AUTODESK

Autodesk Vision Series 2024

Integrating a Content Creation ML Department in a VFX Studio

Rising Sun Pictures | Session Transcript

Image courtesy of Warner Bros. Pictures

Integrating a Content Creation ML Department in a VFX Studio

It is no small feat to integrate a new content creation discipline into a VFX pipeline, especially one as potentially disruptive and transformative as artificial intelligence. Head of Technology, Alex Meddick, and Machine Learning 2D Supervisor, Robert Beveridge, take you on a journey into how Rising Sun Pictures embraced machine learning to give storytellers more creative freedom. Discover their successes, challenges and lessons learned as we all continue to embrace this exciting future in front of us.

Watch the recording

Alex Meddick, Head of Technology at RSP

Alex delved into the world of visual effects in 2006, joining RSP as a VFX Editor on Charlotte's Web, over time progressing to roles that saw him lead teams that covered Data I/O, Editorial, Render Wrangling and Frontline Tech Support. During those years Alex also graded several feature films including the critically acclaimed Red Dog and Rolf de Heer's Charlie's Country, and standardised RSP's VFX colour pipeline to ACES. In 2017 Alex's role expanded to include the Production Technology teams, who's Software Developers, Pipeline TDs and Machine Learning Developers create the software and tools that add capability, efficiency and scale to RSP's cutting edge visual effects work.

Robert Beveridge, Machine Learning 2D Supervisor at RSP

Rob's career in visual effects spans over a decade. Coming from a compositing leadership background, he began working closely with the machine learning team when they first collaborated with RSP in 2020. As the demand increased for ML solutions, Rob was promoted to the role of 2D Machine learning Supervisor. His role is based around bridging the gap between the creative & technical potential of machine learning, ensuring collaboration with more traditional VFX departments on project driven deadlines. Most recently he has provided visual effects supervision & worked with clients in pre-production to reach creative solutions utilising machine learning.

AUTODESK

Session Transcript

Alex Meddick (00:15):
Thank you for joining us in today's session entitled "Integrating a Content Creation Machine Learning Department into a Visual Effects Studio". First, I'll set the scene and talk a little bit about RSP and how we use machine learning. Then we'll spend the bulk of the talk showing you what we do and how we find existing patterns to solve many of our integration challenges. Finally, we'll cover how we're growing ML and the ways we empower people to use it. But first, some quick introductions so you know who it is you're listening to and how we fit into this story. I'm Alex and I have 23 years experience in media and entertainment and almost 18 years of those at RSP. I started as a VFX editor and a colorist, and in my first decade or so I managed several of the departments that support production. From that vantage point, I was constantly impressed with the technical artistry of every single person, but surprised by how often the gaps between people cause confusion and slowed things down. From there, I developed a passion for workflows and pipeline and in early 2017, I was invited to temporarily manage our production technology team where I could more directly impact the way RSP worked. I enjoyed it so much I asked to stay.

Robert Beveridge (01:29):
My name's Robert Beveridge and my role at RSP is Machine Learning 2D Supervisor. My path into working with a machine learning team is a little bit unconventional. I come from an artist background of over a decade, and I was a lead compositor here at RSP before I took on the role. So the collaboration and the role came about as a way that we could bridge the gap between the machine learning team and the other creative departments of the studio. Through my artist experience working in show production, I'm able to work with the machine learning team on a day-to-day basis to define workflows and help with their creative direction. My most recent contributions revolve around working with the clients in pre-production to help define how they could use machine learning to solve their creative problems, as well as providing overall visual effects supervision on one of our recent projects at RSP.

Alex Meddick (02:10):
So who is RSP and how do we use machine learning? At RSP, our vision is to create the most memorable screen moments. One of my favorites being the Quicksilver Kitchen scene from 'X-Men Days of Future Past'. We do this by being a trusted creative partner with our clients, working collaboratively to solve creative challenges, giving storytellers more tools to bring their visions to life. The following is a short 30 second showreel that takes you from the very first Hollywood film we worked on 'Red Planet' shot in the South Australian outback to some of our more recent work on movies like 'Indiana Jones and the Dial of Destiny'.

VIDEO: Rising Sun Pictures Showreel (02:53)

Alex Meddick (03:19):
Formed in Adelaide, Australia in 1995 at a pub of the same name, the Rising Sun Inn, RSP's first challenge was our distance from Hollywood. This led us to create cineNet to increase transfer bandwidth and tools to minimize the amount of data sent by identifying and sending only the difference between versions of DPX frames. Communicating and capturing creative feedback was also a challenge. In 2005, we created a tool called cineSync that helped us synchronously playback and annotate on QuickTime. cineSync was honored with a Oscar in 2011 and we feel privileged to see the impact it has had across the industry today. Then in 2006, we signed an agreement with Don Parker and the team to help co-develop a piece of software, which at the time was called Shotgun. It had the promise of retiring hundreds of manila folders with printed information that used to be passed between artists.

Alex Meddick (04:14):
This was asynchronous and time consuming. We were the first customers of Shotgun and I was one of the first end users placing editorial information against shots. Of course, we now know that software as Autodesk Flow Production Tracking. RSP currently employs a little under 400 people in two Australian offices in Adelaide and Brisbane. RSP is one of four brands in the Pitch Black company, bringing together world leading visual effects studios, including FuseFX folks and El Ranchito Together, we have 13 locations around the world, and this place is Rob and myself in that timeline. I started back in 2006 right at the end of 'Charlotte's Web' and moved into the production technology in early 2017 just before Rob joined and our ventures in machine learning began. Why machine learning? Finding innovative solutions has always been part of RSP's DNA and our clients were coming to us with creative or logistical challenges that were hard or impossible to achieve with traditional VFX workflows.

Alex Meddick (05:17):
We believe that machine learning could be used to remove some of those constraints and enhance storytelling if it could be done at the fidelity and quality required of feature films. We'd been in contact with the Australian Institute of Machine Learning (AML), or its previous incarnation since the early two thousands talking rig removal on 'Queen of the Damned' or stereoscopic conversion on 'Superman Returns'. Then in 2018, we took them up on an allocation of credits, which gave us access to the R&D time of doctors, John Bastian and Ben Ward. Given our strict content security requirements and the need to closely collaborate, we set up a space for John and Ben to work part-time in-house at RSP. Right around that time, we had the opportunity to do some R&D on Baz Luhrmann's 'Elvis', a show we already had an award for to see if we could place Austin Butler's likeness onto Elvis Presley in some of his iconic movie rols.

Alex Meddick (06:10):
A global pandemic put 'Elvis' on hold as Tom Hanks contracted COVID while he was filming the movie. Here in Australia, this allowed 'Shang-Chi

and The Legend of the Ten Rings' to be the first movie released with RSP's REVIZE™ technology used to create its visual effects. John remembers our head of 3D at the time, Jason Scott, asking if we could try ML on what had been bid as more traditional comp based face replacements. Our initial tests went to the client who commented they were more convincing than some of their finals. We knew we were onto something now. After finishing those first few films, we began to grow the team. We are currently 10 ML devs supported by several others like Rob. In 2023, we launched our REVIZE™ toolset and this year have shared these capabilities with other brands in the Pitch Black company.

Alex Meddick (07:03):
So, what is machine learning at RSP? Machine learning sits within technology next to many of our software developers under our R&D supervisor, Troy Tobin, but the ML team is really an additional creative department that works collaboratively with other disciplines at RSP. It's a very tight feedback loop with compositing and relies on upstream inputs from many of our Autodesk Maya using teams like assets, layout and animation. We combine machine learning and traditional VFX techniques to deliver feature film quality results. Our primary goal was not to be faster or more cost effective, though it has been great when we can increase productivity and spend more time on the creative. What we set out to do was to create new capabilities for solving creative challenges. For RSP, this has only led to more jobs both inside and outside the machine learning team. RSP has become the go-to for many studios for ML solutions.

Alex Meddick (07:55):
REVIZE™ is a suite of software with heavy integration into our pipeline and our workflows. Behind it is a very talented group of real life human beings, not machines, who spend every day problem solving their way through new creative and technical challenges. REVIZE™ allows us to do a lot more than just swapping faces. We can swap and modify entire heads, hair included, as well as bodies in painting as required and modifying performances with both artist driven and synthetic data. The types of shots we are often challenged with are highly dynamic and require us to be able to effectively deal with issues like motion blur, occlusion, object rotation, and constantly changing lighting conditions. These tools change constantly as our clients give us new challenges and we face the inevitable complexities of getting work done at feature film quality. Here are many of the projects our team and the REVIZE™ tools have been used on and some plugs from a couple of our clients. The following teaser contains some of the work we've been able to use REVIZE™ for. Enjoy.

VIDEO: REVIZE™ Teaser (9:00)

Alex Meddick (10:09):
Now we'll take you through some of the integration challenges we faced and how we recognized over time that although there were differences, workflows emerged to fit the familiar patterns of a creative discipline. Such that many of the challenges we faced

with our clients, with other departments, with the scale of the work, its integration into our workflows or how the machine learning developers themselves needed to learn to be artists, could be adaptive from what we already knew.

Robert Beveridge (10:34):
Thank you, Alex. So today I'm going to expand on integration across the company with machine learning, starting with the origin of every show that we work on, which is working with clients. So first and foremost, I want to talk about the creative challenges and our goals when we engage with clients. There's four main points I want to highlight here, and it starts with collaborative pre-production. It's super important that we can collaborate in these early stages of projects with our clients. That's to determine what their goal is and why they want to use machine learning. We really want to highlight the aspect to use this for enabling storytelling and also solving creative limitations due to either physical limitations or any other possible limitations. We want to be able to push creative boundaries and create outputs that just weren't previously possible before. If we can take a look at the example on the right hand side, we can look at this clip from 'Elvis'.

Robert Beveridge (11:21):
Now, the creative challenge with our client here was that they wanted to use archival footage of Elvis from a music video and they wanted to insert Austin Butler's face into this original footage. That was our creative challenge of the client and you can see the result here on the right hand side. So the next clip is another great example of a completely different use of our ML capabilities. The idea behind this one is that the show had a finished shooting and there was this very large shot. There was no potential to go and reshoot this. The studio wanted to replace a line that the character of Gail, the woman in the middle, they wanted to replace her line here to something else. So this breakdown shows the original performance on the left-hand side and the audio and replacement is our end result

VIDEO: Clip from 'The Fall Guy' (12:09)

Robert Beveridge (12:25)
This obviously showcases the idea that even when logistically it's not possible to be able to get everyone back there to do a reshoot, it's possible to make these sort of revisions afterwards utilizing and leveraging this type of technology. So this point of collaboration is one of the most important aspects of using machine learning and VFX, and that's capturing the right data. So the data itself drives the success and determines the quality that we're able to achieve in the end result. We've been lucky enough to work with and collaborate with both the onset teams and drive capture sessions with the actors or collaborate with some of the best established capture companies to be able to capture both data in whatever form it may be to inform the project. We've been able to also help standardize and develop some of these processes as the process itself develops.

So the reasons we want to do this is because the more predictable that the input is, it conforms to structure in a post-production sense. So if we can determine some specs and standardizations around this process, we're going to make our lives easier when it comes to visual effects and post-production, especially when we consider machine learning. So one of the other things to keep in mind is that each of the data sets that we could capture for these shows are linked to their own respective show. So there's no crossing of data sets between respective productions. All of the approvals and transparencies are with the talent at the time of capture. So to expand on this data capture process, we began to run into a bit of an interesting problem. So when we first started using machine learning solutions, we often came on board a project quite late in the day.

Robert Beveridge (13:56):
So this made it pretty hard because often the shoot had already been finished. That meant that our data sets were whatever had been shot at the time that we were not involved in. This took away from that idea of conforming to structure and meant that we often found ourselves wanting to look to quantity where we would ask for everything that the production had available for data and we would ingest all of that data into our pipeline. This always gave us the mentality that more data was better, but we would find out that as we progressed into developing this process more and coming on board quite early in the processes, that it really was a question of quality over quantity. There would be terabytes and terabytes of data that we had to deal with at time, that would mean we had to ingest and it would create more work for us when the data wasn't necessarily processing.

Robert Beveridge (14:45):
This data wasn't necessarily something that was even going to inform the end result. So if you have a look on the right hand side here, it gives you a bit of an idea of what some of the plotting and analysis tools that we've developed around data sets look like. So what you're seeing is a comparison between a plot of training data and a plot of required data to complete a given task. So we can see on the far right that in the plot of required data, there's a real hotspot in the middle here, but we can see on the left hand side that our data coverage is quite a hotspot on the far left and right, but not so much in the middle. This can inform us where we need data, where it's missing and we can convert these diagnostics back into a show production to inform what we need to complete the task.

Robert Beveridge (15:30):
We quite often refer to this as finding a data needle in a haystack, and when you're dealing with terabytes and terabytes of data, this is something that becomes more and more important. So this takes us to tech tests. So tech tests are a big part of our collaboration with clients in the early days of a project and it often happens in pre-production. So we work closely with our internal sales department to understand what the client's goals are, whether their approach may be based on an existing

techniques of something we've previously achieved or it's an entirely new concept in itself. Often this will be testing before the show material even exists. So as you can see on the right hand side here, we have an internal test with both Troy and Tony from RSP. So the idea of this particular test was to see if we could recreate a performance down for an entire head, down to fine hair detail. So that takes us to the next point of evaluating whether these tests are something that maybe only exists and can be completed in a very controlled test environment or it's something that can allow for many variables and it can be considered production ready.

Robert Beveridge (16:38):
So here's the full screen finished tech tests. In this case, we achieved what we set out to do, which was replicating the performance of the source character for an entire head replacement, including hair at a high resolution. So I'll let this play through for a second. So next I'd like to show another tech test that's using the same material, but this time it's for dubbing. So the idea here is that we can inherit an original performance and drive the character with an entirely separate performance. This follows the idea that we can support something like an international dubbing pipeline as well as other dubbing use cases. So I'll play a short clip here of the results.

VIDEO: Clip of Tech Test with Rising Sun Pictures (17:56)

Robert Beveridge (18:13):
As we can see here on the left hand side, this is the original performance. On the right hand side, it is the driver of what we want the audio to be, and in the middle is the combined result of the two, back on the original character.

VIDEO: Clip of Tech Test with Rising Sun Pictures (18:26):

Robert Beveridge (18:34):
So now I've talked about integration with clients and I want to move on to integration with artist departments and production. So where does machine learning sit in the pipeline at RSP? So it's quite a unique department when comparing it to traditional VFX departments due to the idea that it can both inform and be informed by data departments across the company. So for example, in a machine learning standpoint, we could be both providing out pixel outputs to the compositing team or we could be taking diagnostic outputs and feeding them back to the assets team. It could also be in the form of tools or pixels that end up in final results.

Robert Beveridge (19:10):
So to expand on that previous slide, there are two main contributors of machine learning at RSP. So firstly, it's building tools to assist workflows. We need to ensure that these tools are user friendly, they fit into established workflows that have been defined over many years at RSP, and we also want to stick by the idea that there's no assumed knowledge when using these tools. You don't have to have any familiarity with machine learning to be

able to operate the tools. So they should be presented in ways that are quite familiar to the artists. So if we look to the right hand side here, we're seeing one of our in-house machine learning tools. The task itself being that we want to extract the balls and charts from, which is a pass that it's done for every visual effects shot. In a traditional sense, these would be extracted by the lighting artist manually and overlaid over all of their CG renders as part of the QC process. So what this allowed us to happen was to bypass that by detecting all of these aspects out of the plate and then overlaying them. So not only is our machine learning process just to establish tools, it is also creative outputs as you were seeing previously with both the dubbing and the head swap. So the idea of when we have these sort of creative outputs, we need to be able to unbox these to be able to enable the ability to creatively iterate and keep the tools in the artist's hands.

Robert Beveridge (20:29):
So I'd like to use that example to showcase a bit of a problem that we ran into with this process. If we can consider the face swap example that we saw previously, this was what the workflow was looking like when we first started utilizing machine learning at RSP for creative outputs. So every time we landed on a review, we would have to loop all the way back around to the very start of the process. This meant that we started to see bottlenecks, especially in the machine learning department, which was a much smaller department than other artist departments RSP. So the proposed solution to the problem involved the export of the model itself that the machine learning team were training and handing this over to a downstream department. This solution allowed the creative feedback loop to be kept towards the end of the pipe, especially for more simple updates and this sort of echo, the more traditional VFX departments where for example, lighting renders would be handed over and small tweaks would be taken care of at the very end of the pipeline.

Robert Beveridge (21:26):
So as we handed over these models, we started to realize it was actually a very useful tool for stress testing the machine learning models, finding their pitfalls and gathering information that we could feed back to the machine learning team. So we found a wide amount of benefits for this. It has enabled efficiencies in finding solutions, faster creative iterations, as well as batch running models and seeing how they would react to various cases. So if you have a look over onto the right hand side here, here's a very simple example of something that can be very quickly diagnosed outside of machine learning. So for example, if the input was a dark shot here, and you can see in this case the model's not quite resolving or understanding this input very well, and then in the bottom right we can see that the corrected shot with an exposure tweak applied gives us a completely resolved model.

Robert Beveridge (22:12):
So you can imagine how this reacts in a software when we have this model handed over, we can do any form of pulling this model around on the input to see how the output reacts much quicker than can be done in the machine learning team, and then we can take that

information and roll it back into the machine learning model through the department itself. There are some drawbacks with this sort of methodology. The idea that if we sway too far away from what was intended into the original model, we end up with untraceable results and we really need to take into account managing what has happened to a model downstream and tracking that so that it can potentially be rolled back in upstream if needed. So what does that look like? So I want to draw some parallels here to existing CG render workflows for this comparison to put a bit of a visual towards those words.

Robert Beveridge (22:59):
So if we look at the CG render on the left hand side here, just consider for a minute that this was a flat RGB layer and it was baked down and that would severely decrease the amount that you're able to do with it, right? And the amount that you're able to iterate and creatively move around. If we had a note from the client for example, we would have to send this all the way back to that lighting department every single time as it just wouldn't be feasible to fix downstream. So that's what we were seeing in our early stages of our machine learning workflow where we were given an input here, we would receive a baked machine learning flat RGB result, and it wasn't something that we could manipulate as a post-process very easily. It would always have to flow back through the department. So how do we deal with that in a traditional way?

Robert Beveridge (23:42):
So in 3D renders, what we would do is render a multi-pass that stores all of the lights and utilities as separate passes to make this process possible. So if we look at the example on the left, we can see that we can isolate that particular light. Let's consider the idea that we had a note to increase the light on the left-hand side and make it purple. So we isolate the light, treat it to make it the desired result, and then we see our recombined result. So we wanted to replicate a similar idea when talking about machine learning outputs. This is what it looks like in machine learning form. Using that simple exposure tweak as an example, we can see that the machine learning model itself is handed over and we could see a flow that reads much closer to our traditional departments. So we have input here, any given treatment, and then the machine learning inference model applied afterwards. To demonstrate an extension to this process, with another practical example, we can see our test footage and swap here, but if we look to the left hand side, we can see the desired result.

Robert Beveridge (24:32):
This is just an extension to the idea of the flexibility of being able to tweak the model. So here we can see a more complex case where we want to shift the eye lines of the model and you can see the tool is presented with the idea that we spoke about previously where the tool is flexible, it's something familiar to the user, they don't need to understand the backend, but they are provided user-friendly controls that allow them to creatively iterate. So we can see in this case that I can shift the eyes around, I can shift

the eyebrows and the expression around, and I can also shift the mouth regions. These are essentially limitless controls with the idea that if the model understands something, we can utilize this to creatively direct it.

Robert Beveridge (25:25):
So I've spoken about integration with clients and integration with artist departments. So now how do we bundle this all up and apply it at scale, especially in cases where we might have hundreds of shots. So firstly, we need to align with production. One of the biggest advantages that we've had in the machine learning space compared to others is that we're backed by an existing visual effects company that's built a pipeline and developed tools over 25 years. So the more we learn about how machine learning fits into the studio, the more parallels we start to draw to existing processes and help make it fit. There are still unique challenges that we might see, and in this case, I want to talk about traditional resourcing and scheduling. This can be quite difficult with the non-linear progression that we see in machine learning. With more traditional departments, it's quite often a known quantity of what's going to happen when handed downstream.

Robert Beveridge (26:14):
It's also the expectations of what are going to come out of the renderer are quite known, whereas machine learning still exists as a little bit of a black box where we can make an educated guess, but we're not necessarily a hundred percent going to confirm the predicted result is going to come out of the model. After training, we make this a little easier on ourselves utilizing Flow Production Tracking. So with this, we can track all of our data ingestion and as I said, there can be terabytes and terabytes of data on any given model. So having a place where we can store all of this is a massive help. We can also track our machine learning models as assets within this workflow and track things like machine learning machine GPU allocation to tell us the amount of slots that are available at any given time to train unique models on.

Robert Beveridge (26:59):
So I'd like to jump to a bit of a practical example as I think it's a great example of us implementing this type of work successfully at scale. At RSP, we contributed heavily to George Miller's 2024 film 'Furiosa: A Mad Max Saga'. I think it encapsulates everything that I've spoken about today when referring to integration across the company. So the creative challenge here was utilizing machine learning to make the director's vision possible. What he wanted in this case was to help create a more recognizable Furiosa when she was younger in the film. So Anya Taylor-Joy played the older version of Furiosa, and Alyla Browne played the young Furiosa. So we had to make that transition of the younger to older stages more seamless. So a couple of stats about the show. So we joined the project and from early tech tests to final spanned around one year total.

Robert Beveridge (27:46):
We had 300 shots on the film and 150 of these shots were machine learning based. We had five machine learning devs and 25 compositors. I think this is an important point, that directs back to the idea that we need to make these tools available to artists because we have a massive pool of artists that we're able to use and they can apply their artistry with these tools. So we had 21 dedicated NVIDIA GPUs all up to train models on at any given time. There was 900 terabytes of disc space total created at the end of the project. 50 terabytes of this was training data to inform the models. We had 1500 machine learning outputs and 3000 comp outputs at the end of the show. I'd like to share a short clip of Anya herself talking about our work on the film.

VIDEO: Kelly Ripa (Live with Kelly & Mark Interview) (28:29):
We talked about how they were able to find a young actress that looked like you.

VIDEO: Anya Taylor-Joy (Live with Kelly & Mark Interview) (28:36):
George Miller, our director, had very clear in his head because he was already getting an audience to accept a new Furiosa that whoever the two actors that had to play her that transition had to be seamless. So we don't actually look that much alike. It's AI, which is really crazy. So they were able to composite our faces together and it works on a scale. So at the beginning of the movie, it's 35% my face and then it goes all the way up to 80.

VIDEO: Kelly Ripa (Live with Kelly & Mark Interview) (29:04):
Wow, and did that freak you out at all?

VIDEO: Anya Taylor-Joy (Live with Kelly & Mark Interview) (29:07):
It's strange to see your eyes and mouth on somebody else's face, but to be fair, this is what George wanted and it is seamless.

Robert Beveridge (29:16):
So to conclude, I'd like to quickly touch on some of the challenges that we faced on Furiosa. It's our largest machine learning show at RSP today, and we had a couple of key takeaways once we reached the end. So firstly, it reinforced how important it really was to be able to creatively iterate and keep that door open. So that meant that we needed to make sure that our tools were available and in artist hands. The machine learning team and the compositing team worked in hand here to identify and create these tools. Our main ones in this case were eye and gaze redirection. This allowed us to very quickly hit notes in this regard. At the very end of the pipeline, we also used facial segmentation to isolate various parts of the face and manipulate them as needed. It also proposed the idea of quality and quantity.

Robert Beveridge (30:03):
We needed to manage models across, as I said, the first hour of the film, which meant very many different variables such as environments and how the models machine learning

models handled these were different in every case. We had to make sure that there was continuity across all of these sequences. We also had to make sure that this sort of outputs held up at IMAX quality. This film was displayed on some of the largest screens in the world, and there were some cases where young Furiosa was about three quarters of the screen height. So we had to make sure that we were hitting levels of quality and resolution outputs from the models that weren't previously seen before. So now I'm going to hand it back over to Alex. So I've talked about integration with our clients. I've talked about integration with artist departments. I've talked about integration at scale. So Alex is going to touch a bit more on integration with the pipeline team at RSP.

Alex Meddick (30:53):
Thanks, Rob. The reason pipeline exists in the VFX studio is to manage the tension between creative freedom and consistency and efficiency at scale. As our use of machine learning grew, so did our need to integrate with our pipeline. REVIZE™ combines the work of many disciplines. Our ML devs create the tools for processing data, delivering creative results, and creating efficiencies. The team integrates third party software and develops in-house tools. Our software devs create frameworks that run ML processes at scale, and along with our TDS and artist teams, create the integrations into our pipeline tools and artist DCCs. You would've got the sense from Rob, but managing data is still one of the biggest challenges. We've built several applications that ingest, create, process, manage, and introspect the vast amounts of data we received. These have moved from scripts running people's home directory to plugin frameworks with deep integration into many of our systems.

Alex Meddick (31:50):
Our render farm integrations allow us to run models across lots of shots to figure out what needs attention and to iterate quickly and try various approaches. To allow our machine learning team to align with other disciplines, we needed to empower them with the same capabilities like version control publishing and movie generation to be able to submit their work to Flow Production Tracking and get one iteration closer to done. Our ML devs have a blend of machine learning development and artistic skills. This combination is not something we've been able to hire for, so we have focused on building the team from within. None of the ML developers had worked in VFX before starting at RSP, so we spend a lot of time explaining the role in job interviews. Most ML developers expect it to be developing and deploying technology, creating generalized models to be by others that will work well in many situations,

Alex Meddick (32:40):
But in the creative department, success is determined by the outputs created. So ML devs end up constantly problem solving, trying lots of different things to get a creative result, training bespoke models and developing techniques to deal with issues like motion blur, occlusion, lighting changes or alignment. Creating models that may only be used for a few frames of a shot.

Our machine learning developers have learned to be artists creating content, going to dailies, giving and receiving feedback, knowing what they need from upstream departments to help them achieve their work, and knowing when it's good enough to move downstream when it does make sense to fix it in comp. Machine learning is not the first time that technologists have been the initial content creators in a new discipline. Our industry has been built upon the technical innovations in computer graphics and the people behind them.

Alex Meddick (33:32):
While some disciplines have remained quite technical, the tools and others have developed to such a point where those with artistic expertise can comfortably create. We can already see this shift happening in machine learning, though the daily technical challenges we face currently do not align with the accessibility of other ML tools like ChatGPT or Stable Diffusion. RSP's almost six year journey in content creation with machine learning has been one of the biggest changes to how we do VFX in my time at RSP and we are far from done. The large uptake in machine learning across our industry and pretty much all industries has not only validated our efforts, but inspired us to achieve more. We will continue to grow to meet the demand and ensure capacity to R&D new capabilities. As we take on more work across the Pitch Black brands, the efficiency of our workflows at scale becomes even more important.

Alex Meddick (34:24):
But as always, the most important factor in the equation is the people. As such, ensuring everyone feels empowered is and must be a high priority. Better understanding has led to better and more efficient collaboration, which means we can get to the right result faster with fewer challenges. It could be taking our clients on the journey with us, helping our production team schedule ML tasks, our VFX supervisors give clearer notes or our ML devs collaborate more effectively with other departments. We'll also continue to work towards greater standardization, both internally and how data is captured on set. We'll continue to enable more people across RSP to utilize machine learning, putting more tools in the hands of more artists, not just the tools we create, but we're excited to see what our DCC partners like Autodesk are bringing to their applications that will help us do new things or work more efficiently.

Alex Meddick (35:18):
I sit very close to our ML team and I wish I could let each of you spend a day watching the team work to get a better understanding of the constant challenges they face, whether it is managing and processing huge amounts of data or shot by shot issues that require bespoke solutions. We pull every trick in the VFX handbook and often create new ones. We bring people together from all across the company problem solving with meticulous attention to the smallest details. The adoption of any new technology is not straightforward, although we've identified existing patterns to help our integration challenges.

The process has been anything but straightforward. The feeling of not quite knowing how you're going to pull off something you've never done to meet a show's deadline has been and still is very real. Every aspect of what we've spoken about today is under constant evolution and refinement. I hope we've painted the picture that machine learning at feature film quality and scale that enables our clients to put images on screen that would otherwise be difficult or impossible is no less of a challenge than integrating any of the technological innovations that have come before it. Thank you.