>wired_

SEGA

THE
ZINE

2ND ED.

INTERVIEW WITH CANVA MACHINE LEARNING ENGINEER (EX-MAC PRESIDENT!)

With AI content EXPLODING across the world, more and more people are curious about the world of machine learning.

If you ever wondered what it's like training AI models at one of the leading design platforms in the world, then join us as Sai Kumar shares his journey from a student leader at MAC to a Machine Learning Engineer at Canva.

Q: Please introduce yourself and tell us a bit about your background!

I'm Sai, a new Canvanaut and I work as a Machine Learning Engineer. I've been here for a month and I was formerly at Monash, previously the ex-president of MAC (Monash Association of Coding). During my time at Monash, from 2020 to 2023, I studied Applied Data Science, [which] involves Maths and Computer Science. [The degree] was a blend of both things. [In my] final year, [which was my] honours year, I took a research project that was based off my internship at Canva, [which] involved photo AI.

Q: Can you expand more on Photo AI?

So it could be anything when I say photo AI. It could be like looking at an image and saying: "Is this a dog or is this a cat?" Or, you know, picking out a human. It's like computer vision because you have cars that are trying to see people. That's one AI related thing. Then you have image generation in particular. So how do we turn a text description from a user into an image? And so this is called generative AI. It's all the buzzwords lately. But in particular, I looked at how we can look at the representation of people in Stable Diffusion.

Say, a photo of a CEO, and because the training data — which is the Internet — is composed of lots of men, it turns out that the representation of a photo of a CEO — say 100 photos on average — tends to have 90% of them being men, maybe ten of them a woman. So how could we make this more equal in terms of, for example, gender or ethnicity or any other group?

For what I do as a Machine Learning Engineer, I'm specifically in the Photo AI Team. We look at cool photo effects. So, you know, if you've used Photoshop once in your life, or maybe even Canva, you may want to erase people out of images, for example, or remove the background image. We do all that sort of stuff, so we create models that are trained to remove backgrounds, edit objects, and basically make your images look better.

My line of work specifically looks at, hey, we have this object here, this is a model that can turn one thing into another. So if you give it a text description, it'll turn into an image. The user may want to generate, say, a photo of a bear dancing in a circus. Then you'd basically ask for it in the Canva app. My job is to wire all of that together and bring it to the user. So we look at making the model good to use, safe to use, and fun to use.

Q: How did you end up working at Canva as their Machine Learning Engineer?

So during the year of 2022, I was looking around for internships, so I looked for internships in my penultimate year – that's the year before your final year.#

In particular during September, a Machine Learning Engineer internship opportunity opened up at Canva. I decided to interview for it and I got it. So over the summer I interned and I was looking at that project where we were trying to make Stable Diffusion more fair to everyone. So it was a three month internship, 12 weeks, [and] they gave me a return offer [after the internship finished]. So it was like, you can come back here and you can continue the work that you're doing. That means adding all these magic palette things, [so] we call "Magic" at Canva.

Basically integrating all these Machine Learning products, researching stuff. The sort of stuff like removing a background image comes from research. So it's like, hey, we have an idea in mind for messing around with this stuff. So there's a research side of things and then there's the engineering side of things.

I'm currently somewhere in the middle, so I do a blend of research work and Engineering work.

Q: Do you prefer doing research or engineering work?

Some people really like engineering work because it's very deterministic in the sense that: "Oh, well, if you write this code, it's like building a house, you're laying bricks." It makes a lot of sense.

Research — on the other hand — with machine learning models, is a craft. It's empirical. You've got to build off your intuition of how these models work, but basically the way it learns is something like how humans learn. So that's the classic analogy we use for them — Call them neural networks based off the neurons in your mind.

Learning to understand how they work and having the insight and experimental ability to run experiments, to discern results and be like: "Hey, I've [got some] into intuition about this, this could lead somewhere," is very frustrating at times, but it can lead to really good outcomes because you find something that works amazingly.

It's great, but a lot of the time you'll be hitting dead ends. So there are points where engineering work is: "Hey, you did this, that's like completely solid." It's [more] quantifiable and it's not as hard to do as research generally.

Learning to understand how they work and having the insight and experimental ability to run experiments, to discern results and be like: "Hey, I've [got some] into intuition about this, this could lead somewhere," is very frustrating at times, but it can lead to really good outcomes because you find something that works amazingly.

It's great, but a lot of the time you'll be hitting dead ends. So there are points where engineering work is: "Hey, you did this, that's like completely solid." It's [more] quantifiable and it's not as hard to do as research generally.

Q: You mentioned getting an internship in your penultimate year. Did you receive this through the IBL program or through your own sources?

Yes, so my course is technically not [involved with] the IT faculty — it's under the science faculty. So I don't actually have access to the IBL course [which] is restricted to [the Bachelor of] IT, Computer Science and Software Engineering students.

In my case, there is a replacement called Work Integrated Learning. So some other courses have this as an optional thing where you take an unpaid internship sometime during your second or third year, basically your penultimate year, and that's basically credited towards your course. So Science has [a] form of a science internship. You can do it with commerce as well, and a few other faculties.

In my case, [Work Integrated Learning] was basically mandatory. It was baked into the course. I did that over six months, so that was in semester two of 2022. It's a sort of capstone unit in that sense.

Q: Would you still have sought out your internship by yourself even if it wasn't mandatory

Yes. This is a pretty complicated question because there's always some risk and reward. You have to venture out, and as a first year or second year student, you're not going to have a good idea of which is the better option. I'd say the general rule of thumb is to take any internship opportunity you can. In this case, there wasn't any overlap. But in the case of IBL, typically, if you're following a three year progression, you're going to be doing that at [around] the start of your final year and that will overlap with an penultimate internship opportunity. You have to choose one or another.

This isn't a problem with other degrees. So if you're doing a double degree, you're going to have to do IBL at a different point and you can still do a penultimate internship. **So it really depends on what your course progression is like. internship now or sooner.**

underloading, that means you might do IBL at a different time, and you may be able to do an internship now or sooner. So there is no one size fits all answer.

There are good companies that are doing the IBL program and there are really, really good internship opportunities. But having any sort of opportunity is really good. So internships on your resume are amazing, especially if you go for graduate job hunting. A lot of these places do offer internship return offers, so you may [be able] return to as a graduate engineer.

Not all of them [offer it to] everyone, [but] it's always worth having that sort of experience because it's very different to [being in] university.

Q: You mentioned being the ex-president at MAC, do you think extracurricular activities helped you in your internships application?

Yes! Let's talk about extracurricular opportunities and the impact that they have on your candidacy. At the first stage in the resume screening process, they see: "Oh, you're involved with the local coding club or, you know, relevant IT clubs. It means that you are involved in leadership or potentially helping out the community. It's a good signal!

I've been told by recruiters that they look for that sort of stuff. There's some sort of involvement that you're doing more than just your coursework, right? So that's the first step. Then, when it gets to interviews, you'll [often have] the behavioural interview.

[This involves] pulling out [your own] experiences about times you've resolved conflict, [worked in a team, etc.] This could come from your university coursework, but this is very standard, right?

You have more opportunities to pull from if you're involved in extracurricular stuff. So if you're involved in a committee, that means you could be involved in the planning, [for example, industry night]. Or, if you helped organise a hackathon, that means you were definitely involved [and shows that you're capable]. There are a lot of stories to be formed there that you can tell and it blends in naturally because that sort of stuff is relevant to the field you're doing. So that sort of work [you do] with others has a good impression on the people who are interviewing [you], because when it comes down to it, they're trying to see if you're a fit for them. So that means: Can you work with others?

If you're working in a committee, club or student teams, it means that you are already probably good at interacting and working with others.

Q: What is your favourite part of working at Canva?

Free food every day for lunch! But … [honestly] they really care for their employees. That is actually a good part of it. But even aside from that, I [love] worki(ng) with really interesting stuff.

Everyone uses Canva … I know many societies that use Canva. It means I'm always talking to someone who's like: "Oh yeah, I use that product!" So that in itself is pretty satisfying because I worked on that feature and work on the photo editors stuff. So if you have ever edited a photo, you may have used a feature that I might have been involved in. That's awesome.

The other part of it is [that] it's interesting and challenging enough that I want to keep tackling it. It's not stuff that I'd get bored doing. [You don't] do the same thing over and over. It's new [and] refreshing challenges that I really, really enjoy.

Q: What drew you to this field?

To give some background, I started coding around year nine, year ten. That was general stuff to do with modding Minecraft, and I didn't really understand what I was doing. At some point I got into reverse engineering mobile applications. So if you've ever played all these games like Genshin Impact, I've tampered around with those games. So that was a lot more to do with cybersecurity and reverse engineering at the time.

One of the clients who bought my services regularly suggested I just try out Machine Learning and data science because that's where it's at these days. So I gave it a shot and I was really, really obsessed with the idea of generating text at the time. Basically one thing I did back in the day was take a bunch of Percy Jackson books and feed them into a model to try and generate new Percy Jackson books. This is something everyone knows nowadays as ChatGPT – this is way back before these things existed. So seeing the technology evolve with you was really something to behold. So I kept chipping away at it, even though I didn't understand anything at the time.

When university came around, I decided [to do] Applied Data Science. I'm really interested in this sort of stuff, which was more or less a blind decision because that was the first year the course was running, and it paid off really well. It was a gamble, but it worked well. Ever since [then], Machine Learning is all the rage these days, everyone's in it.

But underneath all of that gunk is a lot of really, really interesting problems that just require a lot of experimentation to play around with, and I think there's some beauty in that.

Q: Talking about passion, you mentioned in one of your articles "primetime" and how continuously fostering your passion will help you succeed in your field. There are now a lot of people coming into the IT, Computer Science and Software Engineering field, and how do you think someone, who – for example – may not have the time to expand on this specific passion, competes in the market?

Currently, we are in a really bad market for jobs. [A lot of] graduate roles have been slashed, and [more] people have been laid off en masse. It is not a good time because less places are hiring and that means people are willing to put in every last quarter to get that job. If you're aiming high, then you should be prepared to put in the work.

With that being said [though], I think it is very healthy, very good to have hobbies outside of your [studies]. Have a hobby that is separated from a screen because once you start working, your opinions on coding in your free time may change.

When it comes to internship seasons [though], I think more time should be spent to prepare for it. So perhaps the year before, or the summer before, you start doing a coding problem a day. Or maybe if you're really passionate, work on a project that interests you. I don't recommend working on a project just because you want it on your resume and you can talk about it because you're going to hate it a lot. Do something you're passionate about. Incorporate it into your grind and into your preparation.

For example, some people enjoy doing Leetcode. Some people don't. It's one of those things you just got to swallow. You have to get through it. So having some amount of grit is immensely useful. I can't understate this enough. You shouldn't be spending all your time dedicated towards grinding that is unhealthy and unproductive, but you should be able to — first off — understand yourself and how you learn. That comes through trial and error. Secondly, come to a point where you are satisfied with how much you're grinding.

You need to understand [your goals]: Is this enough for me? [Is this what I want?] Set goals and aim high. Test your limits. If you're aiming for those places, you need to be putting it down like 100% and a bit more.

If you're not as invested and you just want a job, I think it's fine to do a mild amount of preparation. So across all facilities, you want to do your technical preparation, you want to do your behavioural preparation, and you just want to get your resume out.

Basically, you just need to gauge how far you're willing to go.

Q: You mentioned the market being difficult to break into right now. How difficult was it for you and how difficult do you think it will be?

This is [actually] an interesting thing because I've seen a lot of people struggle at the moment and it's somewhat unprecedented because we had a really good era around 2021-2022 where companies were hiring like crazy because of the COVID pandemic. There were more job opportunities. So anyone was willing to hire.

But as we have come to what looks to be a recession [because] companies are becoming more conservative with hiring. This means that the justification for hosting an intern is weakened because taking on an intern, or even a graduate engineer, for the first six months is a negative. It's a loss to them because they have to train you up. You're not contributing to it.

With the more conservative hiring policies, they are not willing to take as much risk, it means there are less spots. Around the time I [was] interviewed was where it started going downhill, honestly, and I got very lucky on that front. One of my friends was interviewing for [a company] at the time he got in and then immediately after they froze headcount. What that means is that they weren't hiring any roles after, so that was a really close call.

Even at my company while I was interning, they had what you call a soft freeze on hiring. After that, I kind of had a sense that it was going downhill. But this sort of thing only lasts a few years, so you want to be vigilant about it and just apply everywhere – you will eventually end up somewhere.

You just need to be super vigilant and proactive in doing so. The market is bad, but you can do a lot to stand out from the average applicant. It's more important to focus on your projects, your extracurriculars, and just really build up your portfolio to stand out.

Q: You mentioned selling your services to a client. Is that your first ever job?

During high school, I made a bit of money here and there, basically selling services [like for example] run(ing) missions for my client on their game account and I'd get a bit of money from that. It was a monthly service sort of thing. I wouldn't call that an official job. It did make me money and I bought a PC with it.

At the start of university, I was tutoring mathematics, and I did that for a year or so before I decided to focus a bit more on university. Around the end of the second year, I had the chance to do research work. So I [became] a research assistant, and I basically [looked] at some deep learning work and helping out a professor with a publication or project

In terms of my first official job, I'd say Canva is my first [actual] full time job.

Q: What soft and hard skills do you think someone who is trying to break into your industry should have?

As a general Software Engineer, you need to be able to talk to people — [the job] has a stereotype of you don't have to talk to anyone and just write code all day, but a majority of your work will be asking for clarification from others about, for example, what this code does, asking for permission to do something or getting your code reviewed. A lot of that requires some amount of decent interpersonal communication skills. So you need to be able to talk to other people, work with other people and collaborate. That [definitely] is a bare minimum.

As for hard skills, it is harder to pinpoint what is necessary [because] those can vary a lot, so I can't give you specific advice in this case. For Machine Learning, everything is done in Python, but you need to know your theory really well. It is one of the fields where being well versed in your theory is much more crucial.

Q: Do you think networking is important to breaking into this industry?

Networking is always a good thing, because even if my own story suggests the contrary, it is always invaluable to have. So in my case, I got into Canva without a referral. People commonly network or make friends with people at other companies to get a referral set out.

But when it comes to networking, it is more [important] that you get connected with students who are like minded. For example, if they're also interested in applying for these companies –maybe they're already interning at these companies – you want to be able to network [and get to know them].

I don't think there's as much value in networking with companies because they basically tell you about the intern [and] grad program and you don't find out too much more about them. That's in the case of a recruiter.

When it comes to an engineer, there are a lot of good questions you can ask! You can ask questions about the challenges they face, interesting problems they solve, [and] some people have very interesting answers about that. It'll give you insights about the company.

I'm also just good to know lots of people because the more you know about jobs, the better off you are.

Q: You actually have your own website. Would you like to talk a bit more about that? Is it your personal project?

Shameless plug — my website is https://www.saikumarmk.com/.

So why did I make the website? It's because I had a lot to express — I had a lot to talk about. Though it doesn't seem [to fit the] stereotype [of STEM students], I've written more [on the website] than my honours thesis (which was around 12,000 words). But my "Graduate/Intern Playbook" came out to be around 15,000 words, so I definitely spend a lot of time doing that. It is a good skill to be able to communicate clearly and well. That's why I do it so often, because I think there's value in sharing this information with the community. There was [also] another maybe selfish part of me where I wanted to see Monash succeed. As far as big tech companies go, and quantitative trading firms, they hire a lot of students from Sydney specifically.

What if we just got all the information and advantages [the Sydney students] had and just gave it to students in Melbourne. I want to see representation go up and some of my friends also share that sentiment. Working on that sort of stuff, [making] this information available to others is invaluable because everyone benefits from it.

[For example] someone sees my playbook and they're like: "Oh, now I know what to do in interviews." And it happens a lot. I helped someone [during my time at MAC] and they [actually] got into Google. [For me,] it's really heartwarming to hear.

Q: Do you have any long term career goals, for either your career or for yourself (i.e. your website?)

I was always interested in research [so] there's a possibility I might do a PhD. That's something still up in the air because I love Canva and I might stay there for a while.

As for goals around the website, I've basically published my entire Playbook series so I'm not actually sure what I'm [going to] doi now. I think I might do something related to my university experience [and] writing up a memoir of how I found university, [and the] challenges I faced. I think that'll probably be my next major thing, but other than that, I don't really have any major goals on the site.

Q: When you are working in your desired job, do you think you should do something that is outside the scope of your industry? Should you be doing your own projects? What is your take on it?

For most software engineers, upskilling usually can happen within the workplace. You should be constantly looking for new opportunities, new areas of growth. You want to be working with a lot of people, talking to them. Networking is invaluable [because] there [may be] really good opportunities and let you know when [the company they are working at] has a vacancy. It may interest you [and allow you to get your job easier].

When it comes to technical upskilling, projects of the personal sort don't really cut it compared to actual production work. So your mini to-do list one compared to like the massive code base you see that, for example, powers Canva, [and I think] working at that scale is in itself going to upskill you a lot. It's going to prepare you a lot.

With that in mind, not everyone works at places that have amazing engineering practices. So you may want to go upskill somewhere else, [and] I'd recommend looking around for other jobs that may interest you.

Additional Footnotes: Sai's Tips On Understanding Recruitment

[For] anyone still at university looking for internships or graduate roles, recruitment can be confusing. If you get ghosted, remember that recruiters are dealing with a lot of applications. For example, Canva gets about ten thousand applicants. Some places get even more. Imagine ten recruiters sifting through all that—it's going to take time.

So, be patient [and] apply like crazy, but don't put all your hopes into every role or one specific role. Just go for everything.

Once you get experience, it's good to move around, but as an intern/graduate, you're trying to prove yourself with no substantial, tangible proof. Even with committee experience or side projects, it's more about showing potential. Recruiters have to take a leap of faith [that] you can do the job.

The whole recruiting process is somewhat involved. I've written about it on my site [because] there are a lot of aspects that students don't know. For instance, in technical and behavioural interviews, you get marked off a rubric contain[ing] things like problem-solving skills, coding ability, and communication. . So [in the] case of the technical interview, you want to be talking to your interviewer as much as possible, because that way it's demonstrating that you can work with someone like when [doing] peer coding.

Part of it is investigating and exploring yourself. Looking at online job boards and finding opportunities on your own definitely sets you up. Having some level of autonomy and the ability to do things on your own is crucial, especially in a job. Being able to do stuff independently is amazing and you will learn a lot from trial and error.

Special thanks to Sai Kumar, Rohan and RAD-MON for facilitating this interview

Did you ever find yourself on Midjourney trying to make a prompt for the site to give you the best image of your wildest dreams? Did you ever have a favourite character nobody makes art for that you are so deprived of content that you turn to an AI image site and ask "Please, AI, please give me Natsuki in a [REDACTED]" (please don't ever do that)

The recent revelations of AI image generation has been taking the world by storm. It has become more and more accessible to generate art simply with a few words and lines of prompts, and a knack of creativity. A few seconds of processing a user's request, and AI images like DALL-E are able to spit out an amazing image that adheres to the prompt given.

The possibilities are endless, and as long as the resources exist, AI images will continue to develop. Now you can make anything from an old fashioned Capybara dress from the late 1800's to hugging your favourite animal in a city in Mars!

But you might be thinking. What actually happens when you ask a computer to generate AI art? After all, it isn't something like generating an answer to a difficult coding task, where the model goes through many, many sources it was trained on in order to create a specific answer catered to your query(Guinness, 2016). However, would that type of structure be the same for generating images, per se? After all, generating an answer based on articles would be different from going through images, after all. That introduces some more questions. What is so different about generating AI images compared to generating AI responses? Are there any limits to its capabilities? What problems could we potentially face? Well, we'll find out here!! 😜🔥

A BRIEF HISTORY IN AI FROM GOFAI, GAN TO STABLE DIFFUSION

The evolution of AI image generation has come so far. We start out by the earliest record of AI image generation, called GOFAI, short for "Good Old Fashioned Artificial Intelligence"(Poltronieri & Hänska, 2019) back in the 1960s. Pretty cool right? Hobbes' goal

by creating GOFAI was to create some sort of cognitive machine. He wanted to create cognitive results that were able to be on par with human behaviours, such as the ability to write, read, draw, etc. While the overall concept and program had its fair approaches and ambition, there were too many limitations during the current time period. Even after decades of commitment into trying to develop this AI technique, it was said that it had reached its unfortunate dead end.

But the evolution of GOFAI didn't stop there. According to (Poltronieri & Hänska, 2019), it provided the opportunity to expand on AI image generation using its current framework. Harold Cohen, an artist during the 1960s tried to utilise AI in his creative process. Cohen would then continue to delve into his research and create his own software which emulated how an artist would typically think when creating their own form of art. This software would be called AARON, Cohen's own version of creating AI art!!

However, due to constant limitations (AARON wasn't able to comprehend figurative images, as it was too complex for the software at the time) and constant criticism, Cohen decided to retire his pursuit of optimising figurative art and would dedicate his life to abstractionism.

The first colour image created by AARON at the Computer Museum, Boston, MA, in 1995. Collection of the Computer History Museum, 102741168.

And now we are at the stage of GANs. Before that, we must understand the concept of machine learning first. Although these machines don't necessarily 'learn', so to speak. What actually happens in machine learning is that they will try to apply mathematical

algorithms over a collection of inputs. Unlike its predecessors, machine learning typically doesn't require humans to supply all the knowledge it needs. It tries to gather knowledge and experience (despite limitations that it may have)

GAN is an example of an unsupervised learning architecture, meaning that it uses machine learning to analyse and train itself over unlabelled datasets. GAN attempts to utilise deep neural networks to generate realistic data and images based on the numerous datasets it has been given. It uses a concept that can be similar to creating art and cross referencing. Say that you are trying to make yourself a painting based on an idea. You try to find numerous paintings in an art museum as a reference for your artwork, while your friend tries to search on the web to find the real image versus a counterfeit. This is the same for GANs. The generator generates an image by taking samples from numerous inputs, while the discriminator would be the one to differentiate which image was generated by the model, and which image is the real deal. The result of an image generated by GANs would usually be ones which have sharp images close to resemblance to the original image. GAN would continue to be used for various AI image generating models. Dall-e, Imagen, and Midjourney are some models that utilised customised GANs to create their models.

While GANs were the most relevant AI image

they had setbacks that made them difficult to train. One of these problems would be how the generator would only stick to a few outputs that have been able to fool the discriminator. And while GANs were dominating the AI industry as of late, Stable Diffusion made a breakthrough in creating realistic images that were both revolutionary and high quality. Diffusion models tried to use a different approach by deconstructing data structures through a forward diffusion process, then restoring these structures back to create a feasible generative model of data (Faroqui, 2022) Diffusion models also utilise an autoencoder in a small latent space, and decompresses the final latent representation and produces an image. As this process is repeated over and over in the diffusion process, diffusion models were seen to be more cost and time efficient.

🔥 SO HOW DOES STABLE DIFFUSION WORK?! 🔥

With Stable Diffusion getting more and more popular recently, it begs the question of how much goes into its abilities to create such high quality images. But how does it work exactly?

To explain the ways of AI generating using Stable Diffusion, it is easier to break it down into a few simple steps:
> 1. Text Tokenizing, where a text is broken down into a text tokens, word per word in order to be processed by the model
> 2. Image Refining, taking these text tokens to create an image
> 3. Image Upscaling, taking our finished image and upscaling it to become more high quality and refined
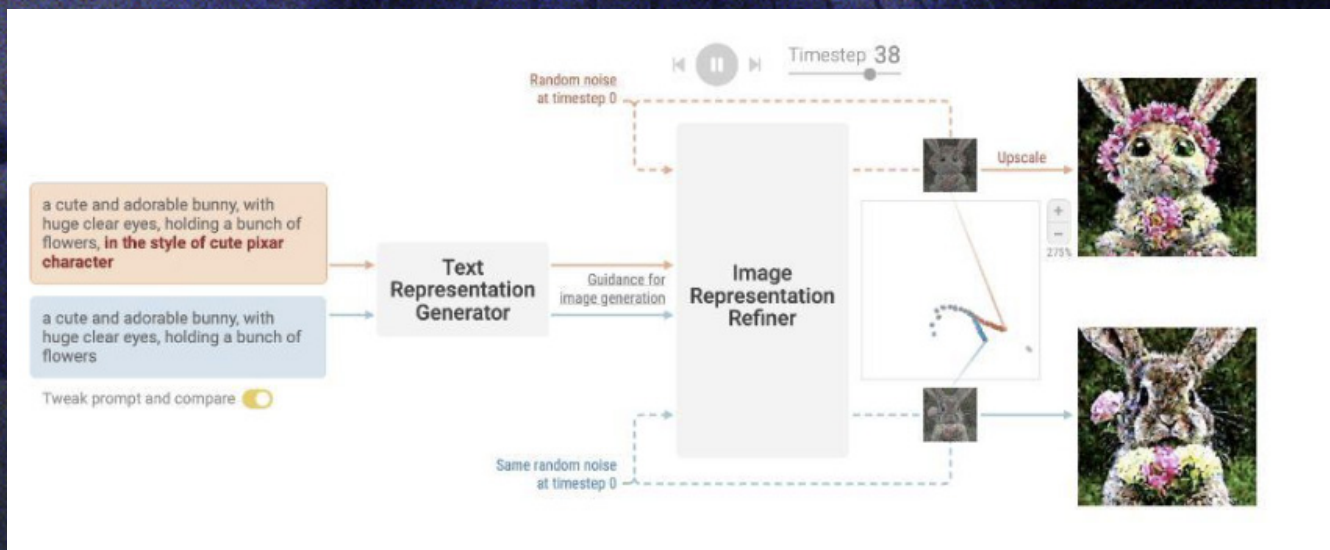
1. TEXT TOKENIZING

When Stable diffusion receives a text prompt, the first step it will do is convert your sentence into a text vector representation in order to make the generation process easier. This stage is called tokenizing. According to (Lee et al., 2023), tokenizing is a method to standardise a text prompt, and break it down into a series of tokens. Example, let's say we type in a phrase such as "A warm pepperoni pizza". What Stable Diffusion will do is take those inputs and split the words into "A", "warm", "pepperoni", "pizza". That way, the words will be easier to be processed by Stable Diffusion. These tokens will be encoded into a neural network called CLIP. In order to generate a specific image, CLIP will project the text prompt given into a text-and-image embedding space (Dehouche and Dehouche, 2023) It will then find a rough image that's closely related to the text prompt given. However, this image wouldn't closely resemble the text prompt just yet. Rather, it is an image made up of noise, and it must go through some steps before it creates a full blown image!

2. IMAGE REFINING

So now that we have a noise image that closely resembles our initial text prompt, Stable Diffusion uses a network architecture called U-Net. U-Net assists the AI image generating process by first downscaling the image and predicting how many timesteps it will take to denoise our noisy image and create an image representation of the current timestep, and then finally upscaling the image later on. First things first, U-Net takes three inputs. These three inputs would be an image representation of the current timestep (which would be our rough, noisy image), the text prompt that was given, and a timestep which indicates the amount of noise remaining in that current image representation. U-Net will decide how many steps it will take in order to make the noisy image go from a very unclear, grey-ish image to a high quality image that the model will spit out that will closely resemble our text prompt (Lee et al,.2023).

In addition to U-Net denoising the image itself, the text to image generation available on Stable Diffusion allows users to be able to input arguments such as sampling type, output image generation, and the seed value to create a certain type of image at a certain step of its denoising process. After calculating just how many steps it will take to turn our noisy image into a generated one that adheres to the text prompt, Stable Diffusion will go through the process of 'de-noising' (Dehouche and Dehouche, 2023). The process of de-noising is done by gradually removing the noise from the noisy image that we had initially per step, refining the image and gradually creating details that produce a much more detailed and complex image than what we had before (Lee et al., 2023). Every step gets subtracted after it's done, and after going through all the established steps, we now have somewhat of a refined image!

3. IMAGE UPSCALING

After the highly taxing process of image generation through de-noising, we now have a good image that resembles our text prompt. Stable Diffusion uses a neural network called Decoder to upscale the current image we had into a higher resolution image. This newer or refined version retains all the elements of its downscaled image thanks to the use of text representations, and the image should now be similar to the text prompt!

After going all through that, we now have an image that (hopefully) resembles what request we gave the model. However, if it doesn't quite resemble the model, there are a few things that could be done. One of which would be making more detailed descriptions in our text prompts to help the model decipher more and more specific details and generate a more detailed version of our request. That way, we will get more satisfying results!!

FLAWS IN AI IMAGE GENERATING???

Nothing can be without its flaws, and AI image generating is certainly one of them. Even with Stable Diffusion, there are still many challenges the industry is facing in order to make image generating more reliable and time efficient.

As of right now, models require a lot of GPUs for it to run as well as they do, and it contributes to incredible energy usage and produces a lot of heat, which isn't the best for the environment (Placido, 2023). Five to fifty questions can result up to gallons and gallons of water being used. The energy that a model takes to analyze the data it was trained on is too much, as revealed by recent environmental reports of AI sites' water usage and semiconductors (Ligaspi, 2023). If there is no further action, it is predicted that the demand for water may double over the next decade.

Another problem with AI image generating is the output itself. As of late, I bet a lot of people have seen some flaws if they've ever tried to make a model do some kind of prompt.

Mismatched fingers, misconstructed buildings, objects that don't make sense, or lighting that doesn't seem to follow the course of nature. There are so many problems. But why does that happen? Isn't AI supposed to be able to correlate text to images with accuracy?

Well first off, AI does not process things like us humans do. They are trained off the perception that they were given. This happens to be based on internet articles and media. This has the potential for AI's perception and understanding of certain phrases or words to be misinformed. When it comes to AI art, the model tries to take in the words given and find similar images based on the data it was trained off of. However, these models have a different perception of reality from humans depending on the data it was trained on. Otherwise, AI would be able to produce more accurate looking images. To solve this, so far we have to be extra specific with our requests and text prompts so AI can grasp the concepts that we are looking for, and the model will do its best to replicate.

🤖 HOW COULD YOU LEARN AI IMAGE GENERATING FIELD? KNOW THE BASICS HERE 🤖

1. Learn the basics!!

Before diving into the world of artificial intelligence, your life will be much easier if you have prerequisite knowledge first. Learn the fundamentals, such as a proficiency in Mathematics (it is essential) as well as proficiency in programming to learn the essential AI tools and packages needed for this particular field. You should also learn how to interpret statistics and be familiar with different datasets you may encounter.

2. Remember that the AI field is always changing

AI is a fast-growing field, meaning that any information will be up to date as soon as the newest breakthrough comes. Meaning, there will be a lot of learning, and learning, and learning! It is important to note how there are so many AI revelations these days, and one thing you can do is keep up! But don't worry. You got this

3. Just do it

While Deep learning and machine learning may look difficult, the only way to counteract it is to start learning and doing. While it may be challenging at the start, it will only get easier as you learn. As you read through paper by paper, your understanding will grow.

REFERENCES

Computerphile. (2022, October 4). *How AI image generators work (Stable Diffusion/DALL-E)*. YouTube. https://www.youtube.com/watch?v=1CIpzeNxIhU&t=543s

Gonkee. (2023, June 27). *How Stable Diffusion works (AI image generation)*. YouTube. https://www.youtube.com/watch?v=sFztPP9qPRc

Garcia, C. (2019, September 24). *Harold Cohen and AARON: A 40-year collaboration*. CHM. https://computerhistory.org/blog/harold-cohen-and-aaron-a-40-year-collaboration/

Khadatkar, R. (2024, February 20). *The Evolution of AI images: From pixels to perfection*. The Magic Studio Academy. https://magicstudio.com/blog/how-ai-images-have-evolved/#:~:text=The%20early%20AI%20image%20generation,laid%20the%20foundation%20for%20GANs

Lee, S. (2023, November 10). *Stable Diffusion explained and visualised for everyone*. Medium. https://medium.com/polo-club-of-data-science/stable-diffusion-explained-for-everyone-77b53f4f1c4

Legaspi, A. (2023, September 11). *AI technology guzzles water: Enough to fill 2,500 Olympic-sized pools*. Rolling Stone. https://www.rollingstone.com/culture/culture-news/ai-chatgpt-increased-water-consumption-environmental-reports-1234821679/

Nehme, A., & Crabtree, M. (2023, November 22). *How to learn AI from scratch in 2024: A complete expert guide*. DataCamp. https://www.datacamp.com/blog/how-to-learn-ai

Placido, D. D. (2024, May 9). *The problem with AI-generated art, explained*. Forbes. https://www.forbes.com/sites/danidiplacido/2023/12/30/ai-generated-art-was-a-mistake-and-heres-why/?sh=1fd853563fef

Poltronieri, F. A., & Hänska, M. (2019). *Technical images and visual art in the era of Artificial Intelligence*. Proceedings of the 9th International Conference on Digital and Interactive Arts. https://doi.org/10.1145/3359852.3359865

IBM. (2024, April 4). *What is unsupervised learning?*. IBM. https://www.ibm.com/topics/unsupervised-learning

AI SPORT COACH AND NEW GAME TACTICS

Significant differences from a fraction of a second or an inch of improvement, can distinguish the best athletes from the rest. With the eagerness for athletes to constantly seek ways to elevate their ability and performance, the integration of technology has become an apparent instrumental in this pursuit. Since the rise in popularity of AI, sports has revolutionised traditional training methods and offered innovative solutions to enhance athletic abilities as athletes today are undoubtedly stronger, fitter, and are also able to recover faster than their counterparts from decades ago.

Athlete Sports Tactics

Significant differences from a fraction of a second or an inch of improvement, can distinguish the best athletes from the rest. With the eagerness for athletes to constantly seek ways to elevate their ability and performance, the integration of technology has become an apparent instrumental in this pursuit. Since the rise in popularity of AI, sports has revolutionised traditional training methods and offered innovative solutions to enhance athletic abilities as athletes today are undoubtedly stronger, fitter, and are also able to recover faster than their counterparts from decades ago.

AI Coaching in sports represents a cutting-edge fusion of technology and sports science, revolutionising the way athletes train, perform, and reach their peak potential. At its core, AI Coaching harnesses the power of artificial intelligence to analyse vast amounts of data, pinpoint patterns, and provide personalised insights and recommendations tailored to individual athletes' needs.

A notable example of AI's impact on sports is the use of Statcast, an AI-powered tool in Major League Baseball. Statcast meticulously monitors and evaluates various aspects of player performance, such as movement patterns, bat speed, launch angles, and pitching mechanics (Forbytes, 2024). By leveraging this advanced technology

Source: Rocketmakers - Applications of AI in Sports

coaches and players can access detailed performance metrics, allowing them to identify areas for improvement and fine-tune their skills effectively. This data-driven approach enables athletes to optimise their performance and enhance their competitive edge on the field.

A notable example of AI's impact on sports is the use of Statcast, an AI-powered tool in Major League Baseball. Statcast meticulously monitors and evaluates various aspects of player performance, such as movement patterns, bat speed, launch angles, and pitching mechanics (Forbytes, 2024). By leveraging this advanced technology, coaches and players can access detailed performance metrics, allowing them to identify areas for improvement and fine-tune their skills effectively. This data-driven approach enables athletes to optimise their performance and enhance their competitive edge on the field.

Source: MLB - Analytics on baseball player

Furthermore, aspiring athletes can have an AI Coach through Tonal; a smart fitness gym that integrates AI technology with electromechanical weights. Tonal's AI-powered system offers a range of cutting-edge features designed to enhance athletes' performance and facilitate their fitness journey. These features include:

- **Strength Assessment:** Upon analysing and assessing the athlete's strength levels and their individual capacity, Tonal recommends personalised workout routines and weight adjustments to optimise their performance.
- **Adaptive Weights:** Using athlete's real-time performance, Tonal's AI ensures that it will dynamically adjust its resistance levels suitable to the athlete. This ensures that their
- **Full-Body Workouts:** Tonal provides a comprehensive range of exercises targeting different muscle groups, seamlessly integrated into personalised workout programs tailored to the specific goals and needs of athletes.

Injury Preventions

Players need to ensure that their health is not compensated wherever possible to gain a competitive advantage amongst their revival team. This involves staying healthy, recovering quickly, and optimising performance. However, sports like hockey, soccer, and basketball are particularly injury-prone. According to the National Safety Council, there was a 32% increase in sports and recreational injuries from 2021 to 2022 (NSC Injury Facts, 2023).

The integration of advanced technological tools, such as biomechanical analysis and wearable technology, plays a crucial role in monitoring athletes' health and preventing injuries. Biomechanical analysis helps identify and correct movement patterns that could lead to injury, while wearable technology provides real-time data on an athlete's physical condition. This data includes metrics such as heart rate, muscle strain, and fatigue levels, offering a comprehensive view of an athlete's health.

Source: MLB - Analytics on baseball player

By capturing and analysing this data, coaches and medical professionals can better understand an athlete's overall health and ensure that their diet and training regimens are optimised for peak performance. Real-time data during training sessions can pinpoint when an athlete is approaching their physical limit, allowing for timely adjustments to prevent overtraining and injuries (Catapult, 2024). This data-driven approach not only enhances performance but also extends the athlete's career by reducing the risk of long-term injuries.

Furthermore, utilising sensors and other monitoring technologies enables personalised training programs tailored to the individual needs of each athlete. This customisation ensures that training intensity and recovery periods are aligned with the athlete's physical condition, further minimising injury risks. In essence, the adoption of these advanced technologies represents a proactive approach to athlete health and performance, fostering a safer and more effective training environment (Safety and Health Magazine, 2023).

Ethical Issues

However, despite all, this technological progression also raises questions about fairness and the boundaries of competition as technology in sports. Has technology interfered with what sport used to be?

During a game, AI-powered systems can provide instant feedback not only on the team players themselves but also capture their opponents as well. This ultimately gives coaches the opportunity to make data-driven decisions allowing the ability to adjust tactics in real-time. This is because coaches and analysts now have access to sophisticated tools that can analyse opponents' game footage in great detail, identifying patterns, weaknesses, and strategic tendencies (Folio3, n.d). This dynamic adjustment capability is crucial in the fast-paced nature of sports, where the ability to adapt quickly can be the difference between winning and losing. For example, machine learning algorithms can predict the likelihood of certain plays based on historical data, helping teams prepare for various scenarios. This predictive capability allows teams to be proactive rather than reactive, giving them a competitive edge.

Source: TechXplore - Football player performance

Hence, the ethical issue stems from the potential for AI to create an uneven playing field. Teams with greater financial resources can afford to invest in the most advanced AI systems, potentially widening the gap between well-funded teams and those with fewer resources. This disparity could lead to a competitive imbalance, where success is increasingly determined by technological investment rather than athletic skill and strategy.

Conclusion

In summary, the application of AI in sports is not just enhancing performance on the field but also transforming the way teams strategise and make decisions. By leveraging advanced technology, teams can develop more effective game plans, improve their performance, adapt to changing conditions in real-time, and maintain a competitive edge in the ever-evolving sports landscape. As Rafael Nadal noted, technology is indeed changing sports as we know it, ushering in a new era of strategic innovation and competitive game plan.

CREATING THE WORST GAME EVER

An exploration of dark patterns and design technique by Bonsen Wakjira

Sweat cascades down your neck in a torrential downpour, each drop a testament to your unrelenting effort. You hold your gear close and grip your weapon with the weight of a thousand souls, every ounce of your being focused on this final battle. Although your heart pounds in rhythm with your enemy's, you enter the fight with confidence. You've rehearsed this battle in your mind a thousand times, felt defeat more times than you can remember, but you see an opening, a chance for victory. You've honed your agility and skill for this moment, and as your health drops to the brink of existence, you take your final leap… but in a cruel twist of fate, you miss. Just as the weight of failure sinks in, an ad SHOUTS across your screen: "UNLOCK EXCLUSIVE GEAR WITH THE NEW LEGENDARY BUNDLE! ONLY $19.99!"

Welcome to the dark side of the moon, where games don't just test your reflexes and strategies but also your patience and financial limits. Games like Call of Duty, notorious for their loot boxes and microtransactions, have turned the art of gaming into a business model that profits from player frustration. Every death or setback is another opportunity to push you toward yet another in-game purchase.

You might be thinking of challenging games like Elden Ring or I Wanna Be the Boshy, but they differ in that their difficulty is designed to be fulfilling—an achievement of some sort, not an exploitation. In fact, these games are loved for their tough but rewarding experiences, satisfying a deeper psychological need for struggle and accomplishment (Smith, 2022).

I know I love them, even though I've probably died over 100 times in my first playthrough.

A message I've been seeing… a lot recently - Elden Ring

The **"worst game ever"** isn't about challenge for the sake of challenge; it's entertainment crafted with a "profit-first" mentality. These games are engineered to frustrate:
- Ads popping up every time you fail
- Daily login streaks pushing you to play even when you don't want to
- Limited-time events with rewards that capitalise on your Fear Of Missing Out (FOMO)
- Every reward giving you a massive dopamine boost, which is quickly taken away as you hit another paywall

The **"worst game ever,"** simply put, is one that uses **dark design patterns**.

It can't have always been like this, right?

Well, yes… and no.

The first video games appeared in the 1950s, when universities participated in research projects that had no purpose of entertainment or profit, but served more as academic and promotional devices.

"The games of the 1950s were largely confined to research labs… and were not intended for mass distribution and/or public consumption" (videogamehistorian, 2014).

"Tennis For Two", The First Game Ever Made
(videogamehistorian, 2014)

But this is the nature of evolving technology; it always starts academically with thousands of benefits, but under all those green flags, there are some red flags that always show up at inconvenient times.

Focusing on the advent of arcade games like Pong and Space Invaders, often dubbed as pioneers in gaming (Kaufmann, 2022), we start seeing profit being made from players. The business plan was simple:

1. Invest in an arcade machine,
2. Put it anywhere near a school,
3. Get players hooked so they keep feeding the machines those shiny coins.

This is relatively tame, however. Yes, we can argue that difficulty spikes were carefully designed to make you pay to play more, but there are worse things to worry about.

Like the notorious pinball machines—the "game of chance"—often considered a gateway to gambling (The History & Evolution of Arcade Gaming | M&P Amusement, 2021). They were designed to entice, frustrate, and make money, much like their modern-day casino counterparts. Back then, it wasn't as clear who was winning; you got to enjoy games that weren't available elsewhere, while the companies made a bit of profit from your spare change. It was a win-win for everyone.

But then, BOOM!! A new competitor entered the market:

"It's-a-me, Mario!"

The 1980s and 1990s shifted the gaming landscape with iconic titles like Super Mario Bros and The Legend of Zelda (Wallach, 2022), along with home consoles—no more shoving coins into a slot—just a one-time purchase for endless gameplay and full ownership.

WAIT A MINUTE, THIS IS A GOOD THING!

A straightforward transaction: buy a game, own the game, play and enjoy the game, with no further financial obligations.

Then, there's the turning point. As the gaming industry evolved, so did its revenue models:
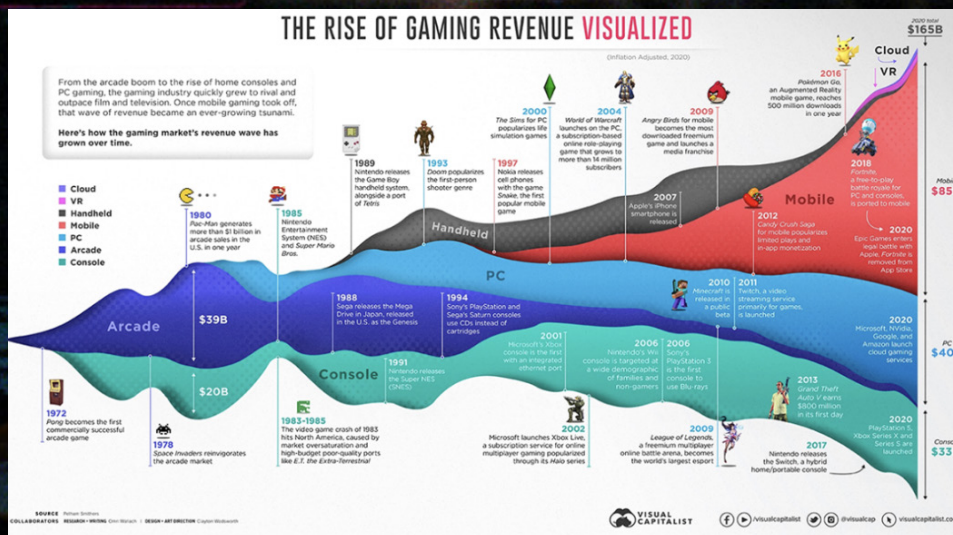
50 Years of Gaming History, By Revenue Stream (Wallach, 2020)

There's a pattern in game design: every year is the same story—record sales, an expanding market, and then CRASH—a new competitor enters.

The allure of quick profit began to overshadow the entertainment industry. With the cost of

development infinitely expanding, the console market crash, and the introduction of PCs, free-to-play models, and handheld devices, companies sought new revenue streams.

Enter 2024, where we have microtransactions, loot boxes, endless DLCs, fake ads, and so much more—all contributing to a $165 billion profit for the gaming industry (Wallach, 2020).

Don't get me wrong; I appreciate the simpler days of gaming and often reminisce about the first games I played. But it's important to acknowledge how the pursuit of profit has led to the proliferation of **dark patterns** today.

SLOW DOWN!! WHAT ARE DARK PATTERNS?

I've mentioned this term a few times already, but I haven't unpacked it yet—right now, all it does is dramatise this article.

Formally, **dark patterns**, a term coined by user experience designer Harry Brignull, refer to manipulative design practices crafted to deceive users into actions they might not otherwise take (Brignull, 2011). More generally, dark patterns are any UI/UX designs intended to trick people into making decisions that benefit the business, often at the user's expense. These patterns walk a fine line between legality and exploitation.

> "Dark patterns exploit psychological insights to manipulate behaviour, creating a façade of user-friendliness while masking underlying exploitation" (Brignull, 2011).

Dark patterns aren't limited to games. In fact, a popular example, "Privacy Zuckering," is named after Mark Zuckerberg. This design technique, often employed on Facebook, tricks users into sharing more information than they intended (Mohit, 2017). This leads users to give up more information unknowingly, often making it purposefully difficult to opt out.

There are hundreds of patterns, and not all harmful design techniques can be simplified into a single pattern. In most cases, software uses a combination of patterns to exploit users.

A list of common patterns can be found at Deceptive Design.
To really explore these patterns, we'll consider three separate case studies, each revealing how design can be weaponized against users to prioritise profit above all.

"MOM! CAN I HAVE SOME VBUCKS???" - EPIC GAMES VS.

Fortnite. A genre-defining game that absolutely revolutionised the gaming scene, redefining what a battle royale meant. Not only was it fun, addictive, and enjoyable, but it also introduced a battle

royale meant. Not only was it fun, addictive, and enjoyable, but it also introduced a battle pass system widely endorsed by game developers. You bought it once, and as long as you didn't overspend your in-game currency, you could keep buying and buying, effectively amassing thousands of in-game

It's surprising (not really) to hear that such a popular game was fined $720 million for using deceptive patterns.

Epic Games, the publisher behind Fortnite, found itself in trouble with the Federal Trade Commission (FTC), accused of employing deceptive patterns that led to unintended purchases and unauthorised charges. The FTC's complaint against Epic Games was based on allegations that Epic had purposefully created obstacles to prevent users from cancelling charges or requesting refunds, especially for in-game items like Battle Passes, Llamas, and Cosmetics (Deceptive Patterns - Legal Cases - in the Matter of Epic Games, Inc., n.d.).

What dark pattern is this?

Epic's refund process was convoluted and involved unnecessary confirmations, whereas there was no confirmation or verification when purchasing in-game items.

And it gets worse… Recent research into monetized video games by King et al. has linked addiction-like behaviours to compulsive spending. A study involving 428 Fortnite players revealed that those who spent the most on microtransactions were likely experiencing symptoms akin to gambling disorder. The study found that social influences—such as seeing a close friend frequently spend money on Fortnite—played a significant role in encouraging spending. Higher spenders also reported feeling a stronger motivation to chase in-game rewards and perceived these items as better value for money. However, this relentless pursuit of digital goods was also linked with problematic gaming behaviours, such as impulsivity and a diminished sense of self-worth when attempting to cut back on playtime.

While there is no single design pattern directly linked to these issues, the implications suggest that monetisation schemes are not only squeezing money from players but also encouraging addictive behaviour.

- **Fake Urgency:** Limited-time events are designed to create a sense of urgency, pushing players to spend money to maximise their progress.
- **Visual Interference:** Ads are frequently integrated into gameplay, disrupting the experience and often prompting further spending.
- **Disguised Ads:** Some ads are so interactive that they blend seamlessly with gameplay, leading players to unknowingly participate in promotional content.

The allure of these deceptive ads is partly due to their design, which exploits psychological triggers like FOMO (fear of missing out) and instant gratification. Advertisers craft these ads to be visually compelling and frustratingly engaging, often featuring someone failing at what appears to be a simple task. This tactic not only frustrates viewers but also entices them to download the game to see if they can succeed where others failed (AppSamurai, 2024).

So, is this kind of false advertising illegal? Technically, no. Evony exploits a loophole by including various mini-games that more closely resemble the ads they produce. These mini-games are easy to showcase and don't necessarily break any laws, allowing advertisers to sidestep regulations while still misleading potential players.

"MONKEY EATS BANANA" - BANANA NFTs??

Have you heard of "Banana"? Maybe you've heard of the prequel "Apple"? The sequel "Melon"? No? Well, I know you've heard of the crossover "Fruits".

Who said AAA game companies were the only ones allowed to use dark patterns? Indie companies have just as much right to employ them! Banana is a clicker game that appeared out of nowhere and took over the Steam market. It uses a technique I'd describe as a "Crypto Ponzi Scheme," where players collect different types of bananas that can be traded on the Steam Marketplace for real money. Common bananas go for as little as 3 cents each, with developers (and Valve) pocketing two-thirds of that amount.

If you thought NFTs were finally over, then you're dead wrong.

The concept might sound enticing—you play the game and earn a little cash, it's just like a job! In reality, however, this is just a Ponzi Scheme gamified, benefiting only the early adopters (and developers) while leaving the latecomers out in the cold.

At its peak, Banana hosted 700,000 users, most of whom were not real players but crypto bots

designed to grind (Carpenter, 2024).

Now, there isn't a specific dark pattern to describe this game design (which is why I dubbed it a "Crypto Ponzi Scheme"), but it does fit the definition of dark patterns as it tricks users into benefiting the business. So, how does this game actually profit the creators? With a player base of 700,000 and trades as low as $0.03 and as high as $1,378 (Carpenter, 2024), the developers get a portion of each transaction. Even a tiny percentage of profit per banana adds up. For example, if each banana trade nets the developers 1 cent, and there are 700,000 trades, that's $7,000.

On a side note, you'd think Valve would do something about this and ban all types of "fruit clickers" on Steam. But let's not forget that Valve also takes a percentage of earnings, so they're not completely innocent in this.

Even though these games don't involve blockchain technology and aren't NFTs in the traditional sense, their design mimics some of the same exploitative tactics. By creating an illusion of endless potential earnings, they attract players into a system where the true financial benefits are heavily skewed in favour of a few.

EITHER DIE A HERO OR LIVE LONG ENOUGH TO BECOME A VILLAIN

Isn't it every gamer's dream to design their ideal game, jump on a game engine, and start coding their AI-based open-world FPS roguelike MMORPG? Oh wait, that actually requires time and money?

No worries! Let's craft a profit-maximising, user-exploiting game together that's not only easy to develop but will make us RICH!

STEP 1: NO MONEY, NO PROBLEMS

Why spend time, energy, and money creating your own 3D models, game art, or in-game assets? Instead, grab some premade assets, or if you don't want to be an asset flipper, use generative AI to create assets for you.

Your game might not look the best, but as long as the cover art is engaging, who cares?

STEP 2: IMMEDIATE PROFIT

Don't have money to market your game? No problem! Combine your premade assets into a quick showcase and launch a Kickstarter campaign, pitching a game so revolutionary that they MUST

give you their money.

For an added bonus, hire someone to create a basic game foundation. The only important functionalities are:
- Player moves left and right
- Player can get upgrades
- We can place ads anywhere

STEP 3: LOOT BOXES & PAYWALLS

Let's make the first 5 levels free, because we are so generous. Afterwards, the difficulty spikes so hard that players either have to revive by watching countless ads or pay for in-game "cosmetics" that totally do not affect the game in any way. Totally.

Also, make players buy loot boxes to unlock "cosmetics." Of course, we won't show the true chances of getting legendary items; everything will be rigged!

STEP 4: COSMETICS WITHOUT MULTIPLAYER?

Who wants to buy cosmetics in a game without multiplayer servers? Not me! But hosting a server costs time and money, which we don't have. So, let's hire another developer to create fake bots that just walk around and say random stuff in the lobby. All these bots will have cool cosmetics that make actual users jealous.

STEP 5: PONZI SCHEME

The Steam market seems too saturated with fruit-based "pay-to-earn" games. Instead, include a mechanic where players can buy virtual land that generates real-world profits. Just make sure the land is expensive and change the mechanic whenever you start losing money.
Or you could just make another fruit-based "pay-to-earn" clicker game, I don't think "Kumquat" is taken

Consider the controversial example of Twerk Race 3D — Fun Run Game.

STEP 6: MORE… MORE!!

More ads, more in-game currency, more "cosmetics," more abuse, more profit—MORE!

Congratulations! You've just designed the perfect game. Make sure to give me a percentage of the profit. Actually, never mind—you would never do that.

REFERENCES

Brignull, H. (2011, November 1). Dark Patterns: Deception vs. Honesty in UI Design. A List Apart. https://alistapart.com/article/dark-patterns-deception-vs.-honesty-in-ui-design/

Carpenter, N. (2024, June 7). Why clicker game Banana — the "legal infinite money glitch" — is going viral on Steam. Polygon; Polygon. https://www.polygon.com/gaming/24173617/banana-clicker-game-steam-egg

Deceptive Patterns - Legal Cases - In the Matter of Epic Games, Inc. (n.d.). Www.deceptive.design. https://www.deceptive.design/cases/in-the-matter-of-epic-games-inc

Fake Mobile Game Ads - A Closer Look Why Advertisers Use It? | AppSamurai. (2024, February 27). https://appsamurai.com/blog/fake-mobile-game-ads/

Htwo. (2024, June 7). The Insidious World of Fake Mobile Game Ads. YouTube. https://www.youtube.com/watch?v=NhajAqI66nU

Kaufman, A. (2022, September 15). What was the first video game? Detailing the first at-home video game, who invented it and more. USA TODAY. https://www.usatoday.com/story/tech/2022/09/15/what-was-the-first-oldest-video-game/8021599001/

King, D. L., Russell, A. M. T., Delfabbro, P. H., & Polisena, D. (2020). Fortnite microtransaction spending was associated with peers' purchasing behaviors but not gaming disorder symptoms. Addictive Behaviors, 104(1), 106311. https://doi.org/10.1016/j.addbeh.2020.106311

Mohit. (2017, April 10). Privacy Zuckering: Deceiving your privacy by Design. Medium. https://medium.com/@mohityadav0493/privacy-zuckering-deceiving-your-privacy-by-design-d41b6263b564

Murphy, S. (2022, March 2). The App Store's Adware Gaming Plague - Sam Murphy - Medium. Medium; Medium. https://medium.com/@stmurphywriter/the-app-stores-adware-gaming-plague-691eeb448f25

Prema, S. (2024, February 8). What is "brain rot"? Do you Have it? SBS News. https://www.sbs.com.au/news/article/what-is-brain-rot-do-you-have-it/39fexbr4u

Smith, E. (2022, April). Why We Love Hard Games Like Elden Ring And Dark Souls. GAMINGbible. https://www.gamingbible.com/features/why-we-love-hard-games-like-elden-ring-and-dark-souls-20220401

Smithsonian. (2019). Video Game History. Smithsonian Institution; Smithsonian Institution.
https://www.si.edu/spotlight/the-father-of-the-video-game-the-ralph-baer-prototypes-and-electronic-games/video-game-history

The History & Evolution of Arcade Gaming | M&P Amusement. (2021). M&P Amusement.
https://mpamusement.com/pages/the-evolution-of-arcade-gaming?srsltid=AfmBOop MZXvSO8mndwkQHFthrXjCBo8A5IxLC-KSQuj4-dPIHNGjsIy2

videogamehistorian. (2014, January 23). The Priesthood at Play: Computer Games in the 1950s.
They Create Worlds. https://videogamehistorian.wordpress.com/2014/01/22/the-priest hood-at-play-computer-games-in-the-1950s/

Wallach, O. (2020, November 23). 50 Years of Gaming History, by Revenue Stream (1970-2020).
Visual Capitalist. https://www.visualcapitalist.com/50-years-gaming-history-reve nue-stream/