chip-design have allow many players to enter this space, very few possess the correct mix of understanding and practicality that allow for the creation of commercially viable products necessary to prepare your company for the challenges of the next 20 years. Hive is one company that does.

We have demonstrated that our position as a full-stack machine learning company gives us a unique advantage in developing innovative and cost-efficient classifiers. Our performance rivals that the largest players, but we possess a sense of focus that gives us the edge.

To prepare your company for the increasingly competitive and rapidly evolving economic climate, incorporation of machine learning into your companies needs to be a top priority. Hive is here to help you accelerate this AI-led transformation and either become or remain an industry leader for the coming decades.