endava

How AI can Improve and Streamline 3D Asset Creation in Game Development

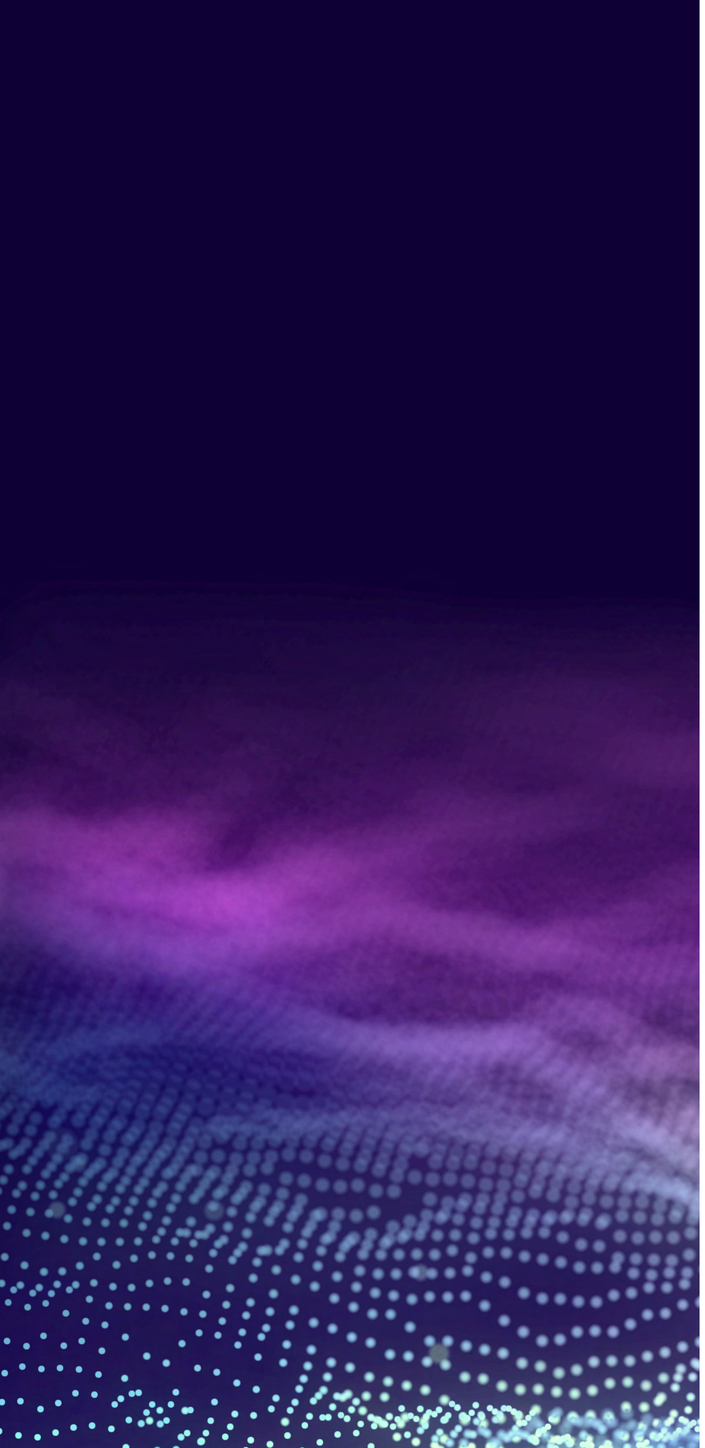

Contents

An R&D project as a showcase

Tech giant Endava has already worked alongside many high-profile and multinational companies on thrilling research and development (R&D) projects in the area of artificial intelligence (AI) and machine learning (ML), but non-disclosure agreements make it challenging to demonstrate their meticulous methods in a public arena.

Consequently, to help illustrate how it takes on the kind of complex commissions its clients ask for, Endava has commissioned an R&D project of its very own. The project outlined in this paper is merely a detailed example of the kind of challenge Endava can – and does – take on, on behalf of its tech-savvy client base.

Whilst considering how much software development has been optimized by automation over the years, Endava noted that art creation – particularly as artificial game worlds get ever more expansive and elaborate, requiring more and more variety, assets, and creativity to fully immerse players – typically remains a manual process, requiring a significant investment in time and personpower to get right.

Through this R&D project, Endava tasked itself to explore further how it could help studios and artists simplify art creation workflows from ideation all the way to the implementation of art assets, as well as cut costs and time. By working alongside artists to stimulate ideation and cut development duration – there are no plans here to replace asset designers, only better support and inspire them – Endava created a faux project brief to support the development of asset creation.

With this task, Endava strives to demonstrate how it deeply researches, validates, and executes the viable strategies it explores on behalf of clients sitting within the highly technical and innovative games industry.

The project: Explore how neural networks can be used to improve and accelerate 3D asset creation

The challenge? We know that developing virtual worlds requires variation and detail to keep players fully engaged with a game's narrative and mechanics, and the cost of game development is ever rising. Right now, asset production – typically a slow and manual process – sees limited benefits from automation, so how can Endava support artistic workflows that can truly save developers money and time without detrimentally affecting the quality of that art?

With its in-house team of UI developers, graphic designers, machine-learning ops, games experts, and data scientists, Endava set out to explore if it was possible to use machine learning and neural networks for 3D mesh and texture generation. The project – which took two months – was divided into five distinct workflow phases: analysis, discovery, proof of concept, evaluation, and product phase.

Before it could begin, though, there was much for the team to consider. How does a 3D artist typically add detail to a basic shape to create unique features? How much influence does an artist need, and how would they interact with 3D models during – and after – machine-learned applications? Do artists feel there is merit in using AI to assist with some of these processes, or are they intrinsically wary of automating creativity?

The team was also keen to assess current tech solutions by exploring existing research and how useful 'out-of-the-box' asset creation tools might be. Could any part of those processes be automated to aid the artist and/or speed up the process? Was it even possible to create an entirely new, bespoke system that could be used to generate assets that offer skin variations to create unique items for virtual worlds?

After its thorough academic research, Endava concluded that as it was using CLIP (Contrastive Language–Image Pre-training) - a neutral network trained on text and images - it could build a cognitive computing solution that "thinks" like a human being, understanding the context in which it is operating. Rather than building solutions for technology-literate people, Endava wants to build people-literate technology.

Endava set out to explore if it was possible to use machine learning and neural networks for 3D mesh and texture generation

R&D: Using Neural Networks for 3D Mesh and Texture Generation

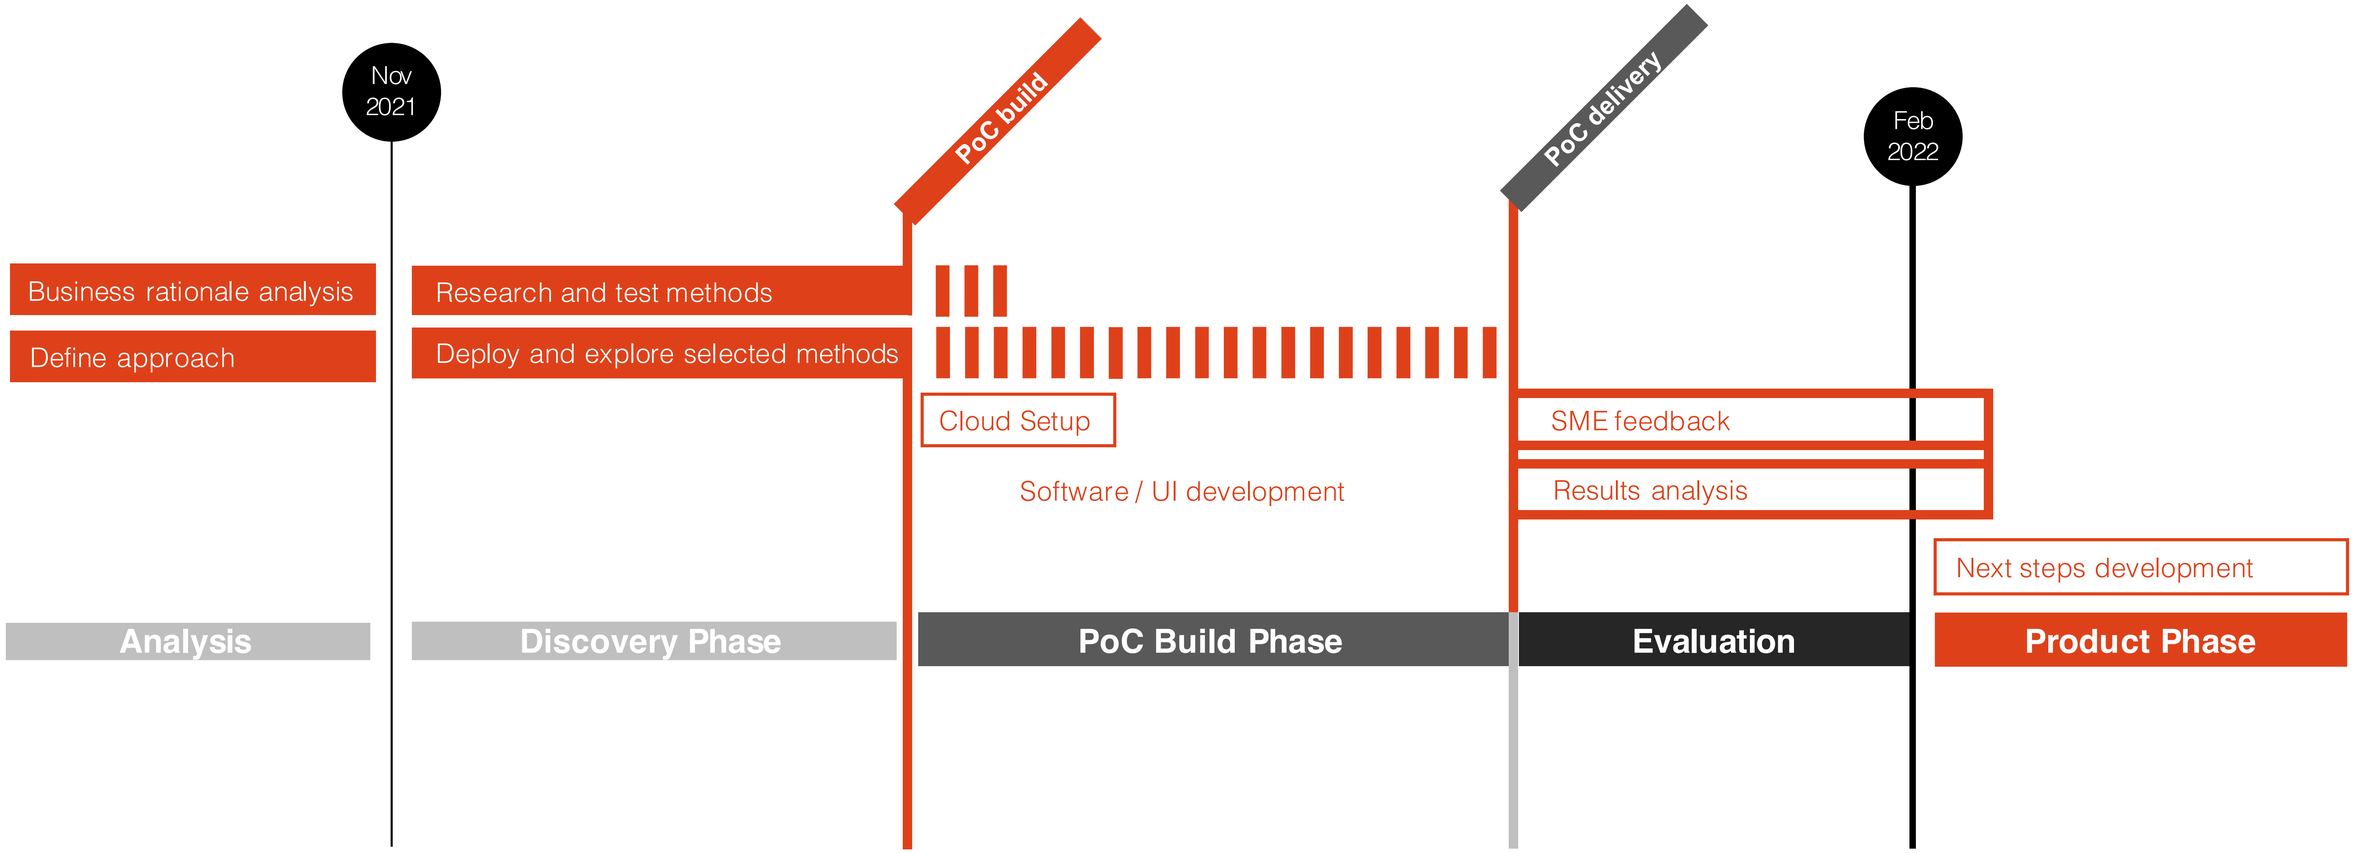

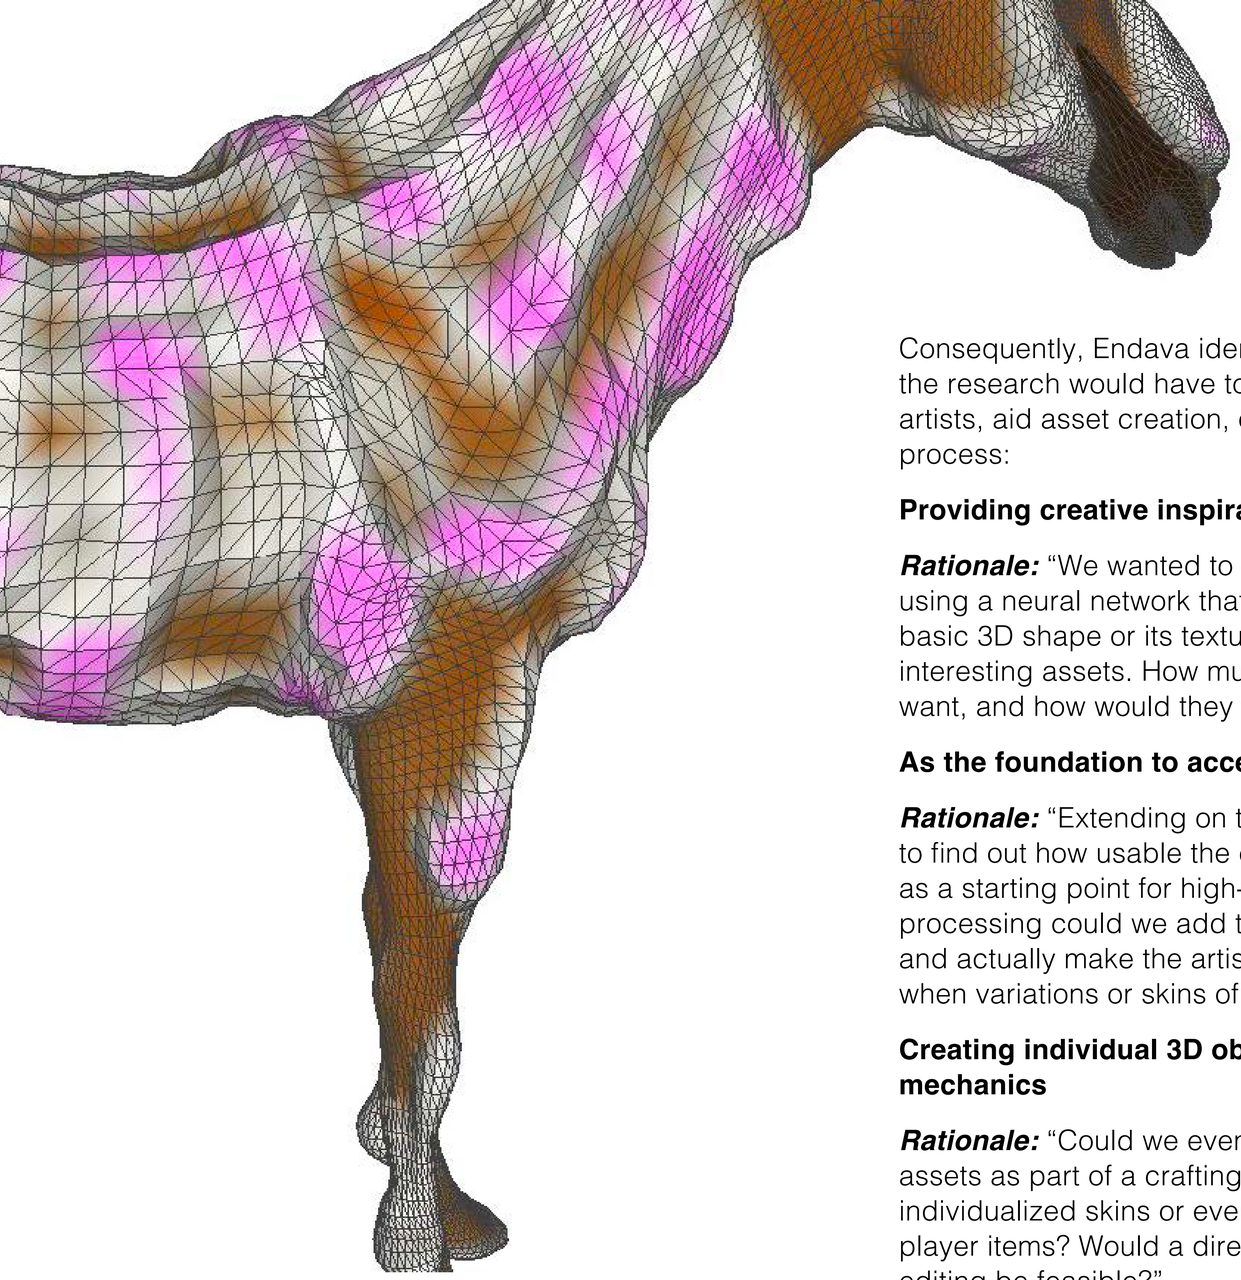

We wanted to explore if 3D artists feel value in using a neural network that specifically adds details to a basic 3D shape or its texture as an inspiration for creating interesting assets

Consequently, Endava identified three key business needs the research would have to address in order to inspire artists, aid asset creation, or even fully automate the process:

Providing creative inspiration influenced by the artist

Rationale: "We wanted to explore if 3D artists feel value in using a neural network that specifically adds details to a basic 3D shape or its texture as an inspiration for creating interesting assets. How much influence would an artist want, and how would they want to interact with it?"

As the foundation to accelerate varied asset production

Rationale: "Extending on the former approach, we wanted to find out how usable the created assets are out of the box as a starting point for high-quality game assets. What auto-processing could we add to facilitate pipeline integration and actually make the artist more productive, specifically when variations or skins of an asset are needed?"

Creating individual 3D objects for in-game crafting mechanics

Rationale: "Could we even come close to using generated assets as part of a crafting or breeding mechanic to create individualized skins or even mesh variations for unique player items? Would a direct integration without manual editing be feasible?"

As well as assessing the tools and innovations already existing in the games space, Endava also performed scholarly literature research to ascertain if there were any solutions to the questions it had posed already available.

Professional tools like Unity's ArtEngine already make use of ML in interesting ways to improve texture work, but Endava wanted to take it a step further and explore mesh creation and alteration. Could ML be used as creative input for artists, or even directly facilitate the creation of 3D assets?

Thanks to that literature review, the team was able to identify that a possible solution already existed, potentially saving hundreds – maybe even thousands – of hours of development time before it had even begun.

Possible methods and software solutions were assessed on two criteria – their usability and suitability for this particular project – and all but three of the 19 possible approaches considered were discarded (the full assessments of which are included in the project's appendices for full transparency) – leaving Endava with the three most promising candidates: StyleGAN, StyleCLIP, and Text2Mesh. From there, the team built a working implementation of an AI-aided asset creation journey, evaluating where this approach was most promising and where there were limitations.

Endava’s preferred solution: Text2Mesh

Text2Mesh takes a 3D mesh input with a carefully constructed topology and strives to augment its geometry and texture to match the textual description provided by the user.

Though it takes the parameters set by the user, it respects the semantic meaning of the object – and it’s constituent parts – as demonstrated in the example below, which shows how from the original mesh, the ML process positioned laces in the correct position on the image, and added spiked soles, as commonly found on hiking boots.

Or, to put it another way: input a plain, ring-shaped asset image, add the text prompt “doughnut with sprinkles”, and the output should, theoretically at least, be a delicious-looking doughnut complete with multicolored sprinkles.

The team concluded from Text2Mesh’s preliminary testing that, coupled with Endava’s experienced data scientists, it could save artists many hours of work indeed. But before the team could check out early and chalk up another successful day at the office, Endava needed to assess Text2Mesh’s full capabilities – and potential limitations, of course – when it came to real-world asset creation scenarios for artists. But how much influence would an artist want, and how best would they want to interact with it?

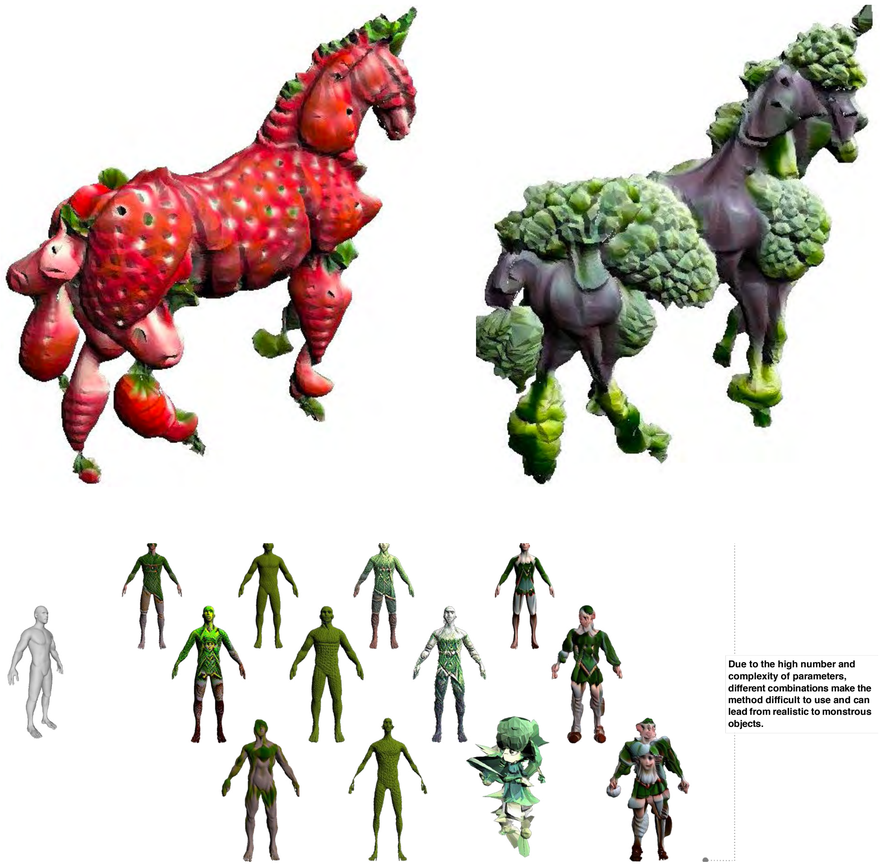

Solution testing and troubleshooting

By using 3D mesh, language-guided 3D augmentation, and post-processing support, the initial results were promising. Experimenting with the parameters and adding different dimensions to the prompts – descriptors such as “strong” or “small”, or even curveballs like “broccoli” or “strawberry” – made for some very interesting variations, some an artist may not have considered at all without the automated process, helping them avoid starting off with an inspiration-less “blank page”.

That said, the procedure was not without its issues, and single-word parameters – in this instance, “elf” – led to some curious combinations, deviating from the original concept and returning some “monstrous objects”, as Endava summarizes in the image to the left:

Similarly, Endava’s testing identified that image prompts – that is, parameters set by pictures and not text – were often not successful, as the image was seemingly translated into a word by the automated process anyway and then, in turn, was used as a text prompt – sometimes erroneously.

The team did, however, identify that the AI had a global understanding of a humanoid figure and was able to apply skin meshes to the human body in a logical and biological way, applying coloration and patterns to arms and legs symmetrically. The one exception was the human head, which malfunctioned often enough to necessitate separating character heads from their bodies to allow for better augmentation.

The AI had a global understanding of a humanoid figure and was able to apply skin meshes to the human body in a logical and biological way, applying coloration and patterns to arms and legs symmetrically

Curiously, the analysis also revealed that different mesh structures – albeit always needing to be "triangulated and normalized" – produced different results, even with the same inputs and parameters.

A significant drawback of the Text2Mesh process, however, was "spiking", an issue where the longer an asset is in the process, the more random spikes are mistakenly generated on the image. Because of the speed and the vast number of machine-learned iterations, the user may not even be aware the issue has occurred and leaving it untreated means subsequent versions will continue to iterate using the impaired image, potentially invalidating the entire process.

Consequently, Endava developed two methods to solve this issue – one based on the Hausdorff distance for 3D spike detection, and one based on computer vision morphology for 2D spike detection – and, now that spikes could be automatically detected immediately when they occurred, the user could be informed and given the opportunity to start over or generate a new object with fresh parameters.

The team then sought to explore how "out of the box" useable assets could be as a starting point for high-quality game assets. What auto-processing could be added to facilitate a pipeline integration that could make artists more productive, especially when variations or skins are required?

Endava then built algorithms into the system to generate more production-ready assets, ensuring that the mesh was fit for purpose and transformational should the user wish to take the generated asset and make further manual adjustments.

A UI prototype was also devised to enable artists to run their own augmentations through the browsers.

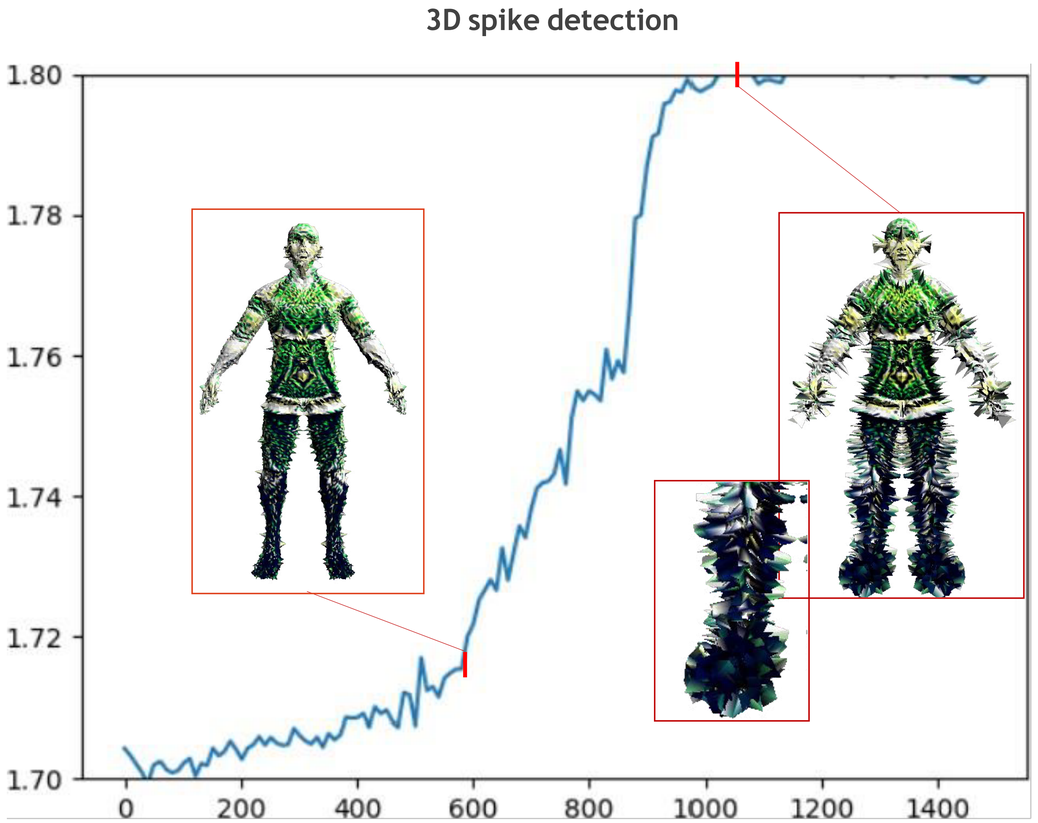

Evaluation

With Endava's bespoke interventions for the Text2Mesh system – such as its aforementioned UI interface – the team was able to generate assets with themes such as 'medieval', 'stardust', and 'colored glass' to create themed objects as well as fine-tune 3D objects to successfully take account of complex 3D geometry such as facial features, applying logical skin variations in certain positions to aid realism and avoid impractical or even invalid results.

Thanks to these adjustments, Endava was able to iterate thousands of sample assets from a single origin image, often in a single day.

The image below painstakingly details how the original picture – a blank, gray image of a generic walking boot – evolved into a viable iteration for any contemporary game with nine hours of further enhancement from an artist.

To evaluate the success of the project, Endava reflected back upon its initial business needs:

PROVIDING CREATIVE INSPIRATION INFLUENCED BY THE ARTIST
Endava's conclusion: Likely usable with additional development
There was a perceived value in not starting with an "empty sheet"
There were good 3D outputs, if not quite optimal
Complicated meshes can still serve as great inspiration
The process could create interesting shapes and extend the mesh complexity automatically
Post-processing for a compatible 3D asset is recommended
There remains the possibility to shape the interface and pipeline to enable an iterative fashion to be more specific and directed
The AI could create a set of outputs mapped to the same style and present them in a fashion that is easily explored by the artist

AS A FOUNDATION TO ACCELERATE VARIED ASSET PRODUCTION
Endava's conclusion: Promising results for smooth geometric deformations worth exploring
Works well independent of attributes, parameters, and meshes
Powerful semantic segmentation – potential in organic and humanoid shapes
Good 3D output foundation that could be automatically post-processed to a certain extent
Working on pre-divided meshes seems more promising
Artist's interpretation and post-processing is likely still needed
There was the potential to set a starting point for a style of asset creation – multiple objects themed in a specific topic
Actual productivity gain and real production applicability to be examined in a wider group

CREATING INDIVIDUAL 3D OBJECTS FOR IN-GAME CRAFTING MECHANICS
Endava's conclusion: Great potential only for a game built specifically upon it
Creation of unique-looking assets to auto-individualize in-game items
High-poly and low-detailed results compared to manually made assets; however, the original number of vertices is maintained, hence premade rigs could be compatible
High-poly models acceptable by game design and architecture
Automatic post-processing to map down a high-res version to be implemented for a higher level of detail
The process could control the attribute space and optimize the input mesh to increase the usability
There is the ability to define crafting ingredients that map to a preferred style and select from multiple outputs to deal with unwanted outliers
Actual productivity gain and real production applicability to be examined in a wider group

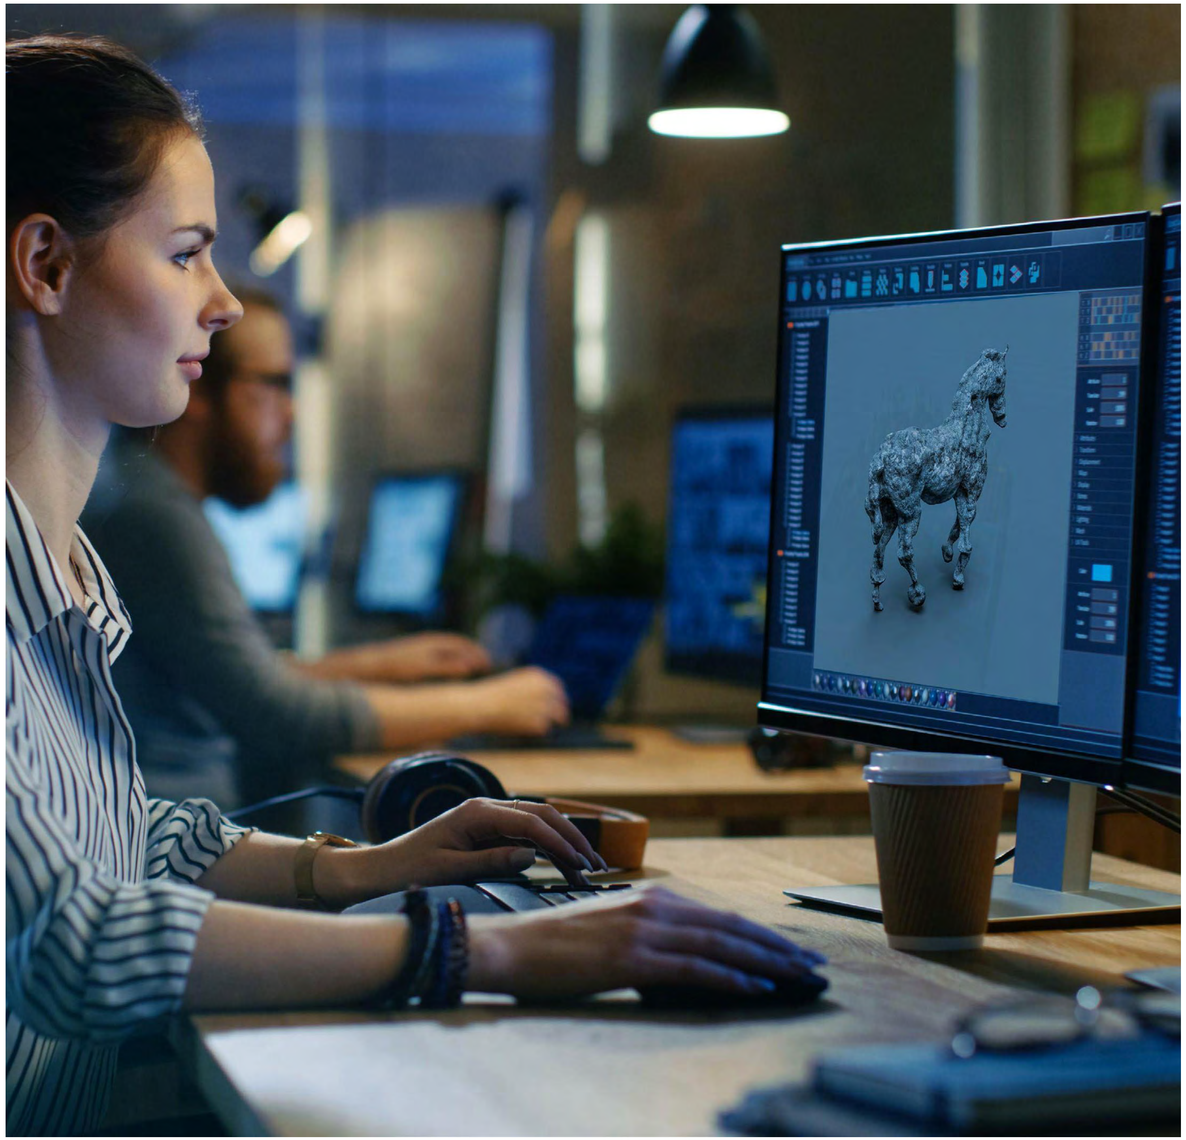

Endava’s solution:

At the end of the project, Endava could see this solution not only working as a tool for asset creation or alteration,but also as a feature that can be directly integrated in a game, such as crafting or breeding mechanics where the player could combine the characteristics of two individual assets into a brand-new form.

That said, in its current state, the present method likely only has limited use, and only if the strange and bizarre assets created fit a game’s broader style and identity.

However, it’s clear that the idea has considerable potential, and a custom solution could bring significant value to the game economy. The possibilities of applying machine learning to game production are endless; not as a far distant innovation, but today, in all stages of software development.

This is Endava: reimagining the relationship between people and technology

You may not recognize the name Endava in the games space just yet, but let there be no doubt: you will soon.

A 20-plus-year-old next-gen technology service provider that employs over 10,000 experts across Europe, the Americas, and the Asia-Pacific region – generating revenue of £446.3m in FY 2021 – Endava is not a newcomer to the blisteringly fast-paced tech world. It has helped some of the world's leading companies accelerate their ability to take advantage of new business models and market opportunities by ideating and delivering dynamic platforms and intelligent digital experiences.

Combining traditional IT services with the capabilities of business and technology consultants and digital agencies, Endava helps its clients find new ways to interact with their customers and userbase, boosting engagement, responsiveness, and efficiency. Scaling Distributed Enterprise Agile, Endava embodies empathy for user needs, curiosity, creativity, and a profound understanding of technologies.

And now, Endava is ready to share that incalculable expertise with the games industry, too.

In 2019, it acquired Exozet and its special focus on the games experience, as well as its capabilities in virtual and augmented reality and gaming technology. Blended with the company's broader tech expertise, Endava is now ready to position itself as a premier game production service provider, offering support across all aspects of games production courtesy of its nearshore delivery teams from right within their clients' own time zones.

This includes end-to-end smart service solutions for its clients from proof-of-concept to production – both during development and at launch – including product strategy, application management, and single-service disciplines like Unity engineering, automation testing, and DevOps. And thanks to its engineering experience, Endava can deliver enterprise platforms capable of handling millions – yes, millions – of transactions per day.

Given its extensive experience in working with Unity 3D and in bringing projects to life, earlier this year the firm joined the Unity Certified Creator Network (CCN), a network that leverages Unity technology to solve customer problems with the power of real-time 3D. It is focused on putting diverse, innovative creators such as Endava at the center of the metaverse economy by fostering innovation through cross collaboration and expanding expert capacity to support Unity to drive the democratization of the metaverse.

Endava is ready to share that incalculable expertise with the games industry

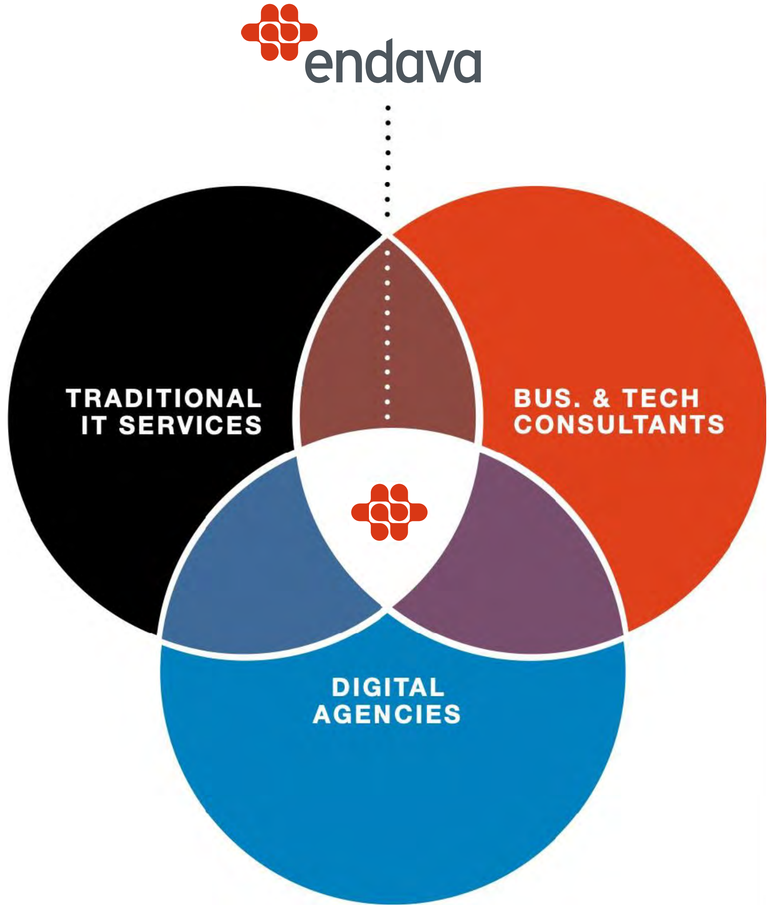

What Endava can do for you

Endava is here to help your studio create better games at all stages of the game production process, from full- service – including ideation, game design, asset and art creation, and animation – to co-development, game porting, backend support, scaling, automated testing, and DevOps and LiveOps. With experience in building games and game services for the past 15 years, it knows game production pipelines are about efficiency as much as creativity.

But it also seeks to help its clients innovate and improve their tools and pipeline, too.

By unlocking the potential of machine learning, Endava believes it will elevate game production to new heights. Thanks to its in-house games and extended reality divisions, the team can also support extended reality development, including virtual and augmented reality and any real-time 3D application, along with user experience and user interface design, too.

Whilst it's true that, on the surface, Endava may be offering the same outsourcing services as its competitors, its strength lies in its expertise beyond the game space.Endava – with its own in-house data scientists, experts, and IT engineers – is uniquely positioned above its peers, as it can bring in its own expertise in artificial intelligence and machine learning to support its clients and develop smart solutions to game development problems.

Recognizing that not all firms can afford to run their own in-house research and development services, Endava's here to do all that for you – and that's in addition to its more traditional outsourcing services. This means all the expertise you need – from art assets to porting assistance to full R&D projects – can be found under one roof at Endava; there's no need to lease the services of external advisers.

To be clear, Endava is not selling the concept or product outlined in this document.

It is merely a vehicle through which the company can transparently outline the methodology and technological approach it brings to all its R&D projects, demonstrating how the team approaches every assignment to support and streamline game companies.

Endava is sharing these findings not to convince you to rush out and try Text2Mesh – as tempting as that may be – nor to try and pitch a similar automated product to you. This paper exists purely to show you the length, breadth, and comprehensiveness the firm brings to all phases of all its clients' projects; be they for new, existing, or even – as in this case – imaginary clients. It was an exciting experiment for the team that showcases how Endava, with its interdisciplinary teams and great expertise both within and beyond core game development, can help its clients with research.

Thanks to its team of in-house data scientists, Endava can apply a vast range of expertise to each and every project and not only identify fresh, lean, agile ways of working, but also identify new and existing workflows, too.

The scope and depth of its preliminary work with literature reviews ensure that time and money are not wasted reinventing the wheel; wherever possible, the team will explore existing processes first and individually test and trial each one to assess its suitability. Comprehensive testing during the discovery phase not only seeks to identify bugs and issues ahead of deployment but also to find and implement solutions for them, too.

But best of all? All of this expertise – the data experts and software engineers and the graphic designers needed to research, test, refine, and implement these projects – are on the books at Endava, expediting the discovery process and minimizing both costs and time, while the IP remains with the client.

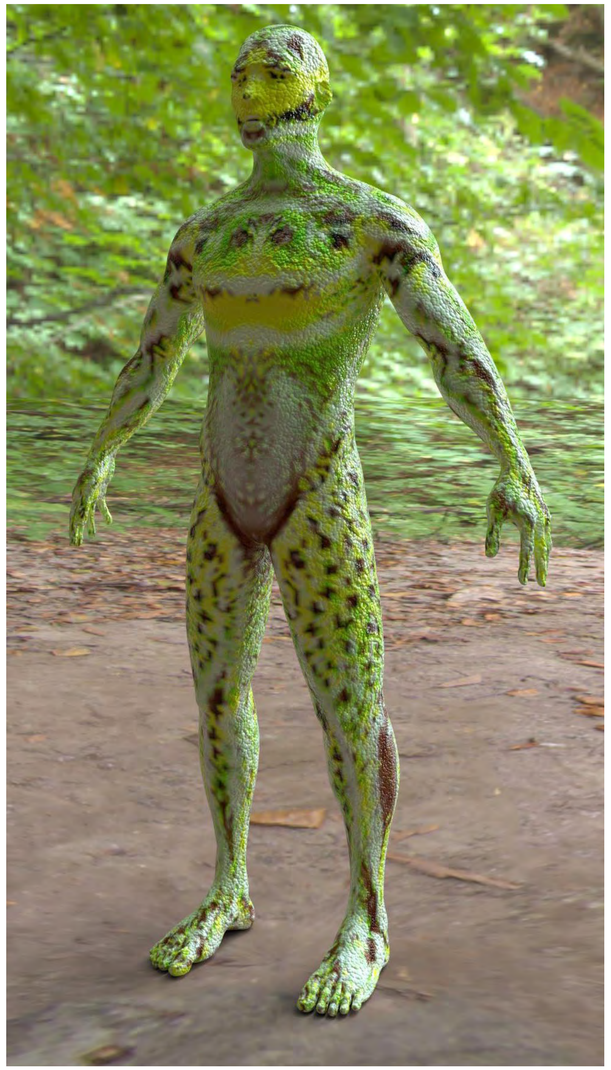


endava