aws machine learning

GOING LONG ON MACHINE LEARNING

How AWS and the NFL teamed up to transform a 100-year-old league



03 **Introduction**
Letter from Michelle McKenna, SVP and CIO of the NFL

04 **What are the odds?**
Opening story

05 **What happened?**
The data behind the stat

06 **The factors**
How is Completion Probability calculated?

08 **The mechanics**
How does Next Gen Stats calculate Completion Probability?

09 **The formula**
How the machine learning models are trained

11 **The journey**
The NFL's machine learning journey

13 **The problem**
Why machine learning?

14 **The end results**
Measuring the impact

15 **Ready for more?**
Taking machine learning beyond the end zones

16 **Get in the game**
Resources

17 **Glossary**
More stats

If you grew up in a football family like I did,

you might already know that the NFL is America's largest sports organization, with over 188 million fans worldwide. We're a big content creator—NFL games accounted for 38 of last year's top 50 telecasts. We're also a big data creator—every week our league generates 3terabytes, equivalent to 1,500 hours, of data.

As CIO of the NFL, I'm responsible for ensuring we leverage our data to create the best and most efficient technology solutions that will evolve our game, engage our fans, and protect and develop our players.

> ***Transforming a 100-year-old league is not an easy job. Likewise, the NFL's machine learning journey has not been a straightforward path.***

Next Gen Stats, one of our machine learning projects—is the product of our working partnership with Amazon Web Services. Working with AWS just made sense because of their flexibility, security, and ability to scale. AWS has the broadest offering of cloud services for our builders to build upon—giving us the ideal platform to grow.

For me, getting started with machine learning was not a question of "why" but of "how." I knew machine learning would transform our company. I also knew our technology projects needed to be supported at the top level of the organization to get both the business and technical teams working together and sharing the same priorities.

Implementing machine learning benefits the entire company, not just the technology department. Machine learning has made the NFL's production teams more efficient—transforming previously tedious roles like video labeling into an automated, streamlined process. Coaches can use the technology to officiate playbook formations and automatically draw out plays, saving them time on the sidelines. With the power of machine learning on AWS, we can better understand fan engagement, how a game is presented, the potential impact of adjusting the rules of play, how the game is called, and player performance and safety. Next Gen Stats allows us to use this real-time data to engage, inform, and empower fans in new and unique ways.

> ***We're excited to share this portion of our journey with you and help you see a little bit of what we've been up to at the NFL.***

We're working to build a better playing, coaching, and viewing experience. Thanks to AWS machine learning, we're revolutionizing a 100-year-old league.





Michelle McKenna
Senior Vice President and CIO, NFL



It was Sunday night, week one in the 2018 season, fourth quarter.

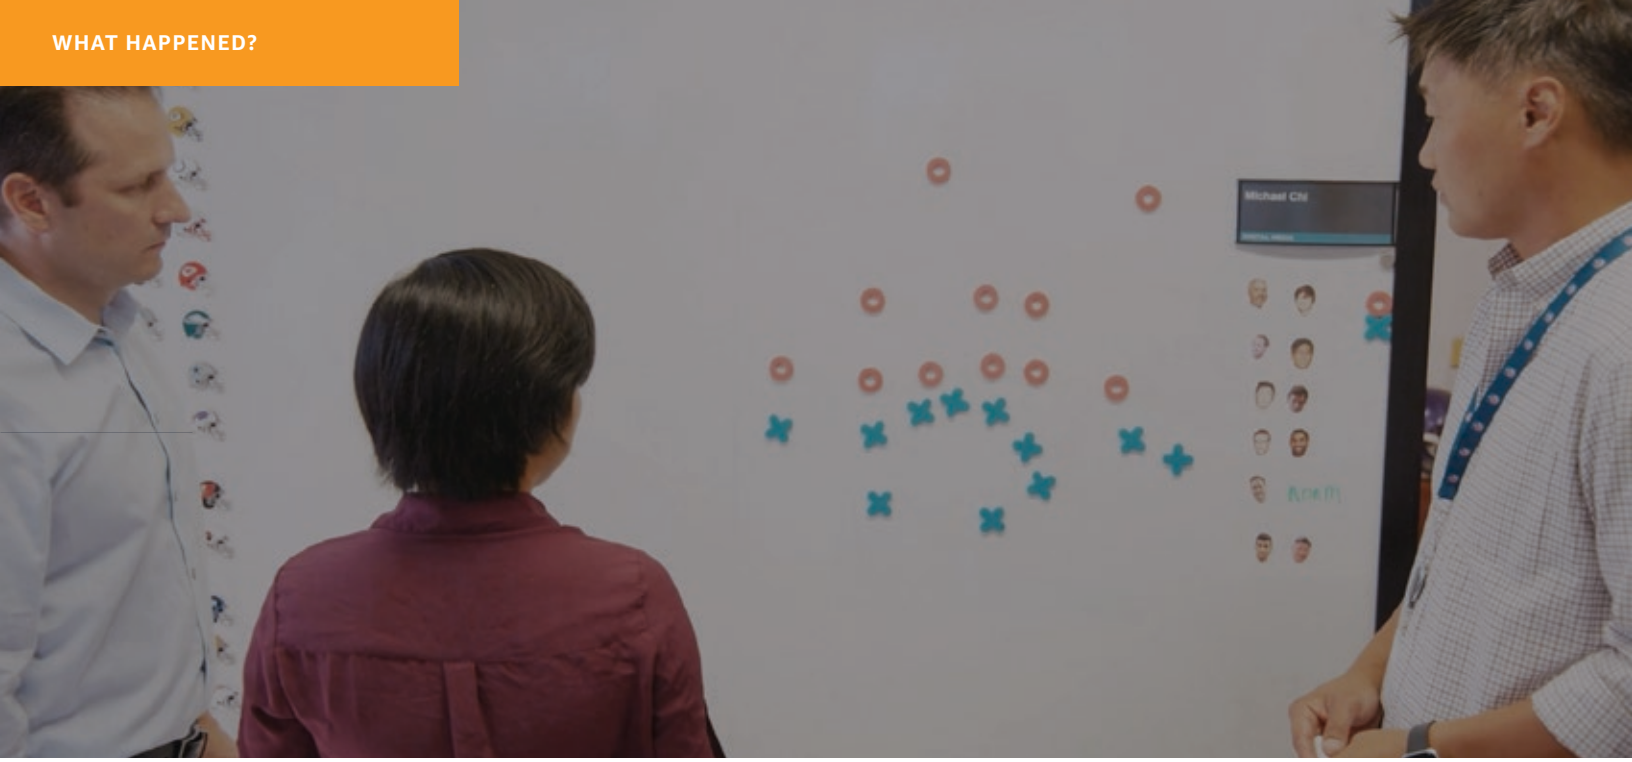
The Green Bay Packers were down 3 to 20 against longtime rival the Chicago Bears. Aaron Rodgers, who sat out much of the first half from a knee injury, was back in the game, but things weren't looking good. Second and 2, and Corey Linsley snaps to Rodgers. Man-to-man coverage. All eyes are on Rodgers, who appears to have time, except nobody down the field looks open. Rodgers steps back and launches.

Not all passes are created equal. When players defy the odds, we are exposed to how talented they truly are. But this often doesn't get represented by traditional box score stats, which would score Rodgers the same whether his pass traveled three yards behind the line of scrimmage to an open running back or whether his pass did what happened next.

The ball sails 39 yards down the field toward the back-right corner of the Bears' end zone, and the Bears' Kyle Fuller is all over Geronimo Allison, the target of the last several of Rodgers' first-down throws. Everyone can see it's an unlikely catch as it arcs toward the pylon. But how unlikely?

Next Gen Stats, powered by machine learning models built on Amazon SageMaker from AWS, had just launched another new metric for the 2018 season called Completion Probability, which leverages tracking data to improve upon the limitations of raw box score stats and add context to each passing play. Next Gen Stats calculated the pass had just a 14.7% Completion Probability—the most improbable completion that week.

Allison leaps with Fuller on his back, in full reach, fingers wide to swat the ball, and misses by what seems like inches. The ball lands right in the pocket as Allison cradles it close and plants two feet with full control before sliding out of bounds for a touchdown. This is the beginning of the end for the Bears, who ultimately lose the game 23-24 to a significant fourth-quarter comeback for the Packers.



The data behind the stat

By Matt Swensson, Vice President of Emerging Products and Technology of the NFL

Completion Probability is measured using more than 10 different in-play factors starting with data transmitted by RFID chips in the football and on players' shoulder pads all collected by RF receivers around the stadium.

In the case of Rodgers, the data shows the pass traveling 60.3 yards in the air from the location of Rodgers at the time of the throw to Allison at the time of the catch. Rodgers had 2.1 yards of separation from Jonathan Bullard when he released the ball, and Allison had 0.9 yards of separation from Kyle Fuller at the moment of the catch.

All of those factors, among several others, had a direct relationship with the likelihood Rodgers' pass would be complete or incomplete. We can evaluate these relationships by plotting each in-play factor against the actual completion percentage to better understand each factor's effect on the outcome of a play and contextualize the difficulty of a throw.

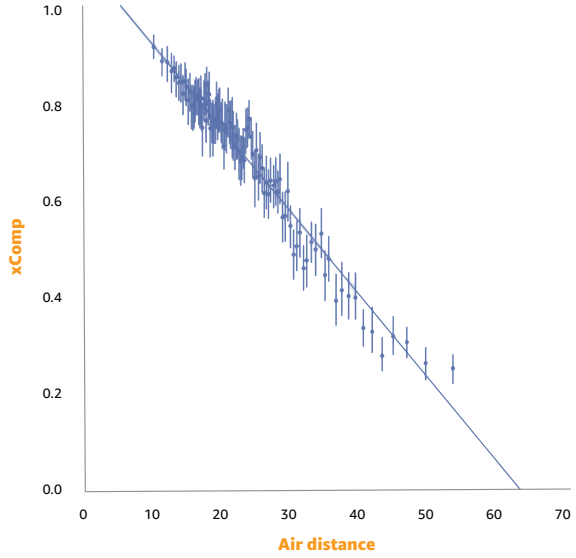
Let's review some of these factors and examine how the predictive models were trained.

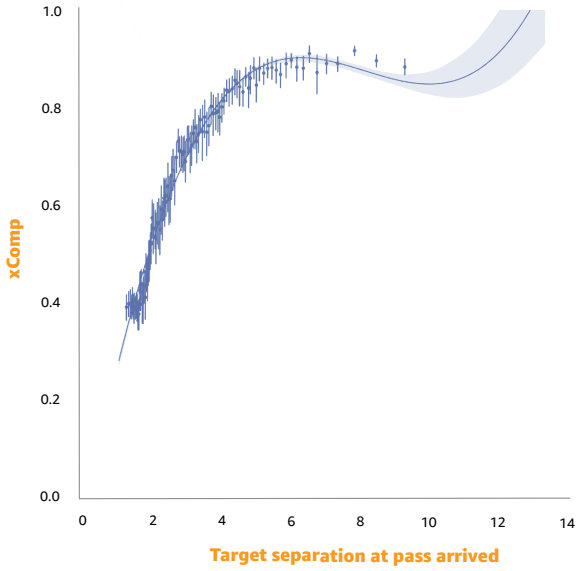
Completion Probability's top factors

1

Air Distance

The further the ball has to travel, the lower the likelihood of completion. This is measured by the air distance – the true distance from the location the ball is thrown to where it is caught. Passes traveling more than 40 air distance yards have approximately 20% chance of completion.





2

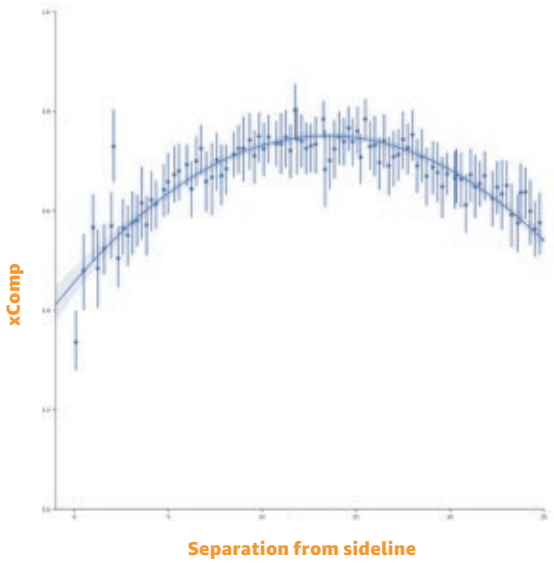
Target Separation

As the distance between the receiver and nearest defender increases, the likelihood of a completion also increases. The larger circles at lower target separation show that it's more common for receivers to have close defenders.

3

Sideline Separation

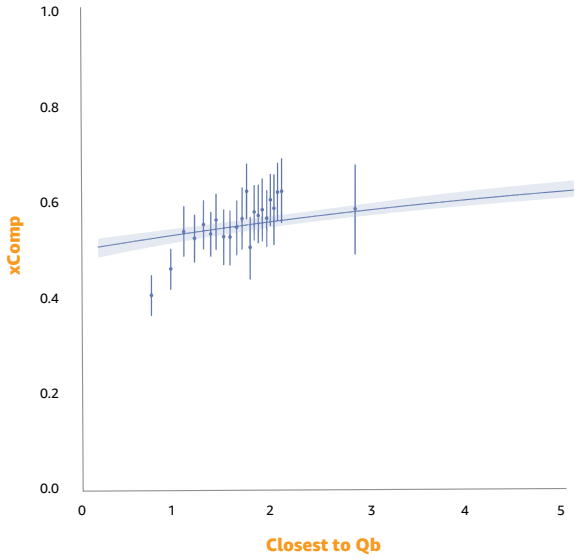
As the distance between the receiver and the sideline decreases, the likelihood of a completion also decreases. The probability of a completed pass decreases rapidly at five yards of sideline separation. Controlling for all other factors, passes to the sideline just inside the white paint have a roughly 30% chance of a completion. After about 10 yards, we see diminishing returns.

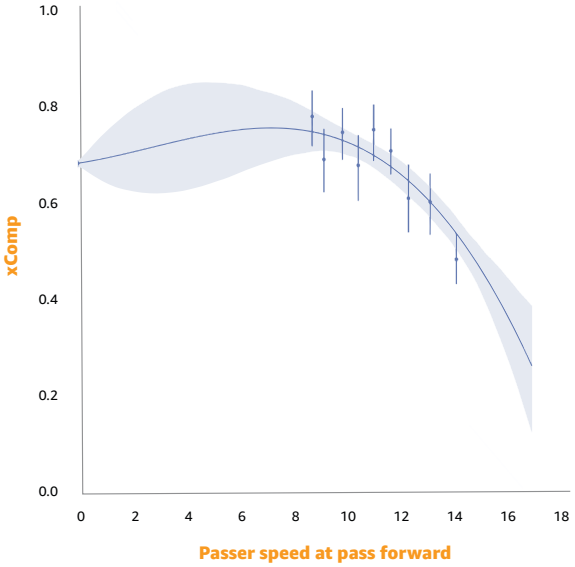


4

Pass Rush Separation

As the distance between the quarterback and nearest pass rusher at the time of the throw decreases, the likelihood of a completion also decreases. A quarterback throwing with no defenders around has a higher probability of a completed pass compared to a quarterback with a pass rusher within a few yards at the time of the throw.





5

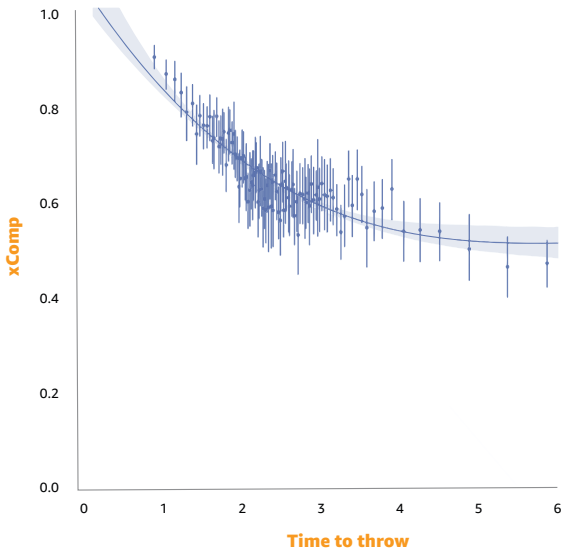
Passer Speed

As the speed of the quarterback at the time of the throw increases, the likelihood of a completed pass decreases. Speeds below 8 MPH have little effect on the probability of a completion. However, as the speed of the quarterback increases above 8 MPH, the chance of a completion decreases.

6

Time to Throw

Most passes occur between 2 and 3 seconds after the snap. As the duration of time increases from snap to throw, the likelihood of a completed pass decreases. The probability of a completion declines significantly after 3 seconds.



These are just a few of the data points measured and fed into machine learning model to develop the Next Gen Stats Completion Probability metric. Next we'll explore why the NFL decided to use machine learning.

How does Next Gen Stats calculate Completion Probability?

By Jarvis Lee, AWS Data Scientist and Tyler Mullenbach, AWS Practice Manager

By leveraging AWS' broad range of cloud-based machine learning capabilities, the NFL is taking its game-day stats to the next level so that fans, broadcasters, coaches, and teams can benefit from deeper insights.

In partnership with AWS's Professional Services, the Next Gen Stats team created and deployed machine learning models using Amazon SageMaker. These models leverage inputs from traditional statistics, player tracking data, and player profiles, and these models are continuously refined by the Next Gen Stats team to reflect the trends and outcomes during any given season.

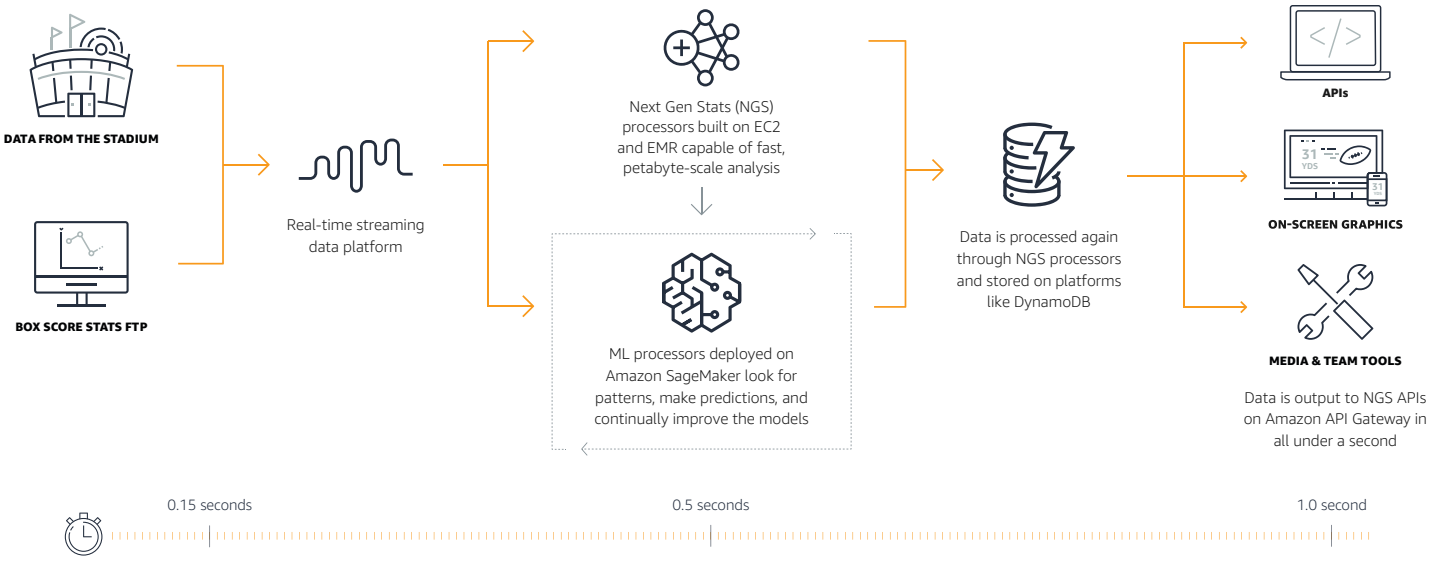
Once deployed, the machine learning models provide real-time inference during games. To obtain real-time insights, data from the stadiums is streamed to AWS. Once in AWS, the data undergoes over 100 processes in under a second.

The results: The NFL is able to produce APIs, create on-screen graphics, give sports announcers more unique data points to talk about, and expand their media and team tools.

Amazon SageMaker

Building and training machine learning models used to be locked in the ivory towers of elite developers and data scientists. But thanks to Amazon SageMaker, developers can build, train, and deploy machine learning models quickly, no PhD required. The fully managed service covers the entire machine learning workflow to label and prepare your data, choose an algorithm, train the model, tune and optimize it for deployment, make predictions, and take action.

Get started with Amazon SageMaker >>

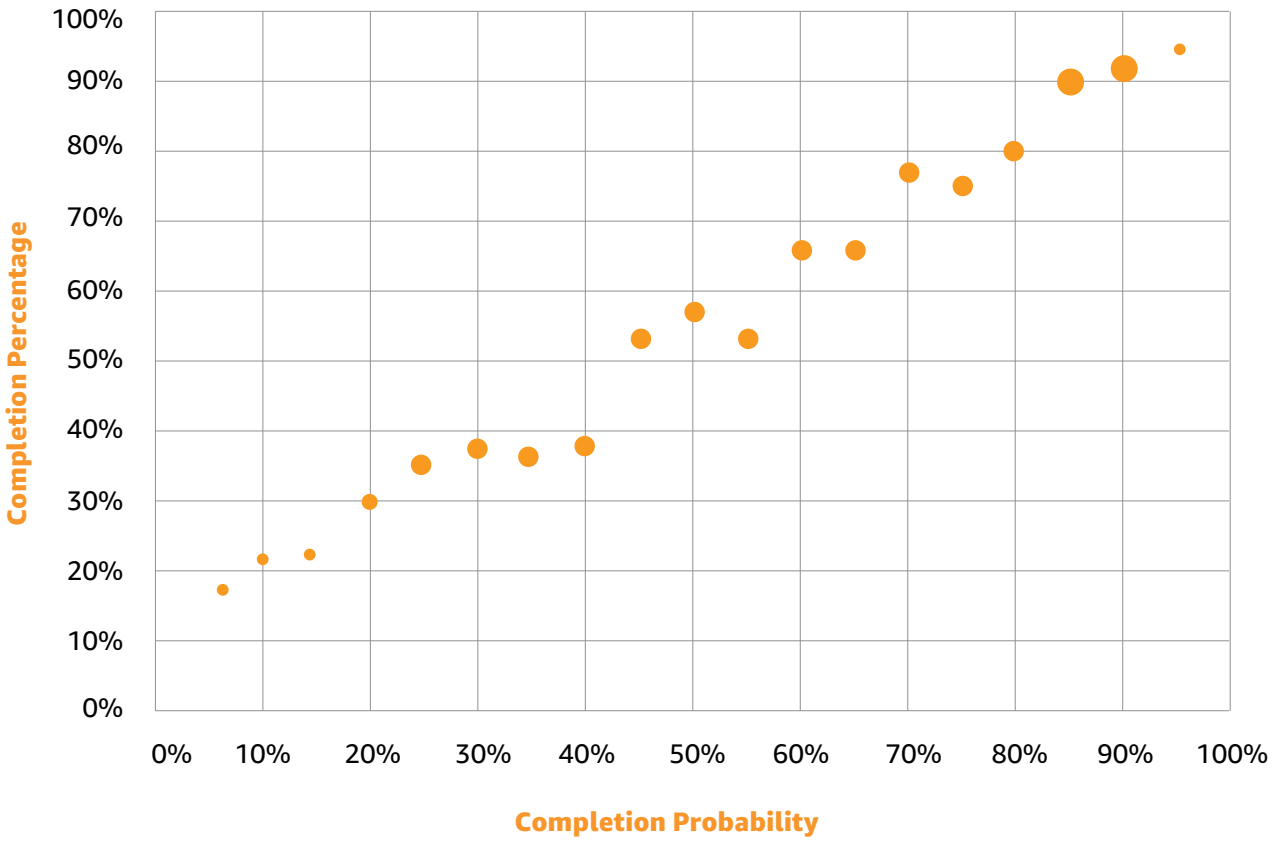


How the machine learning models are trained

Next Gen Stats uses a flexible machine learning algorithm that frames Completion Probability as a linear regression problem. The algorithm outputs a probability of whether the given inputs from any given play would result in a completed pass or not.

The team was able to validate a best-performing model, using SageMaker through a 70-20-10 train/validate/test methodology. This means that 70% of the data is used to train the model, 20% to validate the model based on that training data, and 10% to confirm the model is generalizing as expected. If all three buckets have approximately the same result, then we know we have hit the sweet spot of producing consistently accurate predictions.

Completion Probability vs. Completion Percentage





Q&A with Matt Swensson

How was the first Completion Probability metric built?

The first concept of Completion Probability was built in a spreadsheet. I think you need to truly have a good understanding of what you're trying to solve for before you just automate it.

Why was machine learning the solution for your data collection?

We have so much data—the question was how could we process it all and how could we process it as fast and accurately as possible? Machine learning allows Next Gen Stats to analyze data thousands of times faster than a human while meeting human accuracy benchmarks. It's truly a game-changer.

How do you know you can trust the machine learning models you built?

With machine learning, there's never 100% accuracy, that's why you need to understand your criteria for a model and the level of accuracy you're comfortable with. There's this perception that humans are always the gold standard, but you have to take a step back and consider that computers are better at some tasks.

How do you continue to fine-tune and enhance your models?

Given machine learning's need for iteration, it's important the teams working on the Next Gen Stats models are always communicating, trying new things, and open to new discoveries. Machine learning is not something you do once and then you're done. We have metrics we haven't solved yet, and that's okay. You have to go in with the expectation that it might not be the right tool, or it might not work. You have to trust the process and know the wins will come.

> ***You have to trust the process and know the wins will come.***

The NFL's machine learning journey

Fans only see what's on the field. But there are thousands of split-second decisions made by the coaches and players during a game. In a post-Moneyball world, traditional commentary isn't enough—you need data.

Having seen success with the launch of NFL.com along with club sites and Fantasy Football, Matt Swensson's team was ready for the next big thing. When, in 2012, the NFL wanted to pull data right on the field,Matt jumped on the opportunity and the team rolled up their sleeves.

In a post-Moneyball world, traditional commentary isn't enough — you need data.

In 2009, the NFL first experimented with optical solutions, some of which employed computer vision models. But those solutions did not provide effective results. Computer vision technologies are limited when it comes to separating and recognizing multiple players in scenarios with occlusion – optical tracking wasn't reliable enough to be the primary source. Radar-based solutions also proved limited in what they could track.

The idea of RFID tags had caught the NFL's eye as the technology was becoming usable at smaller sizes.

Not every idea was a winner.

Some of the tags were as large as a cell phone, with batteries that would need to be recharged every game. Others had unreliable reception.

Finally, after a number of iterations, the NFL teamed up with Zebra Technologies, which offered a tag the size of a bottle cap. The tag since shrunk down to the size of a stack of dimes.

The team then needed a way to store the data, and a cloud platform that could scale, and they needed a way to process the vast amount of data they were collecting on a game-by-game basis.

AWS proved to be the best solution. The NFL started with calculating the most basic derived metrics—distances between players, how players were lining up in formation and grew to more complex metrics such as pass rush detection.



> ***It was and continues to be an iterative, side by side partnership.***
>
> ***- Matt Swensson***

"As we continued down this path of stat creation we wanted to do more," said Swensson. "It got us thinking: 'how can we take our stats to the next level?'"

The chicken and egg problem

Where do you start if football analysts don't understand what's possible and your engineering teams don't understand what's useful for those analysts?

You listen to a lot of commentary.

Swensson's team devoured content from greats such as Aikmen, Gruden, and Collinsworth.

"We listen to what our end users and analysts are asking or telling us," said Swensson. "I think the technology has been so new everyone is learning what's available to them."

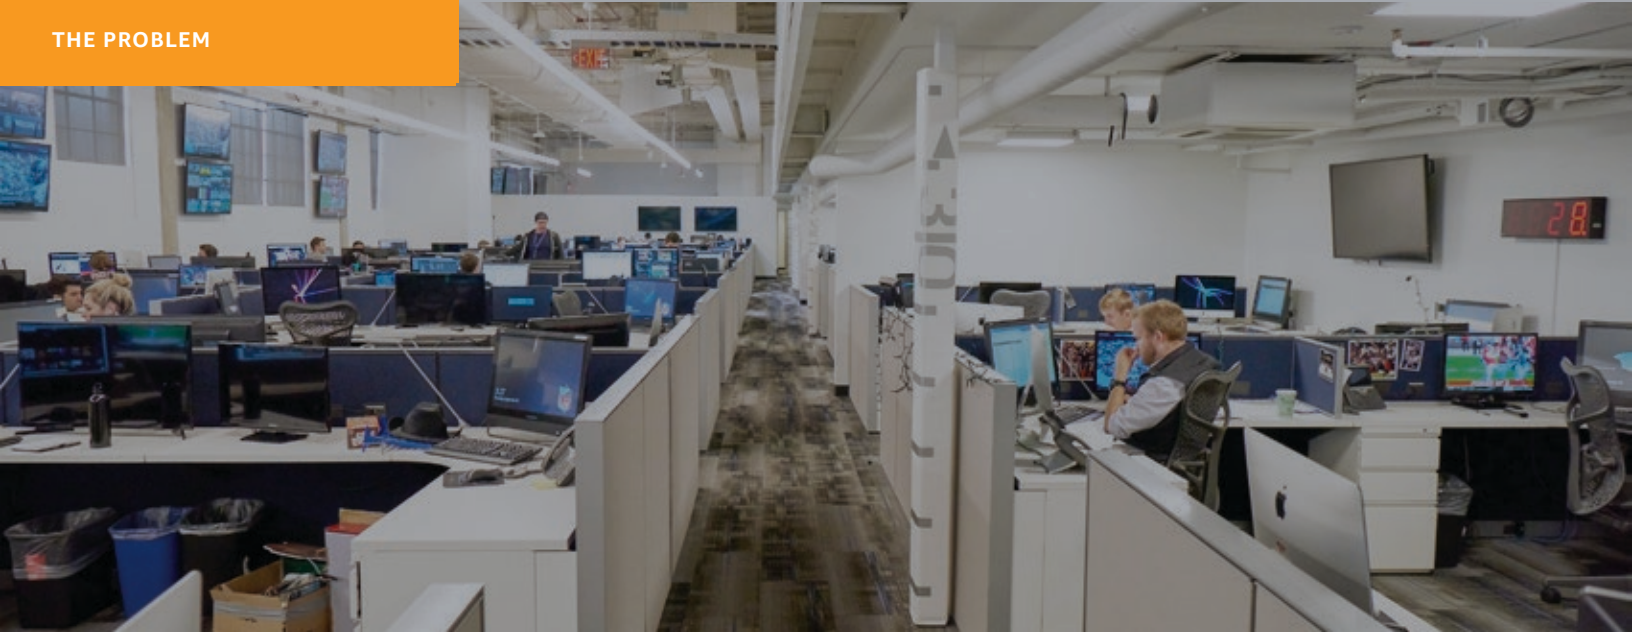
Swensson's team white boarded and brainstormed ideas. What might play well on broadcast? How could a play be better articulated? How risky was that running play or that pass attempt?

It became clear quickly that the most predictive models would need machine learning. And it made sense to work directly with AWS to leverage their expertise.

The NGS team flew to Seattle and worked through their goals and challenges with the AWS Professional Services team. AWS provided an overview of how they could work together. They outlined and workshopped on machine learning solutions and how to work with tools like Amazon SageMaker.

Throughout the process, the NFL's team spent several days at a time in the weeds with the AWS Professional Services team talking about how they'd solve certain problems. They'd prioritize models and work through limitations of the technology.





Why machine learning?

In complex NFL scenarios, business logic specific rule-based systems can start to unravel – there are simply too many variables to process. This is where machine learning comes in to play – it helps calculate complex stats, like Completion Probability, in real-time when there are too many variables to rely on a simple decision tree.

Not all passes are created equal. But with machine learning, the NFL can enhance traditional box score statistics and contextualize events on a per play basis.

"Shortly after joining the NFL in 2007, the implementation of player tracking became a much larger focus," said Swensson. "Once put in place, we started to collect so much data—the question we continued to ask was 'how do we make sense of it all?'"

The NFL didn't need to arm themselves with teams of data scientists—its engineers were able to advance Next Gen Stats quickly through the use of Amazon SageMaker.

> ***We collect so much data—the question was always 'how do we make sense of it all?***
>
> ***- Matt Swensson***









Measuring the impact

By delivering data-rich insights in real time, Next Gen Stats is revolutionizing the way football fans, broadcasters, and players tell the story of each game. Some of the stats translate well into on-screen real-time graphics during games. Others provide valuable color commentary for announcers and broadcasters, source material for fantasy football junkies, second-screen experiences, post-game analysis, and online experiences built for die-hard fans to better understand the performance of their favorite teams and players. And the Next Gen Stats team is just getting started.

Swensson's team has synced the tracking data to video to the hundredth of a millisecond so they can tag key events during plays.

They're also exploring new ideas for automated event detection and new ways to layer other tracking tech with a high bar for accuracy. They're finding new ways to use machine learning in the league to better understand fan engagement, game presentation, modeling the potential impact on adjusting the rules of play, how the game is called, and player performance and safety.

Engagement metrics track direct value—Next Gen Stats social posts can receive over millions of organic views. However, there is indirect value such as operational efficiency gains. Many groups at the NFL can leverage the data and the Next Gen Stats platform to do their job faster – from auto-tagging video with metadata to providing Club personnel with printouts of play diagrams rather than manually charting —the time savings when you add up all of the areas going from hours to minutes, freeing up people's time for higher level tasks, is priceless.

"It's still in the early days," said Swensson. "Our team is excited to see where this journey takes us."

> ***"It's still in the early days. Our team is excited to see where this journey takes us."***
>
> ***- Matt Swensson***

Taking machine learning beyond the end zones

By leveraging AWS' broad range of cloud-based machine learning capabilities, the NFL is taking its game-day stats to the next level so that fans, broadcasters, coaches, and teams can benefit from deeper insights.

"We chose AWS because of its combination of an advanced cloud offering, powerful machine learning capabilities, and experience operating at the scale we need. By powering Next Gen Stats with AWS, we are able to kick off every season with even more impactful and meaningful content, uncovering deeper insights into the game of football than we've ever done before," said Swensson.

One of the biggest benefits of Next Gen Stats is improving the fan experience. In addition to Completion Probability, the NFL has built several new statistical models that can calculate predictions and advanced stats, such as Expected Yards After Catch, QB Risk Index, and Rushing Efficiency.

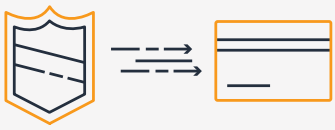
AWS and the NFL are using machine learning to transform how football is analyzed, played, coached, and experienced. But machine learning is changing the game for more than just football.



The manufacturing industry uses machine learning to optimize everything under their command—from supply chain and order management to employee safety and order fulfillment. Fender, a global guitar manufacturer, uses AWS DeepLens in their production factory to choose the best wood to use in its guitars and gain supply-chain efficiencies.



In health care, we're seeing a shift from reactive to predictive care, including the use of predictive models to accelerate research and discovery of new drugs and treatment regimens. Companies like Aidoc, Arterys, and Cerner are taking advantage of AWS' machine learning tools to drive positive outcomes for patients and medical staff alike.



Financial services use machine learning to improve customer communications, predict global markets, and prevent fraud. FICO uses AI services like Amazon Polly to power voice applications that improve the customer experience. NuData Security uses Amazon SageMaker to prevent credit card fraud by analyzing anonymized user data to detect anomalous activity.

Learn how other enterprises are transforming their businesses with the power of AWS >>

Resources and materials to get started with machine learning today

With AWS, you have access to a range of professional services and training programs that let you tap into our experience to accelerate your machine learning initiatives.



Engage With Amazon Experts

The Amazon Machine Learning Solutions Lab pairs your team with Amazon Machine Learning experts to build new solutions for your business.



Choose the Right ML Partner

AWS Machine Learning partners offer a range of consulting services and technologies to help you explore and build the right solutions to take your business forward.



Build ML Skills in Your Organization

Build new machine learning skills in your organization, using the same curriculum used at Amazon—be it for business executives, data scientists, or app developers. On-demand digital training is available at no cost.



Build Your AI and Machine Learning Strategy

This Enterprise AI Guide helps executives tackle implementing artificial intelligence and machine learning in a manageable framework. Supplemented by real-world examples of enterprises that run on AWS tools and technologies, this guide demonstrates the value machine learning can drive in your organization.

Defining the stats

NFL player tracking, also known as Next Gen Stats, is the capture of real-time location data, speed, and acceleration—for every player, every play, and on every inch of the field. Sensors throughout the stadium track tags placed on players' shoulder pads, charting individual movements within inches.

Passing Stats

Time To Throw (TT)

Time to Throw measures the average amount of time elapsed from the time of snap to throw on every pass attempt for a passer (sacks excluded).

Average Completed Air Yards (CAY) and Average Intended Air Yards (IAY)

Air Yards is the vertical yards on a pass attempt at the moment the ball is caught in relation to the line of scrimmage. CAY shows the average Air Yards a passer throws on completions, and IAY shows the average Air Yards a passer throws on all attempts. This metric shows how far the ball is being thrown "downfield." Air Yards is recorded as a negative value when the pass is behind the line of scrimmage. Additionally, Air Yards is calculated into the back of the end zone to better evaluate the true depth of the pass.

Average Air Yards Differential (AYD)

Air Yards Differential is calculated by subtracting the passer's average Intended Air Yards from his average Completed Air Yards. This stat indicates if he is, on average, attempting more deep passes than he completes.

Longest Completed Air Distance (LCAD)

Air Distance is the number of yards the ball has traveled on a pass, from the point of release to the point of reception (as the crow flies). Unlike Air Yards, Air Distance measures the actual distance the passer throws the ball.

Aggressiveness (AGG%)

Aggressiveness tracks the number of passing attempts a quarterback makes that are into tight coverage, where there is a defender within one yard or less of the receiver at the time of completion or incompletion. AGG is shown as a percentage of attempts into tight windows over all passing attempts.

Air Yards to the Sticks (AYTS)

Air Yards to the Sticks shows the number of Air Yards ahead or behind the first-down marker on all attempts for a passer. The metric indicates if the passer is attempting his passes past the first-down marker or if he is relying on his skill position players to make yards after the catch.

Completion Probability

The probability of a pass completion, based on numerous factors such as receiver separation from the nearest defender, where the receiver is on the field, the separation the passer had at the time of the throw from the nearest pass rusher, and more.

Expected Completion Percentage (xCOMP)

Using a passer's Completion Probability on every play, xCOMP determines what a passer's completion percentage is expected to be.

Completion Percentage Above Expectation (+/-)

A passer's actual completion percentage compared to their Expected Completion Percentage.

Rushing Stats

Efficiency (EFF)

Rushing efficiency is calculated by taking the total distance a player traveled on rushing plays as a ball carrier according to NGS (measured in yards) per rushing yards gained. The lower the number, the more the rusher can be considered a "North/South" runner.

8+ Defenders in the Box (8+D%)

On every play, NGS calculates how many defenders are stacked in the box at snap. Using that logic, DIB% calculates how often a rusher sees eight or more defenders in the box against them.

Avg Time Behind Line Of Scrimmage (TLOS)

NGS measures the amount of time a ball carrier spends (measured to the tenth of a second) before crossing the Line of Scrimmage. TLOS is the average time behind the LOS on all rushing plays where the player is the rusher.

Receiving Stats (WRs and TEs)

Average Cushion (CUSH)

The distance (in yards) measured between a WR/TE and the defender they're lined up against at the time of snap on all targets.

Average Separation (SEP)

The distance (in yards) measured between a WR/TE and the nearest defender at the time of catch or incompletion.

Average Targeted Air Yards (TAY)

The average passing air yards per target for the receiver, by measuring the yards downfield at the time of all passing attempts that the receiver is the target. This stat indicates how far down the field the receiver is being targeted on average.

% Share of Team's Air Yards (TAY%)

The sum of the receiver's total intended air yards (all attempts) over the sum of his team's total intended air yards. Represented as a percentage, this statistic represents how much of a team's deep yards the player accounts for.

Yards After Catch (YAC)

The yards gained after a catch by a receiver.

Expected Yards After Catch (xYAC)

The expected yards after a catch, based on numerous factors using tracking data such as how open the receiver is, how fast they're traveling, how many defenders/blockers are in their space, etc.

YAC Above Expectation (+/-)

A receiver's YAC compared to their Expected YAC.

Top Plays

Fastest Ball Carriers

Fastest Ball Carriers shows the maximum speed, measured in miles per hour (MPH), a player achieves on a given play when carrying the ball on offense (rusher, passer, or receiver) or special teams (punt or kick returner). This stat highlights many of the fastest players in the league.

Longest Plays (Ball Carriers)

Longest Plays shows ball carriers that have traveled the farthest during a given play (measured in yards) regardless of yards gained. This stat highlights ball carriers that use the whole field.

Fastest Sacks

NGS measures the amount of time that elapses from the moment the ball is snapped to the moment a defender sacks the quarterback. Represented in seconds, Fastest Sacks highlights defenders with some of the best instincts in the League.

Longest Tackles

A defensive player's hustle can go unnoticed on the field. Longest Tackles measures the actual distance (in yards) a defender covered on a play to make a tackle. This stat reveals players who don't give up on plays and can also indicate players who can truly cover "sideline to sideline."