



CODEx OF THE
FUTURE
SERIES

THE FUTURE OF **GAMING** 2020 TO 2070

UNLIMITED THINKING . EXPONENTIAL POTENTIAL

BY MATTHEW GRIFFIN

**“IF WE’RE LIVING IN A SIMULATION TODAY
THEN THE CREATORS SHOULD BE FIRED. ”**

- Matthew Griffin, Founder and CEO, 311 Institute

EXPLORE MORE

Looking for even more
insights into the future
and deep future?



Explore More

[BLOG](#) • [BOOKS](#) • [EDUCATION](#) • [KEYNOTES](#) • [NEWS](#) • [VIDEOS](#)

www.311Institute.com





UNLEASH YOUR EXPONENTIAL POTENTIAL

Free Resources

 [Explore More](#)

EDUCATION • ENTREPRENEURSHIP • INNOVATION

www.311Institute.com

CODEX OF THE FUTURE SERIES
EXPLORE MORE



EXPONENTIAL TECHNOLOGY CODEX.

Explore the hundreds of exponential technologies that are emerging in detail and learn about their implications for global culture, industry, and society.



HOW TO BUILD EXPONENTIAL ENTERPRISES.

Exponential technologies have accelerated the global rate of disruption, now it's your turn to disrupt the quo.



THE FUTURE OF EXPONENTIAL DISRUPTION.

Explore the forces accelerating global disruption and how close we are to disrupting industries in just a day.

CODEX OF THE FUTURE SERIES

EXPLORE MORE



THE FUTURE OF EDUCATION AND TRAINING.

Explore how we prepare our children and society for a future where Science Fiction becomes Science Fact.



THE FUTURE OF GAMING.

Explore the future of gaming and what happens when the simulations become reality.



THE FUTURE OF INNOVATION AND CREATIVITY.

Explore what happens when creative machines and not humans are the dominant creative force in the world.

CODEX OF THE FUTURE SERIES

EXPLORE MORE



THE FUTURE OF INSURANCE.

Explore the future of insurance, and the dangers of a future where global risk becomes systemic.



THE FUTURE OF SMARTPHONES.

Explore the future of smartphones and smartphone formats, and discover what's around the corner.

ABOUT THE AUTHOR



Matthew Griffin, an award winning futurist and author of *Codex of the Future* is described as “The Adviser behind the Advisers” and a “Young Kurzweil.” Matthew is the founder and CEO of the World Futures Forum and the 311 Institute, a global Futures and Deep Futures consultancy working between the dates of 2020 to 2070.

Regularly featured in the global media, including AP, BBC, CNBC, Discovery, RT, and Viacom, Matthew’s ability to identify, track, and explain the impacts of hundreds of revolutionary emerging technologies on global culture, industry, and society, is unparalleled.

Recognised for the past six years as one of the world’s foremost futurists, innovation and strategy experts Matthew is an international speaker who helps governments, investors, multi-nationals and regulators around the world envision, build and lead an inclusive, sustainable future.

BE BOLD. MAKE FIRST CONTACT.

mgriffin@311institute.com
+44 (0) 7957 456194

www.311Institute.com



AIRBUS



BCG
THE BOSTON CONSULTING GROUP

centrica

DELL EMC

Deloitte.

DENTONS

Discovery
CHANNEL

Deutsche Bank



Goldman
Sachs

Google
LUNAR X PRIZE

HM Government



IBM

J.P.Morgan

KPMG



LLOYDS
BANKING
GROUP

McKinsey & Company



pwc

QUALCOMM



sky

sopra steria



UBS

Willis
Towers
Watson

& ...

CONTENTS

Third Edition published Apr 2020. To request this Codex in an alternative language please contact the author.

12	...	INTRODUCTION
14	...	DECODING THE EXPONENTIAL FUTURE
18	...	DECODING EXPONENTIAL DISRUPTION
26	...	BUILDING EXPONENTIAL ENTERPRISES

Learn how to build your own disruptive Exponential Enterprise.

36	...	MEGATRENDS AND STARBURSTS
----	-----	---------------------------

Get a birdseye view of the latest Megatrends and Exponential Technologies re-shaping our world.

44	...	EXPONENTIAL IMPACT
50	...	THE FUTURE OF GAMING

Explore what happens when the simulations become reality.

62	...	THE FUTURE OF GAME DISTRIBUTION
80	...	FUTURE ARENAS AND EXPERIENCES

94	...	THE FUTURE OF GAME PRODUCTION
144	...	THE FUTURE OF GAME CONSUMPTION
146	...	THE FUTURE GAME EXPERIENCE
152	...	FUTURE MARKET SIZING AND PLAYERS
158	...	CONCLUSION





INTRODUCTION

If I posed the idea that we were all living in a simulation in the fifteenth century then I might have been burned at the stake. If I'd posed the idea in the early nineteenth century then I'd have likely been labelled a crackpot, and if I'd posed the idea in the 1960's I'd have been labelled a science fiction writer.

Matthew Griffin.

But if I posed the idea today, as a number of billionaire polymaths have recently, it would be discussed and debated by global audiences, and investigated by leading scientists. And that, right there, is a stark example of just how far we have come technologically and as a society in the comparative blink of an geological eye - insane and surreal sounding concepts are no longer as insane or surreal.

Irrespective of where you sit on the subject though today, right here and right now, I can show you the technologies that will make this a reality, albeit a simulated one, and furthermore as we race towards this future from a gaming perspective at least we'll be able to use these same technologies to create the ultimate unlimited gaming environment, of course, that is provided we aren't already in one,

eXplore more,



DECODING THE EXPONENTIAL FUTURE

WITH ACCESS to the right breadth and depth of insights, or dots, as I call them, putting the big picture together and predicting what the future will look like, and, perhaps more importantly, when it's going to arrive, isn't as hard as many people think.

After all, as they say, the future is hidden in plain sight, and sometimes it's just a simple matter of expanding your horizons.

In order to predict the future as accurately as is practically possible I do my best to work with what I call full network insights. That is to say I work with the inventors, academics, entrepreneurs, investors, multi-nationals, governments and regulators who, in one way or another, are all discovering, combining, building, testing, adopting, deploying, scaling and regulating tomorrow's technologies, products and services, or as I'll call them from here on in, for simplicity's sake, "Concepts."

It's this rich tapestry of contacts, that cuts across every geography and industry, combined with a deep understanding of hundreds of exponential technologies and their impact on culture, industry and society, that allows me to piece together the jigsaw that is the future with a high

degree of accuracy and precision. But it's no easy feat.

However, while technology plays, needless to say, a very important role in helping shape the future, viewing it in isolation in order to try to decode the future is a mistake, because, inevitably the success of these new concepts, and their tipping points, depend on a variety of interconnected factors that include, but are not limited to, accessibility, affordability, cultural and generational bias, investment patterns, maturity, and the regulatory, macro-economic and socio-political environment.

Should any of these fail to align correctly then tomorrow's Earth shaking new concepts can quickly turn into elegant failures whose potential is never realised.

With so many different emerging technologies on the horizon it's inevitable that some of them will compliment each other, and some won't, and that some will be more impactful than others.

Furthermore, when these new technologies do finally emerge from the labs then it's down to you and I, and increasingly our capable synthetic counterparts, the Creative Machines, to combine them into tomorrow's must have concepts.

Notes:

One of the greatest issues for analysts, futurists and industry watchers alike however is the fact that all these dots can be combined in billions of new, unique and exciting ways to create a limitless number of concepts, and trying to pick the winners, and see through the noise can be challenging, and as the volume of new technologies, and moving parts, increases the task is only going to get more complicated.

Personally, and it's more through experience than by design, I've found that the best way to cut through this noise is to divide the universe in twain. On the one hand we have the promising, individual emerging technologies, and on the other we have the new concepts they could be used to create.

Evaluating the technologies comes first though, because unless a specific technology can be bought to market, in the right way, and at the right time, then it follows on that, generally, it will never get the opportunity to be used to create a new concept - let alone a new mass market one.

Then, once we've filtered those it's a fairly straight forward process of ideating all of the different ways in which they can be combined to create new concepts, which, in turn can be evaluated on

their own merits and used to decode future trends and impacts, threats and opportunities.

As you'll see from this codex I've tried to make it easy for you, as easy as it can be under the circumstances, to quickly evaluate the merits and status of each of the exponential technologies I track, after which you should then be able to categorise the ones that you feel are the most relevant to you, or your organisation, so you can begin ideating your own concepts.

As the pace of change continues to accelerate, as the boundaries between industries continues to erode, and as science fiction increasingly becomes science fact, the future will belong to those individuals and organisations that have the foresight to see change coming, and who are agile and strong enough to adapt to it and lead it.

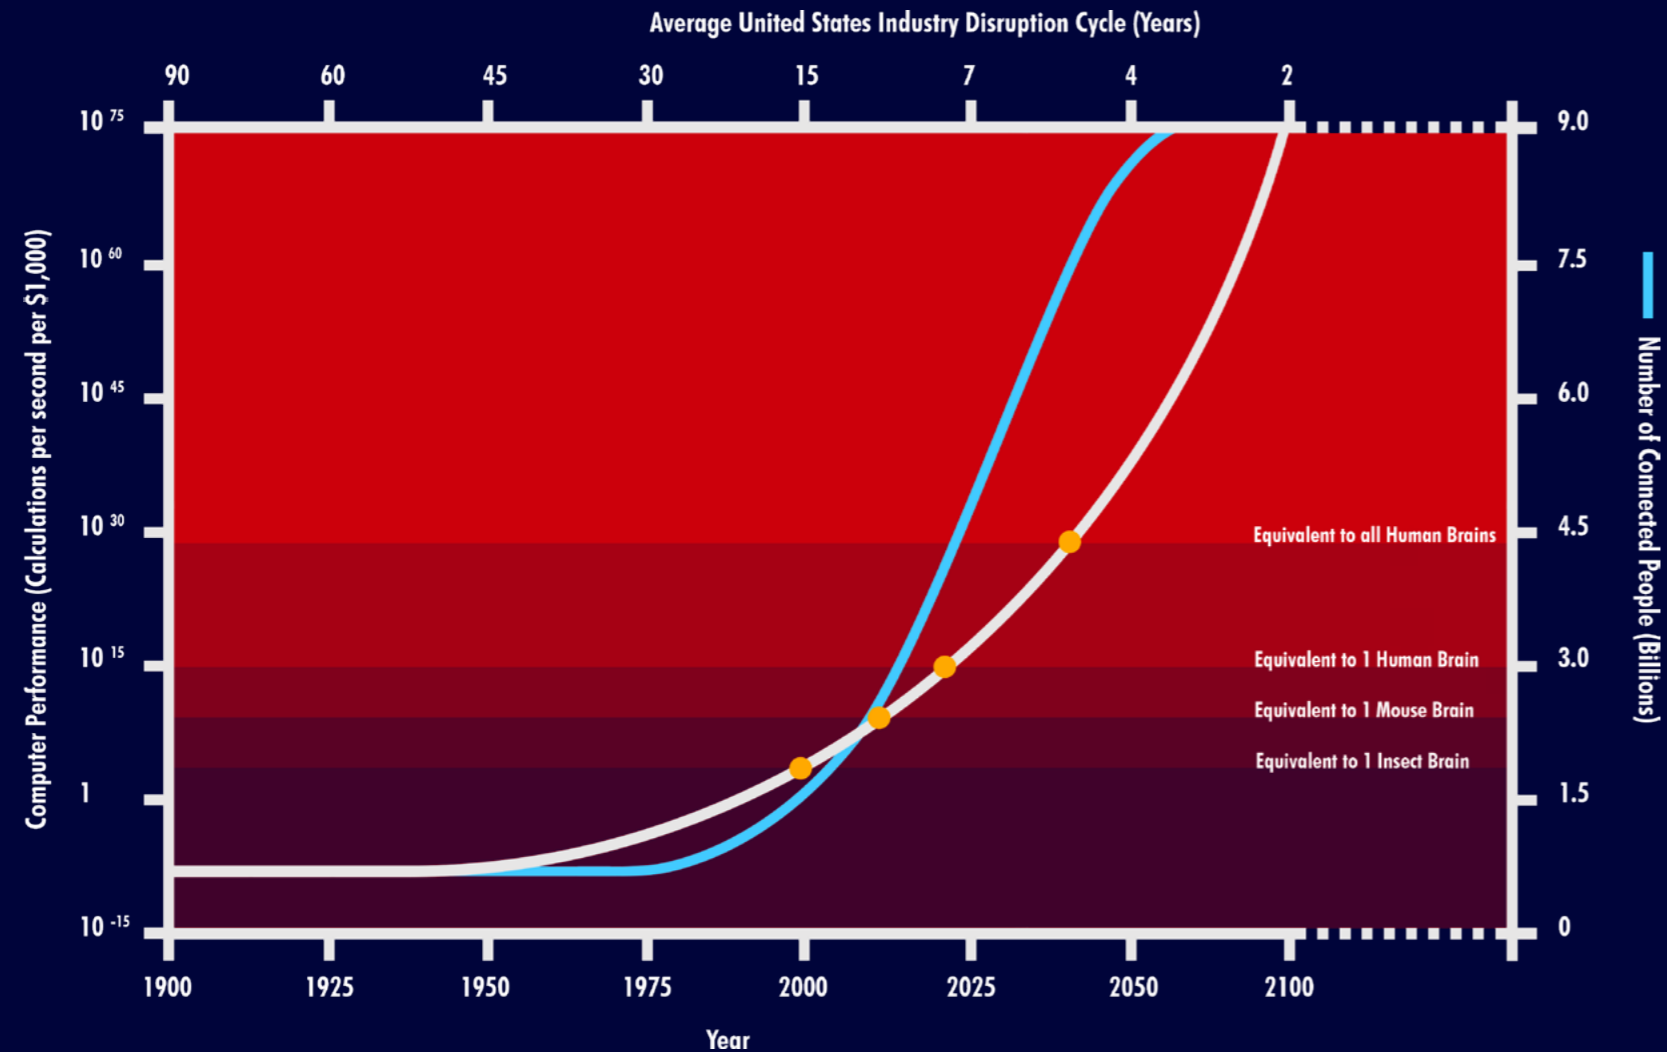
Notes:

DECODING EXPONENTIAL DISRUPTION

IF YOU step back a decade or so ago the word on everyone's lips was innovation and, frankly, if you didn't have it thrust into your face at least thirty times a day by every executive or ad man or woman you met then it's likely because you were in a coma. Or dead. Or both.

Fast forward to today and now they have a new buzz word - Disruption. But is disruption today as commonplace and accelerating as quickly as people will have us believe, or is it just hype and a word that executives and eager Silicon Valley startups throw around with impunity in the vain hope of convincing people that they're innovating at the bleeding edge and pushing boundaries?

Well my friend, let's take a journey together. Let's cut through the marketing fog, summit the hype cycle, and crack open an genetically modified beer while we raise cynical eyebrows and take a deeper look at the world that's unfurling around us.



TECHNOLOGY ENABLED DISRUPTION

As increasingly powerful exponential technologies emerge and are democratized, with computing power being just one example, and as the world becomes increasingly digital and connected industry disruption times plummet.

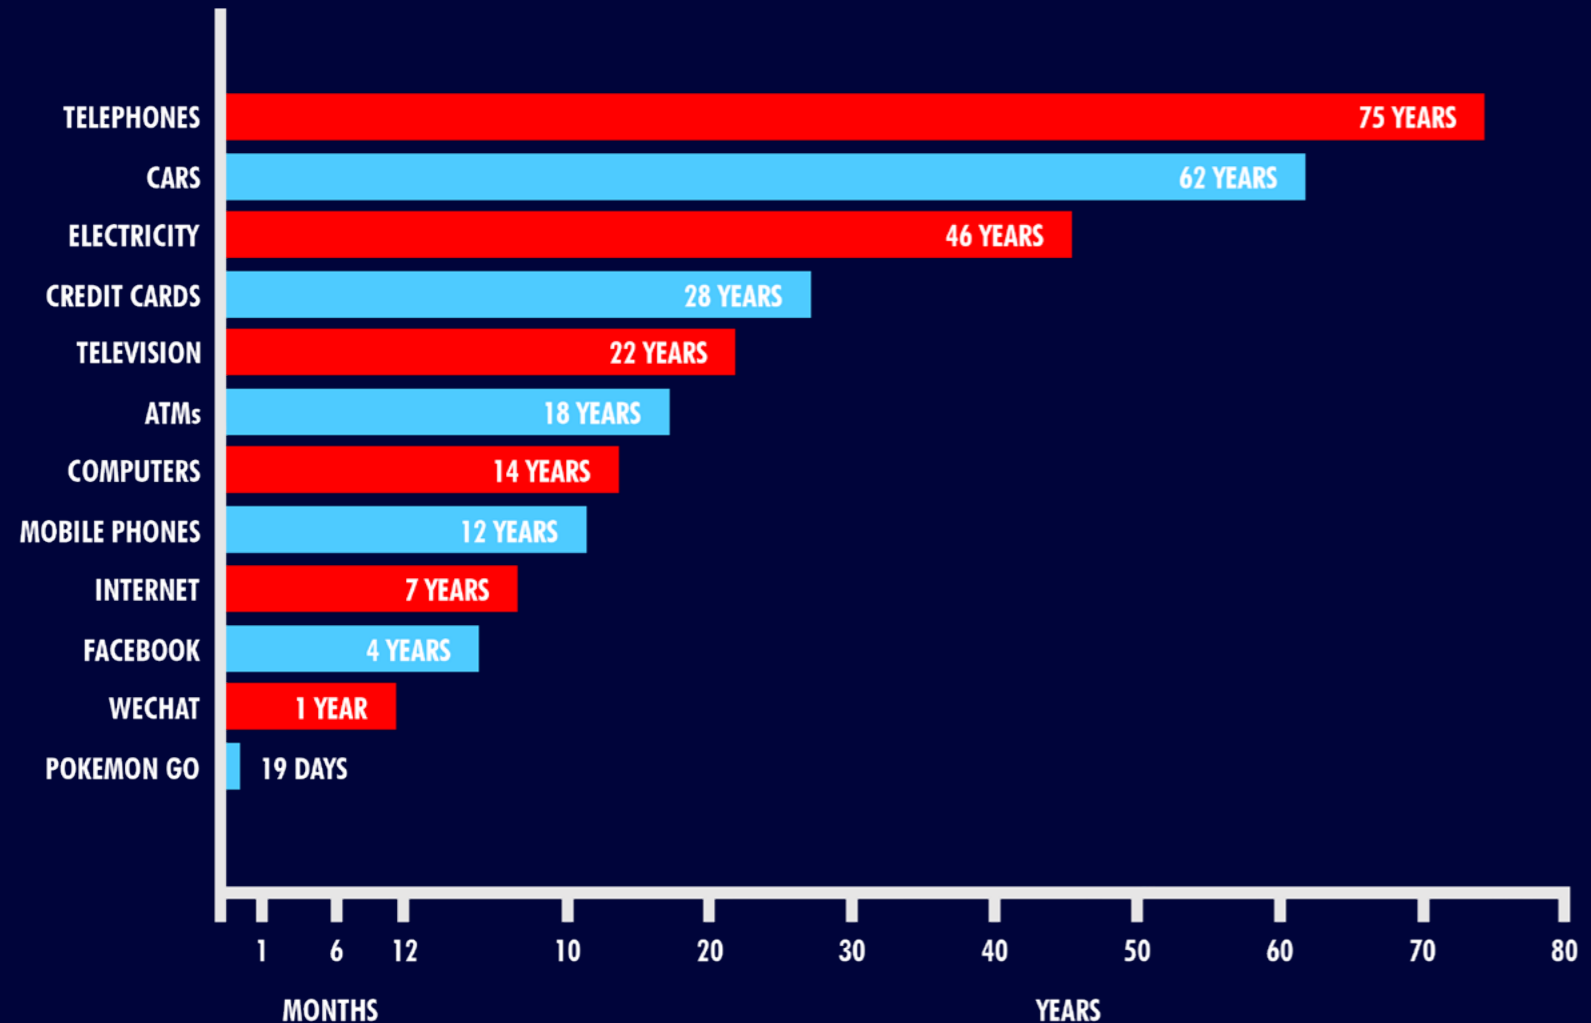
ACCELERATING DISRUPTION.

THE CORRELATION is obvious, but it's worth discussing nevertheless. If you want to disrupt the status quo it's not just good enough to have the ideas, tools, and resources that you need to create and develop your innovative new product, you also need to be able to get it into the hands of as many customers as possible as fast as possible. And in the distant past when products were physical and the only markets that entrepreneurs had easy access to were local ones, trying to disrupt anything at scale was not only an immense challenge, but it also took an inordinately long time - in many cases a lifetime or more.

Today, however, powerful new technologies and an increasingly connected planet have not only changed how we make products, and how we consume them, but they have also made it easier than ever before for innovative entrepreneurs to go global.

Just like their forbears today's entrepreneurs still have to be skilled enough to find valuable problems worth solving, but unlike their forbears they now have access to technologies

and markets that are a match for their ambitions and ingenuity, and as a consequence it is now easier than ever before for one individual to disrupt the status quo faster than ever before.



TIME TO 50 MILLION USERS

As industries become increasingly digitised and as the world becomes increasingly connected it's only a matter of time before we see an industry disrupted in a day and a multi-billion dollar enterprise built in minutes - a trend that is further accelerated by the emergence of Creative Machines.

DISRUPTION IN A DAY. EVERY DAY.

WHEN THE inventiveness of human ingenuity has access to exponential technologies that are powerful enough to help express it and share it on the global stage amazing things happen, and the world is transformed.

In the past, for example, without an effective way to distribute and sell their ingenious products even the world's most inventive entrepreneurs could only at best have hoped to have an impact within their local geography and at a speed that would bore most of today's entrepreneurs to tears.

Fast forward to today though and that dynamic has changed significantly thanks to a slew of advances that have made it easier than ever before for one person to impact the lives of billions of people in a fraction of the time it used to take.

Consequently, when we review some of today's statistics it doesn't take a genius to see that the pace of product adoption, and ergo any potential for disruption, is accelerating.

A scant hundred years or so ago it took over 75 years for over 50 million people to adopt the telephone, but fast forward to today and it took just over 19 days for the same number of people to adopt Pokemon Go.

Q: If a new product can be adopted within just weeks or days by tens or hundreds of millions of people, at what point do we see an entire global industry disrupted in a day?


THE RISE OF CREATIVE MACHINES

Furthermore, when we consider that today it still takes human developers months and years to develop products before they are fit to hit the market, what impact would autonomous companies, powered by so called Creative Machines, machines that can design and innovate new hardware and software products in hours and days today, have on this dynamic?

If they like human entrepreneurs, using in their case a dearth of Big Data,

can identify valuable problems to solve, and then autonomously create and manufacture the products, and build, operate, and scale autonomous enterprises, at what point do we see the time to disrupt industries fall to minutes, and the cycle of disruption accelerate exponentially?

Q: What happens to the rate of disruption when Creative Machines are capable of innovating new hardware and software products in hours and minutes?



FROM MY INDUSTRY TO YOURS

All industries are connected with one another and as digitisation erodes the barriers that kept them all distinctly separate not only do changes in one affect the others faster but it's also now easier than ever before for companies in one industry to enter and disrupt other industries, thereby accelerating the overall rate of disruption.

NO MORE INDUSTRY BOUNDARIES.

WHILE IT has always been the case that changes in one industry would eventually ripple out and affect other industries, when it comes to accelerating the rate of global and industry disruption digitisation simply adds rocket fuel to the fire.

Furthermore, as enterprises and industries accelerate their rates of digitisation one of the most significant impacts of digitisation is the erosion of the individual boundaries that previously kept all of these industries separate and distinct from one another.

Today we see this effect manifesting itself time and time again, where enterprises who've traditionally only operated in one industry vertical are now able to branch out easier than ever before to capitalise on market opportunities in other verticals.

The best examples of this trend typically hail from the technology industry where companies in the so called FATBAG collective, where the acronym stands for Facebook, Alibaba, Tencent, Baidu, Amazon, and Google, now seem to be able to develop new products

and services that cross previously unassailable industry boundaries with impunity.

Amazon, for example, was primarily a E-Tailer, but now the company has interests in everything from finance and healthcare to entertainment. Google meanwhile was primarily a search engine, but now has interests in everything from communications and energy, to finance, healthcare, and transportation. And so the story goes on for all of the other enterprises in this collective.

What most of these enterprises have in common is that they were born in the digital era and so they started out their lives as digital natives. As a result, unencumbered by the need to produce and sell physical products their businesses were afforded a level of agility and fluidity that many legacy organisations, encumbered by physical assets and products, simply couldn't match.

Now though, a decade or so into this new digital era, those legacy incumbents

are trying to catch up by spending hundreds of billions of pounds on wide scale digital transformation programs, and once those programs are complete many of those incumbents are going to have the same opportunity to move into and disrupt adjacent industries with a fluidity and speed that they could only have previously dreamt about. And as a result the pace of disruption will, again, accelerate.

BUILDING EXPONENTIAL ENTERPRISES



CONTRARY TO popular belief, and as obvious as some of this might sound, there are two reasons why individuals and organisations get disrupted.

Firstly, there are the things that disrupt you because you never saw them coming. In short they blind-sided you and, if you have them, your foresight teams.

Secondly, there are the things that disrupt you because even though you saw them emerging and then ascending you never took the necessary actions to counter them. And while the markets and stakeholders will sometimes forgive executives for the former, they rarely forgive them for the latter - especially in a world where disruption is an ever present stalking horse.

Needless to say, disrupting a competitor, an industry, or even a country, is complex, but while many people often like to think of disruption as a singular event it's actually a series of events that, in the majority of cases, have clearly identifiable milestones and markers that we can monitor and track.

However, while everyone agrees that disruption has always been with us and that it can take many forms, from the

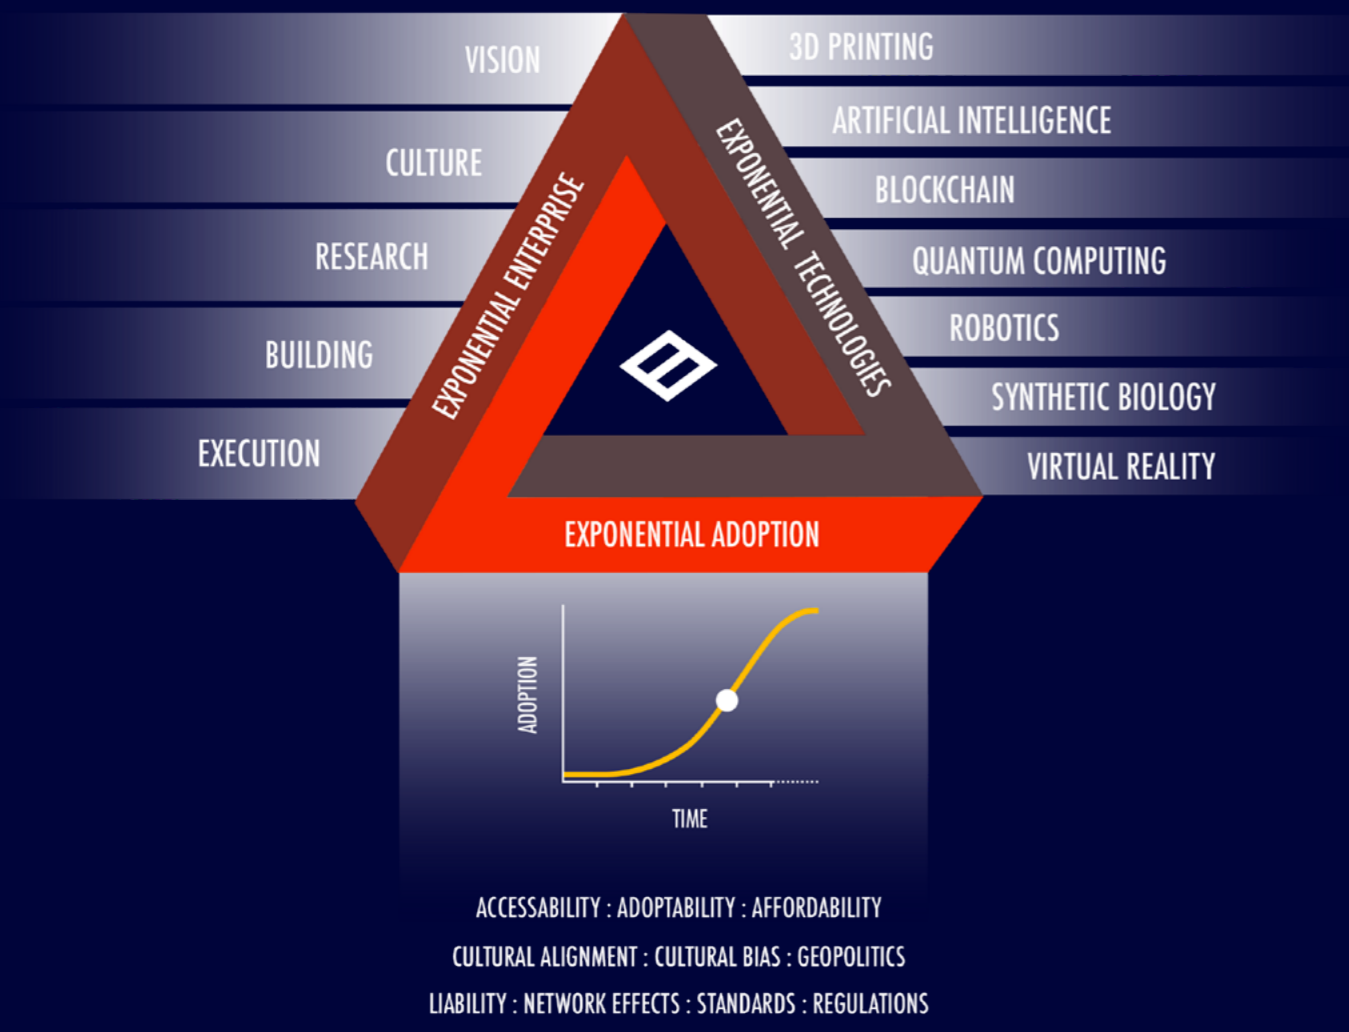
asteroid that wiped out the dinosaurs to the emergence of Netflix who wiped out the video-saurs, one thing that many people still struggle to understand is how the nature of the animal's changed over time and how it will continue to evolve in the future.

Often the reason for this is because sometimes they're looking for disruption in the wrong places, trying to predict it based on historical perspectives, and sometimes it's just because they haven't been exposed to it before. And as for those among you who believe that the majority of disruptions are behind us I can assure you they aren't, and trust me when I say you haven't seen anything yet.

MAPPING THE DISRUPTION LABYRINTH

The process of disrupting anything, whether it be a competitor, an industry, or even perhaps a country, is generally so complex it's positively labyrinthine.

Like all of us though I've lived through many disruptive events and it's these experiences and the impact they had, on enterprises and workforces alike, that drove me to map the labyrinthine-like process of disruption so that companies could understand it, navigate it, use it to



THE DISRUPTION TRIANGLE

The likelihood that a new product or service an enterprise or industry, can be assessed by its progress against three main axes - namely the Exponential Enterprise axis, the Exponential Technologies axis, and finally the Exponential Adoption axis, all of which are intrinsically inter-connected with one another.

Notes:

their advantage, and ultimately come to terms with a world that operates using a new rule book and that no longer behaves like it used to.

As highlighted in earlier chapters, irrespective of how fast disruption seems to materialise it isn't a single event - it's a complex series of events that, in the majority of cases, have clearly identifiable milestones and markers that we can monitor and track, and it's these events that will be the focus of at least part of your discovery process and that will help the vigilant among you identify the next disruptors and disruptions long before they have a chance to wreak their havoc on our companies.

Similarly, these events, and how they combine and the timings of their combinations, also help explain why only a fraction of companies ever make it through the labyrinth to claim cult disruptor status, so let's dive in and have a look at them.

THE THREE AXES OF DISRUPTION

In my experience the likelihood that a new concept will disrupt a market can be assessed by its progress against three main axes as shown in the diagram on the previous page - namely

the Exponential Enterprise axis, the Exponential Technologies axis, and finally the Exponential Adoption axis, all of which are intrinsically linked with one another.

EXPONENTIAL ENTERPRISE

If you're one of those individuals who doesn't want to change the world, and let's face it, not everyone does, and that's fine, then it's unlikely you ever will - at least on purpose. But, if you feel that it's your calling and you can't think of anything else then with the right approach and support you may well just pull it off - never say never, especially in a world where it's easier than ever before for one individual or one company to impact and influence the lives of billions of people.

However, while a determined rebel unit with a disruptive mindset within an enterprise will often be able to change the attitudes and opinions of those who fall within their sphere of influence it has to be argued that true change within an enterprise must be inspired and promoted from the top down.

Over the past decade I've made it my mission to understand what sets enterprises that achieve cult disruptor

Notes:

status, as well as fabled Unicorn status, apart from the rest of the pack and frankly it's a myth that a company's ability to disrupt itself or a market is based on its ability to outperform its competitors in just one single area. In my estimation it's their ability to outperform them in over thirty different areas, often simultaneously, that makes the difference. From the way they build and communicate their culture, values, and visions, to the way they identify valuable problems worth solving and develop their products, ecosystems, and go to markets, and much more, it all counts.

In short, and to be crystal clear, it's not any one thing, it's many, and that's the reality that anyone wanting to build an Exponential Enterprise has to contend with - you're either all in or you might as well go home, anything less and you'll be increasing your likelihood of failure.

Furthermore, it's not simply enough to be moderately better than your competitors, whoever they are and whatever industry they hail from, you have to outpace, outperform, and outthink them all in almost every one of these areas.

Now we've covered the basics let's dive in and have a look at what makes these serial disruptors we're all fond of so special.

In order to make it easier to digest I'm going to divide the DNA of an Exponential Enterprise into five foundations. In order these are Vision, Culture, Discovery, Prototyping, and Execution, and within each of these individual foundations there are at least six main areas that, when performed well and combined, will move the dial in the company's favour.

Firstly comes their Vision, something that conveys a huge amount of information about their overarching purpose and culture, and ultimately acts as their North Star.

Visions and vision statements are normally the aggregated result of a company's ambition and purpose, their discovery and due diligence process, their internal and external deliberations, their framing and the time frame they're working within, and their view of the intersecting trends that they believe will help them achieve their goals.

Generally speaking many of the enterprises that have the greatest impact on the world today and the ones with the greatest disruptive potential are the ones that have bold and ambitious visions with grand aims that, in the words of Elon Musk, get people excited about waking

Notes:

up every morning and feeling inspired by the work they do.

Secondly, and by far the most important of all the five foundations is Culture, which is, among other things, the aggregated result of structural and behavioural company alignment, authentic, inspirational leadership, honest communication, and, again, the company's vision.

We are continuously reminded about the power of culture and it's power to help companies overcome all manner of obstacles. But while creating a winning culture can take years to build and is arguably one of the hardest things for any leadership team to accomplish if you aren't vigilant it can be torn apart in just months.

Furthermore, from a disruptors perspective at least, I like many people have lost count of the number of times I've heard stories about how a company's corporate immune system was responsible for killing the latest innovative concepts - either because they were disruptive to the company's core business, which is obviously laughable under the circumstances, or because of some other political motivation.

Thirdly comes one of the most exciting

foundations, in my opinion at least, Discovery, which is the aggregated result of internal and external conversations, collaborations, and partnerships, exploration, envisioning, and observation, and much more. This foundation is also often the natural home of the majority of a company's entrepreneurs, rebels, and visionaries - the teams of individuals who all too often want to rip up the rule books, go above and beyond, and disrupt the status quo.

And as the rate of disruption accelerates, and as more enterprises feel the effects of disruption on their balance sheets it's no surprise that over the past number of years many of the teams in this space have been the beneficiaries of significant uplifts in funding and new programs as the companies work hard to improve their competitiveness, and defend and extend their customer bases.

All that said, however, it obviously goes without saying that new funding and programs by themselves can't be counted on as magic bullets that guarantee success. Again, it's not one thing, it's many things working in harmony, which, neatly brings me back to the importance of having the right culture and environment.

Fourthly we have the Prototyping

Notes:

foundation, where companies begin to build products that address the problems and opportunities uncovered during the Discovery foundation. This foundation is the aggregated result of conversations, collaboration, and partnerships, experiential and design thinking, ideation and problem solving, to name but a few. One of the most understated areas of this foundation though is the use of beta customers and, where appropriate, the importance of the investors black books - both of which help companies secure early testers and customers that eventually hopefully convert into paying customers and references, with the added benefit that, with the right management these activities and customers will help generate hype around the products that then, in some cases, propel them into the hands of millions of customers.

Fifthly, and by no means least is the Execution foundation that, when done right, which is obviously harder said than done, ensures your amazing new product doesn't get left on the metaphorical shop shelf to die.

The aggregated result of everything from ensuring the right balance of customer value and the right business model and go to market strategy this is where many companies ambitions to disrupt markets fail. As they say - everyone has a plan

until they're punched in the face, or in company speak everyone has a plan until it meets reality.

However, for the lucky companies that do make it past this last hurdle to disrupt a market - whether they're lucky by design or by fluke - this is the stage where all their hard work, everything I've discussed, albeit lightly so far, pays off.

This is also the point at which the incumbents in a market realise that a disruptor has just parked their UFO on the company's front yard, before laughing at it, shrugging it off, and getting eaten by the aliens hoards inside...

Noone ever claimed disruption was easy but throughout my travels and conversations with executives from all manner of industries all around the world it's clear that almost everyone underestimates the complexity and size of the challenge. However, while disrupting any market is difficult it's also clear that the size of the prize, which is often the opportunity to lead and own a market, is worth the effort.

EXPONENTIAL TECHNOLOGIES

Once a company has started its journey

Notes:

to become an Exponential Enterprise and found interesting and valuable problems worth solving next they turn to technology, explicitly combinations of technologies, to develop their products and help get them into the hands of consumers.

And, as you can see from the Griffin Exponential Starburst in the earlier chapters and by reading the other codices in my Codex of the Future Series, there are hundreds of exponential technologies that enterprises can choose from to help them change the economics of their industries, and develop new disruptive products. And more are appearing all the time.

One of the phrases you'll hear me refer to many times throughout this codex is the word exponential, a term that I'm sure you've heard a million times that's often used to refer to technologies that emerge, develop, and mature very quickly, and often at a rate that very few people anticipate or predict.

The term is also a hangover from Moore's Law where Gordon Moore, Intel's co-founder, in 1965 predicted that the number of transistors on a computer chip would double every 18 months, leading to an exponential increase in computing Price-Performance, and

today we're seeing the same pattern emerge in many other technologies - from Artificial Intelligence (AI) and Quantum Computing, to 3D Printing and Gene Editing, and many others.

Although, when it comes to digital technologies, such as AI and Creative Machines, for example, their rates of development even make Moore's Law look positively lethargic, and this is yet another trend that's accelerating disruption.

As the rate of technological development accelerates though there is also another trend you should familiarise yourselves with called "Jumping the S-Curve," and it's important because, in short, it refers to the way that different technologies supersede one another. Furthermore, as the number of exponential technologies that are emerging continues to accelerate and increase this is yet another accelerating trend that you have to take into account when deciding which technologies to use to build your new products and go to market strategies.

The phrase S-Curve refers to the rate of development of a particular technology - like a squashed S first the rate of development starts slow, then it accelerates dramatically, and then it flattens off as researchers struggle to eke

Notes:

out further gains. Furthermore, today, and more so in the future, as the period of time it takes to reach higher levels of Price-Performance accelerates you'll no doubt find that trying to keep pace with all these developments gets even harder.

Jumping the S-Curve then refers to a company's ability to move from one older technology to a newer one, for example, moving from the logic based x86 computers that we use today to tomorrow's ultra-powerful Quantum Computers. Unlike the past though where there were only a few S-Curves to jump now there are potentially hundreds - all of which can be combined in new and interesting ways to further fuel the rate of disruption.

EXPONENTIAL ADOPTION

Of course though, while having an enterprise with the right culture that's capable of identifying valuable problems and opportunities, and which is highly adept at leveraging talent and technology to build great products is a great start the fact remains that you have to get those products into consumers hands.

So, as part of your Execution strategy, it should come as no surprise that there are

plenty of areas left that, on the one hand could stop you dead in the water, or, on the other boost you into the hall of fame. And these areas are so important that I decided to give them their own axis.

While I've already discussed how disruption is a process and not a single event this is the stage where, if you want to disrupt a market, you have to gain as much traction as possible in as short a time frame as possible in order to stymie your competitors ability to counteract you with their own messaging and variants.

Getting your product into the hands, hearts, and minds of consumers though at enough scale to disrupt a market and permanently change the status quo though is obviously difficult. But that said while, yes, you still have to overcome many hurdles, and successfully pull all the right levers you should be able to take comfort from the fact that today, as I've highlighted in previous chapters, it's easier to disrupt the status quo than it ever has been before.

Navigating this part of the labyrinth though is complicated which is why the majority of enterprises struggle to realise their lofty ambitions, and sometimes all it takes is for one key piece to be out of alignment and everything falls down like

Notes:

a deck of cards.

For example, build a great product that the regulators block and you're going nowhere, or build a great product that the regulators approve that is unethical, and yep, again you're going nowhere. And so it goes on - you get the picture.

So, as you can see again gaining mass adoption of your product isn't down to getting one thing right it's down to getting many things right. These include, but are not limited to, your products accessibility, adoptability, and affordability, as well as other factors including cultural alignment and bias, ethics, the geo-political situation, the impact of insurance and liability, network effects, and, of course, standards and the regulatory environment.

Get one of these wrong or get side slammed by one of these and it could be game over for you and your new products.

SUMMARY

Today we live in a world full of opportunity where the rate of change is accelerating every day, and where exponential technologies are global enterprises force multipliers and startups

levellers. And, as a result, yesterday you had tens of competitors in your rear-view mirror, and today you have hundreds. Or more. It's fun to be you.

However, as amazing as all this is it will all soon be eclipsed by an even bigger, and even more disruptive revolution, because a new breed of entrepreneur, one that can out think and out perform humans a million fold to one, and build fully autonomous multi-billion dollar empires within days and months is already emerging.

I am, of course, talking about the rise of Creative Machines, synthetic entrepreneurs if you will, and for those of you who think that such talk of machines that can design and innovate products, and operate and scale companies is far fetched the first fully autonomous enterprises have already been built and they're already operating on two continents.

Today is the slowest rate we will ever move again, but you've seen nothing yet. So pause, take a deep breath, and prepare yourself for what's coming.



MEGATRENDS AND STARBURSTS

EVERY YEAR I publish a new Griffin Exponential Technology Starburst, and a new Megatrends Starchart, that you can see on the following pages, designed to help people envision and simplify the future.

Today, it's no secret that every part of global culture, industry and society are being transformed faster than they ever have before thanks to the relentless, and some would say furious, rate of change that's made possible by giant advances in technology and the megatrends it helps create and drive.

As this rate of change accelerates what you likely won't be surprised by is the fact that, on top of the exponential technologies, such as Artificial Intelligence and Blockchain, for example, that are already here, there are still yet more powerful exponential technologies circling above us like stars in the heavens, that are just biding their time, waiting to fall to Earth where their impact will, over time, be total and irreversible.


What might surprise you though is the sheer number of exponential technologies that are appearing, over 500 by my latest count, with on average more than 60 being added every year.

In the right hands every single one

of these so called "Blank Slate" technologies, so named because until someone innovates on top of them they are just that - blank slates - has the potential to transform either just a part of our society, or all of it.

As powerful as all these individual exponential technologies are though it's when they're combined, to form what I call "Exponential Combinations," that the real magic happens and their power to transform everything is multiplied many times over. And as a result that's what I invite you to examine, and that's why I created the Starchart and Starbursts - so you can join the dots, combine the individual megatrends and exponential technologies together to transform your own futures.

2020



The World's First Living Robot, 2020.
Source: University of Vermont

THE YEAR AHEAD 2020

THE PACE of technology development and disruption is accelerating, and every year that's clearly evidenced by the new additions to the Starburst. In fact, some of the new additions are so amazing that even I'm stunned by their emergence.

Take, for example, this year's addition of Living Robots, not only is this technology positively science fiction incarnate, but these new lifeforms, or Xenobots as they're now being referred to, were designed by algorithms running on supercomputers.

This technology breakthrough is also a prime example of how, increasingly, it's getting easier to spot radical new technologies and technology concepts that help us the S-Curves, by which I mean help us move from one technology paradigm to the next.

As I've always discussed you should think about Exponential Combinations not just simply Exponential Technologies, and as we now race headlong into 2020 I expect the rate of new technology development to accelerate again.

When I first began crafting my Starbursts I would see approximately 40 to 50 new exponential technologies emerge in any given year, and now I'm seeing between 60 and 70, an increase of over 40 percent, and that clearly shows not only that we are accelerating but that we're significantly accelerating.

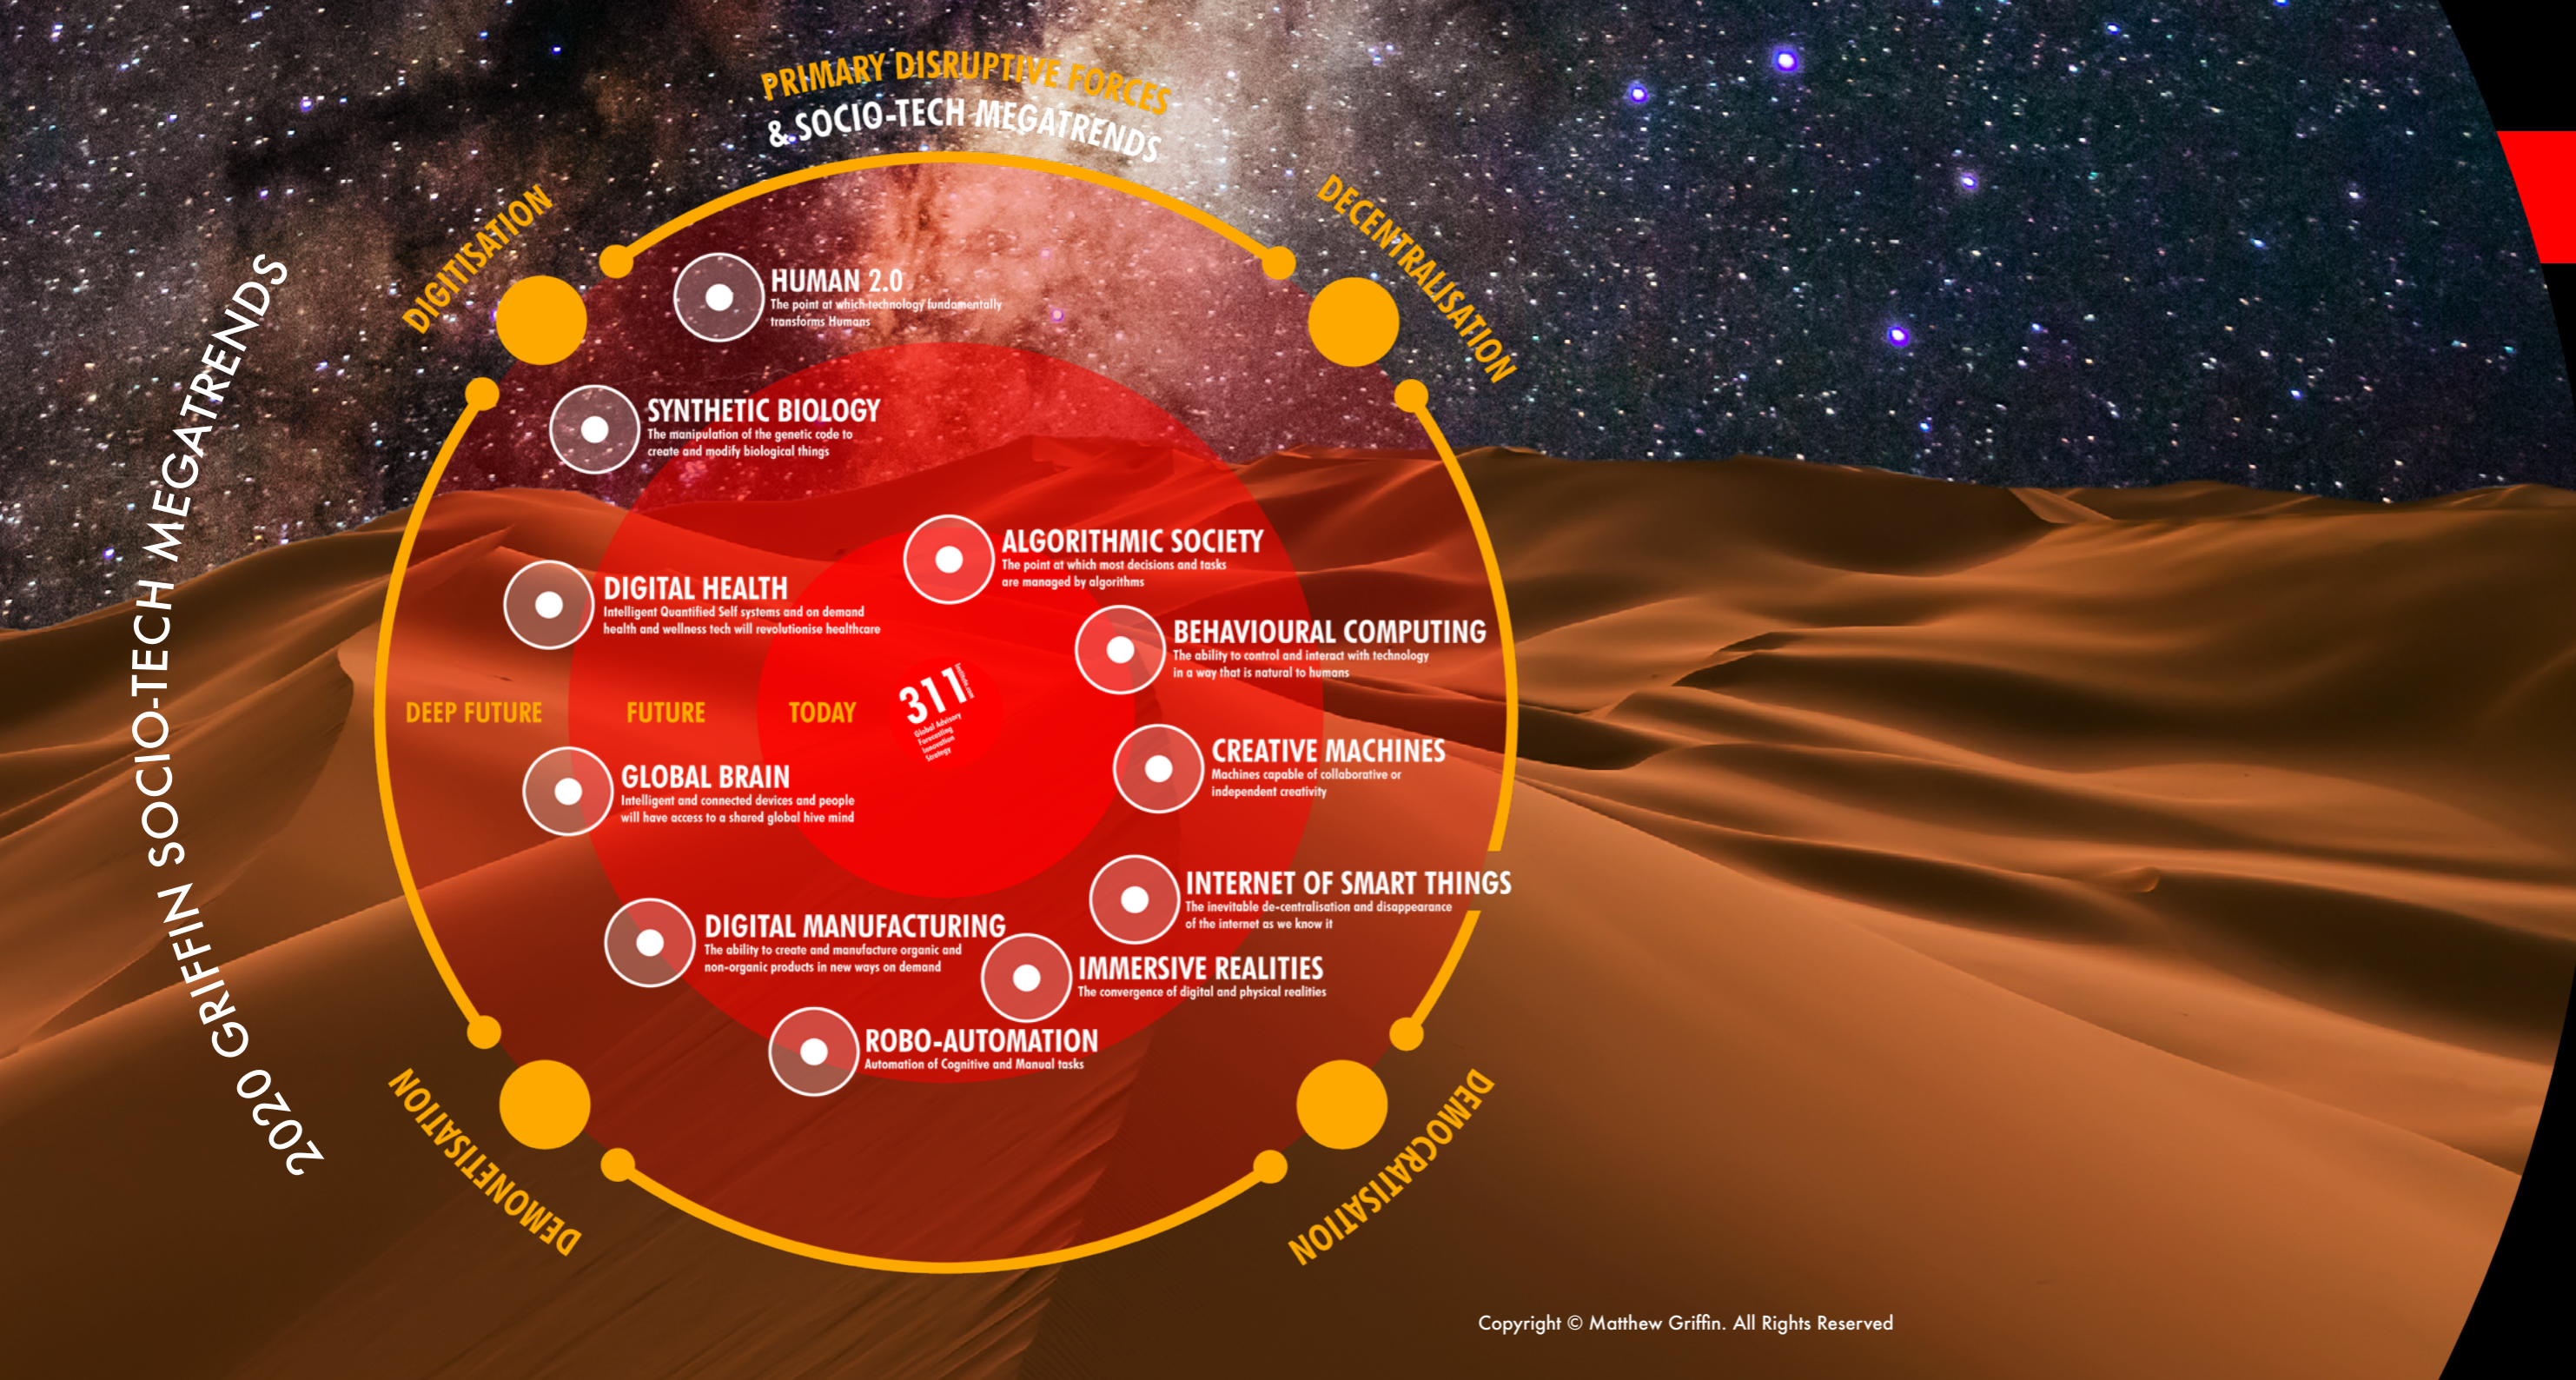
Heading into 2020 I expect that this year's headlines will be dominated again by 5G, Artificial Intelligence, Blockchain, Creative Machines, hypersonic technologies, Quantum Computing, Robotics, and new user interface technologies such as Augmented Reality, Mixed Reality, and Virtual Reality, and that the headlines will be dominated again by talk of automation, DeepFakes, the battle for military supremacy, and the normal iterative tech coverage and its effects on society.

However, back in the labs, there will be developments that make science fiction look tame, from new biological computers and synthetic lifeforms, through to new types of molecular assemblers and neural tech, so be sure to keep checking my blog for updates.

SOME PREDICTIONS

IN 2020 I predict the following:

- 5G and Low Earth Orbit satellite deployments will accelerate.
- Any technology with a sustainability twist will attract out sized attention.
- AI breakthroughs, while still numerous, will slow as new models take more computing power to train, as a result there will be an increased focus on lean AI and zero-day learning AI.
- Creative Machines and Synthetic Content will improve to the point where it blows through uncanny valley.
- Energy technologies will continue to see a Cambrian Explosion.
- Synthetic Biology breakthroughs will accelerate.



MEGATRENDS STARCHART

YOU COULD spend your time analysing and tracking the literally thousands of megatrends that are on the rise, across the Societal, Technological, Environmental, Economic, and Political spectrums, or “STEEP” as they’re known for short.

However, in order to stay focused and keep things as simple as possible for the purposes of this Codex, while still letting you retain the detail you need to compliment your big picture thinking, I find that the best megatrends to keep track of are what I call the Societal-Technology megatrends, or Socio-Tech for short. In short these are the megatrends that will not only have the greatest impact on culture and society as a whole, but also industry, so it’s a good catch all.

EXPONENTIAL IMPACT

HUMANITY'S STORY is one that is inextricably intertwined with technology, in all its forms, from, for example, the early railways that connected our early cities to the telegraph lines that connected our early communities. But, as generations came and went the memory of the power and impact of these early exponential technologies faded, and now they're consigned to the history books and museums as relics of the past.

However, while our memories of those early technologies might have faded their legacies live on, and today the transformative power of the descendants of these and other exponential technologies have become even more impactful, and they're transforming our world in new previously unimaginable ways at an accelerating rate.

The telegraph, for example, was replaced by faster more convenient fixed line telephone systems, which in time were themselves usurped by faster, superior mobile communications technologies.

First came 1G, then 2G, 3G, 4G, and now 5G, and just eight generations on from the original telegraph system that connected people using mechanical clicks and whirs our world lives online, and

**FORGET ABOUT
EXPONENTIAL
TECHNOLOGIES ...**

**... THINK INSTEAD
EXPONENTIAL
COMBINATIONS!**

people have embraced a new type of clicks, and communicate and experience life in bits and bytes in a world where science fiction is increasingly difficult to differentiate from science fact.

However, the transformations we've witnessed over the centuries aren't thanks to the development of any single technology, they're the result of many technologies all working in combination with one another, and this is why individuals, as well as enterprises, must move away from today's rather siloed thinking where we tend to talk and think about the impact and opportunities of singular technologies, and instead think about the impact and opportunities of "Exponential Combinations."

After all, even today's most powerful exponential technologies are simply blank slates that themselves rely on the development of a host of other exponential technologies, as well as an army of human and increasingly machine based entrepreneurs, that prod, shape and combine them to create new amazing concepts, to drive their development and eventual adoption.

It's these combinations, of not tens, but hundreds of exponential technologies, like the ones displayed on my Starbursts, that enable us to transform every corner

of society, from the way we live our lives and how long we live, to where and how we work.

Furthermore, thanks to technologies such as those I mentioned earlier, communities and individuals that were once limited by connectivity and distance now all have increasingly easy and low cost access to a single "global brain" and global resources that can help even the most modest among us transform the world in new and exciting ways, and as these technologies become increasingly digitised and democratised the speed and impact of that transformation, as I also discussed in earlier chapters, will only increase from here.

**BUT... EVERY
TECHNOLOGY HAS
TWO SIDES.**

**BUT... EVERY
TECHNOLOGY HAS
TWO SIDES.**

**GOOD
Y
S
UTOPIA
O
P
I
IS A BLANK SLATE...**

AS WIDE ranging and as powerful as all the exponential technologies that I discuss in this codex are though the fact remains that until someone uses them and combines them together to innovate new products and services they're all just shelfware - blank slates, and technologies without a purpose.

Every technology is a blank slate that can be used for both good or bad purposes. It's down to us to develop and use them in ethical and moral ways that benefit society.

Furthermore, as these exponential technologies and the products and services they can be used to create become more powerful they then give us a moral and ethical dilemma because, just as they can all be used to do great good and benefit society, in the wrong hands they can also be weaponised and cause great harm in a huge variety of ways - many of which we have yet to even imagine.

Take, for example, Artificial Intelligence. On the one hand it has the power to revolutionise healthcare, identify, treat and cure disease in new ways, and discover new powerful drugs and

vaccines, but on the other it's also already being weaponised to create a new generation of Robo-Hackers that can hack and exploit vulnerabilities in critical computer systems hundreds of millions of times faster than human hackers, and that's before we discuss how it's also being used to generate fake content and fake news that undermines our trust in one another and democracy.

These world changing examples are just the snowflake on the tip of the giant melting iceberg, and an example of what good and bad actors alike can do with just a single powerful technology. And there are billions of other examples I could use, including our ability to save lives by using drones to deliver critical first aid supplies including blood and medicines to remote areas, or spray crowds with bullets from drone mounted machine guns.

While this is where I'm going to leave it for now I can spin similar stories and examples for every exponential technology which is why it is absolutely vital that as a society we do our utmost to understand the pros and cons of these technologies and work together to maximise the upsides, while doing our best to mitigate, regulate and police the downsides.

THE FUTURE OF GAMING



TODAY GAMING, in all its glorious forms, is undergoing a transformation unlike anything we've ever seen before. Not only are today's top gamers pulling down seven to eight figure pay days and, like their sporting cousins, being signed up to lucrative sponsorship deals by top brands, but increasingly their adversaries are just as likely to be cloud streamed Artificial Intelligence fuelled entities that, in an increasing number of cases, are mercilessly crushing their human opponents into the ground.

And all of that is before we discuss what it feels like to be immersed in a virtual reality world while wearing a haptic suit that makes sure you feel every kick, punch, and sniper shot hitting your torso in a virtual gaming arena.

So, whether your a top gamer or just someone sitting on the sidelines watching it all the future of gaming has never been more exciting, the experiences as intense, and the stakes have never been higher.



CHALLENGES TO OVERCOME.

WHEN WE have a look at the key challenges that the future gaming world has to overcome, especially as we watch the steady rise of new cloud game streaming services, you can put them into a number of different buckets.

Firstly, there are the cultural challenges of trying to get mass audiences to adopt new styles of game play, whether it's immersive Virtual Reality experiences or getting them to adopt new ways of consuming games, such as the emergence of new cloud game streaming platforms that let them play anything anywhere.

Secondly, there are the obvious technological challenges, as I'll discuss in more detail later in the next chapter, of being able to deliver these new experiences and stream low latency high quality content at resolutions of up to 8K and 120fps that will satisfy even the most demanding of gamers.

And then thirdly, there's the inevitable business challenge of building and scaling the technologies and platforms

to support all of these experiences, the challenge of recruiting and developing the right talent and ecosystems to support their development and run it all, and then there's the ultimate challenge of making the new world profitable - something that so far, on the streaming front as well as the VR front, almost every company has had problems with.

THE CULTURAL CHALLENGE

When it comes to overcoming the cultural challenge it's clear that people are ready for and happy to stream content across all of their different devices provided the content is right, the technology is invisible, and, of course, the price is right. And, let's face it, none of these are small challenges to have.

Furthermore, in today's world these challenges are often further compounded by the fact that many of the enterprises involved in the streaming space feel compelled to own the full stack and be vertically integrated businesses - something that admittedly carries significant rewards but that also, if they

get it wrong, carries significant risk.

THE TECHNOLOGY CHALLENGE

There's a reason why today most games are based locally on gamers devices and machines, and that's because in order to deliver the highest quality gaming experiences gamers insist that their games are full of rich content and are lag free - something that is especially important if they're playing fast paced shoot 'em ups where fast reaction times and timing is everything. And as soon as you try to move those local games into a remote location, such as the cloud, one or both of those experiences suffer.

As a result, as we look at the future of gaming and cloud gaming enterprises that want to dominate the market are stuck between a rock and a hard place because it's generally not good enough, from a gamers perspective at least, to offer one or the other, namely rich content or low latency, companies have to offer both. And that's hard especially when you're trying to cram huge volumes of content and data across variable trans-

Notes:

national or trans-international networks and questionable mobile networks and domestic WiFi connections, and laggy edge devices.

All of this, of course, is even more of a problem when you consider how creating all of these games will change in the future too with the arrival of so called Creative Machines, that I'll discuss in more detail in a later chapter, that can create games and synthetic content, as well as procedurally generated content specifically tailored for each individual's own preferences, on the fly.

However, in spite of the many technological challenges that the cloud gaming market face today over the coming years they'll all be whittled away in the next decade by the emergence of increasingly cheap and performant exponential technologies and tools.

Technologies such as 5G and then 6G, WiFi technologies including WiFi 6 and WiGig, new types of application and content delivery networks, AI codecs, and more powerful edge devices that are crammed not only with more GPU and memory that you can shake a stick at but also AI neural chipsets and chiplets that increasingly make running anything and everything at the edge, without having to run it all through a giant hyperscale

datacenter first, a breeze. And when it comes to generating the synthetic content to feed gamers thirst for richer games and experiences, well, faster cloud datacenter hardware, and incredibly powerful AI's will take that challenge on.

THE BUSINESS CHALLENGE

When you have conversations with everyone about the cloud and related on demand models one of the biggest issues that companies cite when trying to dominate these new market opportunities is the fact that you have to operate a "Build it and they will come" model. And if you do build these cloud platforms and people do come then everything is rosy, but if people don't come, well, I can't tell you how many times I've lost count of companies that bet the farm and lost everything.

As a result when we look into the future of gaming, especially as it relates to streaming, it's therefore no surprise that the companies who are moving first into the space already operate successful multi-sided business models where the profits from one part of the business can be used to sustain what they hope will only be temporary costs and losses from the new venture. In short they operate a blended approach that lets them manage

Notes:

costs and risks much more effectively than any non multi-sided business could.

However, as new consumption and distribution models pop up, and as new technologies continue to make their presence felt in the market, as I'll discuss in the next section, there will be plenty of opportunity for companies to innovate their monetisation strategies.



BUSINESS MODEL INNOVATION

AS THE prominence of the 45 year old gaming industry continues to expand and grow it's inevitable that new business models and monetisation strategies are going to emerge, so in this chapter I'm going to take a quick look at the front runners and how they all compare.

With the growing popularity of E-Sports and mobile gaming, and the arrival of cloud game streaming services multi-national brands especially are now taking a closer look at the gaming market in order to expand their reach and revenue streams.

In 2017 alone the global gaming industry generated over \$120Billion, according to Newzoo, and with the explosion of games like Fornite as well as others, and the introduction of new gaming technologies it's widely expected that revenues will top \$180Billion in 2021.

While companies across mobile, console, and computer gaming have traditionally profited from selling games and devices, new methods for generating revenue are now taking off.

Today free-to-play online multiplayer games have finally come to consoles, while native advertising now appears in some of the world's most popular games, and an increasingly wide range of companies are finding new and novel ways to expand their reach, developing everything from cloud gaming platforms and applications through to developing and investing in full blown global E-Sports leagues.

As the gaming world becomes increasingly popular around the world, and as new technologies help expand its reach and appeal, developers, gamers, and companies alike are all searching for new ways to benefit.

What follows is a quick review of several of the emerging business models that are transforming the industry and changing the way the gaming ecosystem operates.

1. CLOUD GAME STREAMING

OnLive, who went on to be acquired by Sony, was one of the first companies

to offer cloud game streaming, albeit in a limited form with some significant technical limitations, and as discussed that act was quickly followed in a more sophisticated way by a host of other companies including big names such as AWS, Google, Nvidia, and others, with Microsoft and Sony eventually coming together in an alliance to try to defend their turf.

While cloud game streaming is still in its early years there's obviously growing interest in it, in one part because of the so called Netflix effect which has now firmly established itself as the dominant business model that's turned the traditional cable and content company industries on their heads - and there's no denying the appeal of being able to play any game anywhere on anything for a fixed upfront fee.

As a result it doesn't take a giant leap to see how cloud gaming will become the defacto way most games are consumed, and when it comes to monetisation strategies, well, all of the models I've talked about here can be applied to it, but there are also significant upsides

Notes:

when you consider this type of service's ability to combine them all together into one neat, seamless package.

2. CROSS BORDER PUBLISHING

Traditionally, game publishers in the West published in the West and game developers in the East published in the East. However, with growing similarities in gaming style and preferences around the world we're seeing publishers enter new markets abroad.

In 2017, Tencent released Arena of Valor in the US for iOS and Android. This followed the success of the game's popularity in China. Then in 2018, Tencent released the mobile version of the MMO Battle Royale game PlayerUnknown's Battleground. This followed the worldwide success of the PC and console version of the game, and the mobile version became the number one most downloaded iOS app in 48 countries, and still currently ranks in the top 10 for free iOS games in the US.

Furthermore, while there are rumors that Tencent is now looking to back the game's original developer Bluehole Studio, the company has already invested \$15Million in Epic Games to bring Fortnite to China.

3. DIGITAL ASSET EXCHANGES

Today many gamers earn in-game items that increase their performance or improve their character's or environment's aesthetics. Sometimes, these items are issued for in-game achievements, while other times they are purchased as part of in-game micro-transactions even though historically these so called assets are difficult to exchange and often have no second hand resale value.

As a result, and in order to try to tap into this potential second hand market marketplaces that let users trade and sell their digital assets have sprung up.

Companies, for example, like Gameflip and OPSkins have not only built a business around these activities by helping enable these trades but they've found a way to monetise them by taking a fixed fee for each transaction, a business that's so lucrative that recently OPSkins created a separate, decentralised marketplace called Worldwide Asset eXchange (WAX), where gamers can tokenise digital goods to be exchanged for cryptocurrencies. And according to the website, so far over 400 million gamers purchase over \$50Billion in virtual items worldwide.

Notes:

4. E-SPORTS LEAGUES

With the ubiquity of technology today, the definition of sports is expanding to include "games facilitated by electronic systems." Consequently, these competitive games, or E-Sports as they are better known are providing new mediums for monetization on a global scale with major brands sponsoring them and with even Atari recently jumping into the space and creating the first dedicated E-Sports and gaming hotels in the US.

While E-Sports themselves obviously provide a variety of avenues for monetisation, through advertising, sponsorship, streaming, and tickets sales, they also help promote the games played and generate enthusiasm. Additionally, just as is the case with professional sports, E-Sports platforms are now also increasingly being used to promote the activities and games of comparative real world sports leagues, such as the Formula E, NBA, NFL, and NHL, and their respective teams, which only serves to further boost their appeal and revenues.

So far, these leagues have been successful in promoting games like NBA2K20, Madden NFL 20, and NHL 20, as well as their respective sports organisations. As a result these E-Sports

leagues offer traditional professional sports organisations an opportunity to engage fans in entirely new ways, both in and out of season.

In fact the model is now so compelling that even sports stars, such as Shaquille O'Neal, Alex Rodriguez, and Jeremy Lin, as well as franchise owners like Stephen Kaplan, Robert Kraft, Fred Wilpon, and Stan Kroenke, are getting in on the action.

5. FREE TO PLAY MMO

Fortnite, a Massively Multiplayer Online (MMO) battle royale game, has taken the gaming industry by storm. However, when Epic Games first released the title on Xbox and PS4 consoles back in 2017 there was some uncertainty about how it would fare against its traditional Pay-to-Play peers.

After growing its audience to more than 45 million total players across all platforms though, and counting, Epic has been able to successfully monetise Fortnite using so called micro-transactions. Originating with mobile gaming and later adopted by PC games like Dota 2 and League of Legends (LoL), micro-transactions allow players to purchase inexpensive, typically cosmetic,

Notes:

upgrades or items within a game, and while this model has, up until now, been largely lucrative for mobile and PC games it's a new trend for console games.

As a result these micro-transactions allow developers and publishers to lower the price of the game or, in the case of Dota 2, LoL, and Fortnite, offer it for free, but while the purchase of these digital items is optional Epic Games announced that last year alone they managed to take in over \$126 Million, so the model, provided of course you can reach critical mass, evidently works.

6. IN APP GAMES AND INSTANT GAMES

In-app games provide developers and publishers with new mediums to distribute content. They also allow them to reach audiences unfamiliar with mobile gaming or those reluctant to download separate mobile applications.

Both Tencent's WeChat and Facebook's Messenger offer in-app games, referred to as mini-games and instant games respectively, in order to try to keep users within the confines of their so called social ecosystems. Furthermore, the combination of a strong acquisition channel, large user base, as well as the

tremendous appeal of social sharing is also making this a more important monetisation strategy for games developers who want in.

However, while the social media platforms, for example, take a cut of the ad revenue generated by these games the majority of the proceeds go to the developers and publishers, and in the past few months alone WeChat's mini-programs, apps that range from music to e-commerce to games, have topped 170 million daily active users.

Today there are 580,000 mini-programs available online, and the aforementioned mini-game Tiao yi tiao (Jump Jump) reached 100M daily active users alone - in fact it became so popular that WeChat dedicated an entire competitive event to the game.

7. NATIVE IN GAME ADVERTISING

Native advertising, or native monetisation as it's also known, refers to a type of paid advertising that aims to seamlessly blend advertising content in with the existing content and style of a given asset, such as a game, or web page, and it's a strategy that's become increasingly popular thanks to the rise of social media sites social media sites like Facebook and

Notes:

video platforms like YouTube.

Consequently it's no surprise therefore that it's now starting to make an appearance in gaming. Back in 2015, for example, advertisements for Alienware, Dell's computer hardware subsidiary, appeared in Psyonix's Rocket League, a vehicular sports game. While the company claimed it was intended for a private tournament, and not for the general public, it highlighted developers' ability to include advertisements within the game's natural setting.

Electronic Arts, in particular, has also benefited from this strategy due to its broad portfolio of sports games. These games allow for advertisements as they would appear at a real-life sporting event, such as billboards promoting a credit card company. While the company has been leveraging this strategy for over a decade advances in technology and ad targeting now mean EA has a way to display increasingly hyper-targeted, dynamic in game advertising.

Elsewhere Tencent's WeChat has also experimented with native advertising with the in-app mini-game Tiao yi tiao (Jump Jump), in which a player has to jump from one block to another. With over 100 million daily active users as of the beginning of 2018, Tiao yi tiao has

included native advertisements from Nike and McDonalds that appear on blocks in the game.

The majority of other mobile games use formats like popup videos or reward offerings, offered by companies like IronSource. But native, in-game advertising seems to provide a less disruptive experience for the gamer, as it often blends into the game environment. And similar to micro-transactions, native advertising allows developers to monetise otherwise free-to-play games.

8. INTENSITY

As gamers access new sensory experiences via more advanced immersive gaming gear, for example, accessing touch via haptic clothing, or the experiencing the sense of taste and smell via bio-mimetic sensors embedded into their VR headsets, this opens up yet another business model for the industry - the ability to sell subscriptions based on the intensity of the user experience.

If you'd like a tame experience then that's the Bronze tier, but if you want to feel the punches as you get hit in the chest from sniper rounds in the game then that's the intense Gold tier. Now just imagine the rush of that...

THE FUTURE OF GAME DISTRIBUTION

TODAY GAMES are primarily distributed in one of three ways - you can buy physical games online or from a store, you can buy computer games on physical media from online or a store, or you can download games straight to your favourite devices and play them locally.

Now though, as cloud computing platforms and networks get faster and more powerful a forth way to buy and consume games has emerged in the form of cloud game streaming which is nothing less than the gaming world's "Netflix" moment - the moment that changes the entire industry irrevocably and forever. Provided of course that the companies behind them can make the economics work - something that Netflix and all of its many competitors are still having problems with as they all struggle to turn a profit.


However, that asides, from a users perspective provided your networks can withstand the streaming onslaught which, bearing in mind the arrival of 5G today, and then 6G in 2030, and the arrival of domestic WiGig and WiFi 6, all of which will be turbo charged by new super low latency trans-oceanic fiber connections, the only thing you'll have to worry about is what game to play.

As networks get fatter and faster though it's inevitable that gamers will demand 8K and 12K resolution games running at 120fps and above that will suck up even more bandwidth so as a result the companies behind these new cloud gaming platforms will have to devise new tools and tricks to balance latency, quality, and cost - tricks that include new codecs, new powerful computing platforms, and new rendering techniques that, when combined, could reduce future network loads by up to a whopping 90 percent.

INVISIBLE TECHNOLOGY

While game streaming will bring about a wide variety of benefits, from an explosion of new content through to new gaming experiences, its greatest impact will be to make technology itself invisible and let gamers pick up any device anywhere and continue playing from where they left off.

To get all professional on you provided the technology is correctly implemented, which it will be, this will create what's known as the "ultimate frictionless customer experience" and this will drive the adoption of these new platforms into the stratosphere.



A SIMPLE CONCEPT ...

While creating a cloud streaming service sounds simple there is a huge amount of technical wizardry that needs to all come together in order to deliver gamers the experience they demand and crave.



... THAT'S DIFFICULT TO REALISE

But in reality creating a globally scalable platform capable of delivering high quality gaming experiences in extremely difficult and expensive to realise.

STREAMING GAMES FROM THE CLOUD.

THE BENEFITS of streaming games from the cloud are obvious and there is no question that cloud game streaming will be the dominant way that games are consumed and played in both the near and far future, but while the premise seems relatively straight forward there are a huge number of obstacles that have to be overcome before we realise utopia. In fact getting it to work will be a display of technical wizardry and human ingenuity at their best.

ANY GAME ANYWHERE

The greatest promise of cloud gaming is the ability to play any game, on any device, anywhere, and this alone is a big enough prize to encourage companies to steam into the space, but the problems of making this real are very real - even for companies with seemingly infinitely deep pockets and some of the world's greatest experts.

Ironically in a perfect world with the technology we have today none of the problems that are keeping companies awake would even appear on their

radars, but in our imperfect world, where players don't live in the most convenient places right next to fat network pipes and powerful datacenters packed full of powerful hardware, and where they don't all use the same identical and most up to date devices or technologies it's enough to give anyone a headache.

THE DATACENTER PROBLEM

With latency and the quality of the user experience being at the top of everyone's mind the first problem that needs addressing is the issue of datacenters and datacenter build outs, and if you want to dominate the gaming market then you better have deep pockets and be prepared to drop billions of dollars to build out powerful datacenters in multiple countries and places that are as close as feasibly possible to your target audiences.

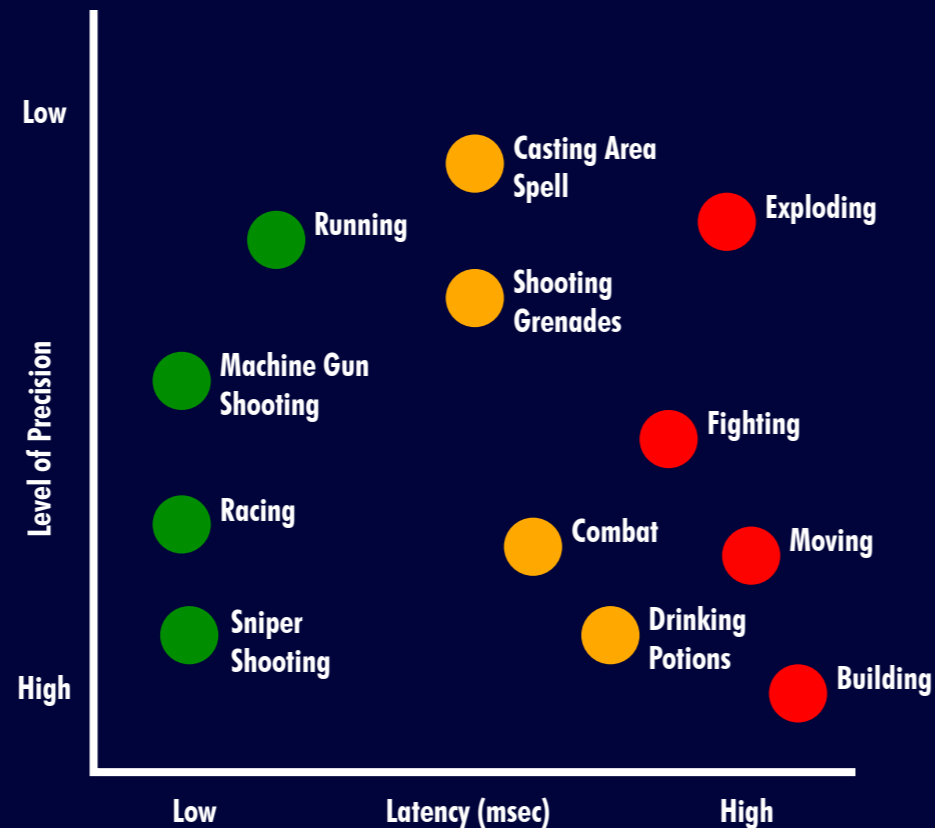

THE ISP PROBLEM

As the amount of traffic streaming across ISP's nodes and networks continues to

increase at an almost exponential rate on the one hand they're going to have the issue of having to invest more heavily in infrastructure, but on the other, and especially if some of the new games streams are encrypted then ISP's could also face the uncomfortable problem of having to play policeman and try to identify malicious individuals who are using these new platforms for nefarious purposes.

THE ECOSYSTEM PROBLEM

When you have a look at all of the companies that are vying to win the battle for our gaming future ironically none of them are gaming companies, for the most part they are all global technology platforms where gaming will be just one side of their multi-sided businesses. But, just as we saw with Netflix, while all of these platform companies today are trying to encourage studios to embrace their new streaming models these same platforms are also buying up game studios so they can own the entire stack for themselves and cut out the middle men.



THE NEED FOR SPEED ...

The richer the user experience the more reliant game streaming services will be on the end to end quality and speed of the hardware and networks needed to deliver it.

... WHEN IT REALLY MATTERS

But the type of game and the activities users are going to be performing while playing them, which will be especially pertinent in multi-user and MMO games, will also have a significant impact on the type of network they need to be able to hook into.

THE ISSUE OF LATENCY.

LATENCY LURKS everywhere, and everything adds latency to your gaming experience. It's inherent in the trans-oceanic network cables that carry internet traffic around the world, it's in the wifi signal in your home and the device and display controllers in your favourite gaming machine, as well as a million other places. It's no wonder then that when it comes to being able to deliver high quality gaming experiences via the cloud it's the topic at the top of everyone's minds and the question on every analysts lips.

When streaming companies talk about latency the most common law they talk about is the speed of light and the fact that, try as they might, they just can't change or beat it - something that's going to remain a fact until society reaches a point where we can communicate via Gamma waves, which, according to friendly physicists are the only things that go faster than light speed.

Until that time comes companies are going to have to find alternative ways to deliver the high quality experiences gamers expect, and fortunately there are

technologies and methods that can help, so let's dive into them.

ONE LAW TRUMPS ANOTHER

When it comes to the future the law that has remained true for decades now, and that is accelerating, is the law that states in the future everything gets faster, cheaper and more ubiquitous - all of which helps us overcome the latency issue.

The compute and network resources in our datacenters will get faster, and in time the silicon circuits that shunt electrons around our computer networks and infrastructure will give way to photonic circuits that operate at the speed of light. Our trans-oceanic terabit cables will get fatter and faster, and be augmented by new space based satellite communications platforms. 4G and LTE mobile networks will give way to 5G and then 6G, and then WiFi and WiFi 6 in our homes will be usurped by WiGiG.

DEDICATED NETWORKS

One of the ways that cloud game streaming companies are trying to reduce latency is by building out their own dedicated global network infrastructure, and boy are these projects big with most companies in the space spending billions of dollars a year building out and developing an intricate web of datacenters, edge nodes, and points of presence that span the globe, touch every country, and reach every gamer wherever they may be. But while having your very own dedicated high performance network infrastructure is one thing, it's only one battle in the war to defeat latency.

NEGATIVE LATENCY

When it comes to discussing the issue of latency and how best to overcome it there's an oxymoron in the market as companies around the world increasingly discuss how best to achieve a phenomenon known as "Negative Latency," which, if implemented correctly, means that cloud streaming services will be lower latency and more responsive

Notes:

than even the most powerful local games consoles.

However, and here's the fun part, there's no such thing as negative latency, and furthermore it's impossible. So are these companies all stuck in some fantasy land or do they know something that we mere mortals don't?

Well, and here's the confusing bit, yes they are all stuck in a magical fantasy land called "Marketing World," but at the same time negative latency is kinda real, and it's also kinda divisive.


Negative latency is the concept of reducing game latency by using Artificial Intelligence (AI) to predict a player's actions milliseconds ahead of time so that cloud gaming servers can pre-generate and pre-load the next sequences, and it isn't a new concept. In fact while implementations vary it's actually quite common practice around the world, for example, in the enterprise storage world where companies try to predict the batches of files that users will want to access next and pre-load them based on previous behaviours. And so on.

While it is increasingly possible to use AI, among other techniques, to predict complex and dynamic player behaviours though to reduce or even eliminate

latency, something that every gamer wishes for, the gaming community are obviously divided on the concept of using AI to predict what a player's next actions will be because, for example, depending on the situation it could give one player an edge over another in say a battle or a race.

Despite these concerns though one thing will be certain, and that's that AI will play an increasingly central role in both the gaming experience and helping tackle and defeat the latency beast, so let the battles rage on.

Notes:



A PEEK UNDER THE DATACENTER HOOD.

WHILE THERE are many different hardware and software components that all have to come together and work in harmony in order to deliver a seamless gaming experience at a high level at least the majority of companies are simply embedding powerful gaming hardware into their cloud datacenters and then using a mix of hardware accelerators and software to overcome the challenges of streaming games from the cloud.

That said though the size of the challenge to overcome varies depending on the screen size players are using because the smaller the screen the smaller the artifacts are that need to be rendered.

THE HARDWARE STACK

When we look under the hood at the hardware that's actually running these platforms you'll be forgiven for thinking that it very closely resembles the hardware found in traditional gaming consoles, and that's not by accident - CPU, GPU, memory, and storage are all obviously still present. But that's where

many of the similarities end because increasingly cloud datacenter hardware is custom designed, and highly tailored and optimised to ensure the services using them meet their service level expectations while all the time balancing resource demand, including energy usage, and managing costs.

As we look forwards in time increasingly more of the CPU's that support these cloud game streaming services will be embedded with intelligence, and, even though Moore's Law is slowing down new 2.5D and even new chiplet architectures could be bought in by the big hyperscale providers in order to help them accelerate specific tasks and workloads that are important to the successful processing and delivery of consumer's gaming experiences.

While Moore's Law might be struggling though when it comes to trying to fit more and more transistors into 7nm and 5nm packages there's no such problem for GPU's which still have a huge amount of headroom left to play with, and GPU's are the real workhorses of the cloud gaming world which gives all of

the companies involved a great level of comfort when it comes to their being able to meet gamers increasingly insatiable demand for higher resolution game experiences.

THE MAGIC OF ELASTIC COMPUTING

Elastic compute is a software based technology capability that's relatively unique to modern cloud environments, and it will play what's arguably the starring role when it comes to making sure that players receive the right user experience at the right time.

Today elastic compute is the technology that literally makes the cloud work, both commercially and technically, because it allows cloud based workloads, from databases to games and everything in between, balance their needs and access the compute resources they need when they need them on demand. It's also a capability that, try as they might, traditional consoles will always have problems leveraging because of the way they're architected.

Notes:

In order to leverage the full benefits of elastic compute for streaming purposes though game developers will need to ensure that their game logics are designed to share multiple compute instances, and are able to balance those workloads across instances that could be thousands of miles apart.

All that said though, when it comes to being able to leverage elastic compute for rendering purposes, such as ray tracing and shaders, that will remain a challenge for a long time, as will the ability to balance new Artificial Intelligence based developments, that include everything from the emergence of AI based entities through to AI's that are creating games procedurally on the fly - all of which are increasingly making their presence felt in the gaming community.

SOFTWARE STACK

The vast majority of software powering these streaming services will continue to be based on open source debuggers, drivers, operating systems, and tools - a hyperscale datacenter trend that is accelerating as companies, who are running tens of millions of systems, try to reduce their commercial costs while still being able to develop feature rich software that meets all of their

performance, scalability, and security needs.

As a result, when we consider the future software stacks of these platforms there can be no doubt that they will continue to be based on a Linux foundation and will lean heavily on tools including DirectX Shader Compiler, LLVM, and Vulkan.

Furthermore, this open source approach also interplays nicely with the game development community's strong history of collaboration, innovation, and shared gains, and means that games developed using these stacks can then be distributed via other platforms such as Steam broadening the games, and by association the studios, appeal and reach.

Notes:

LATENCY

Perfect latency at a very low bit rate

VERSUS

QUALITY

Perfect quality at a very low frame rate

THE ETERNAL BATTLE

As game developers create increasingly high resolution immersive game experiences game cloud streaming companies will have to work hard to find the perfect balance between latency and quality, especially in a world full of disparate fragmented devices and networks, and they'll do that by relying on innovation and new technology developments.

NEW RENDERING TECHNIQUES.

HAVING A blazingly fast connection speed and ultra low latencies are all very well and good but that won't be enough for gamers who are demanding increasingly rich and immersive experiences, whether that's using traditional screens or a combination of Augmented, Mixed, or Virtual Reality gaming, and this brings us to the greatest game streaming battle of all - the battle of delivering gamers with the lowest latencies while delivering the best gaming quality.

It's also an issue that will continue to keep cloud gaming executives awake for a long time as they try to devise new ways to stream even the most demanding game experiences from the cloud to the network's edge at resolutions, for example of up to 16K, that will make even the latest 8K releases look poxy by comparison.

The king of all rendering technologies is obviously photorealistic real-time rendering where gaming platforms are able to render complex high fidelity and high resolution interactive environments in real time, and while the Unreal engine

is getting closer to perfecting real-time rendering at increasingly impressive resolutions the fact remains that for the bulk of non AAA games developers they still have to use a combination of animation, final renders, and VFX to develop their games.

However, as games move into the cloud new tools are emerging, and these existing tools are moving too which opens up a whole new world of cloud based development opportunities.


FOVEATED RENDERING

While there are many rendering technologies available today one of the most exciting developments in the field, which will play an especially central role in Virtual Reality (VR) gaming experiences because of the way games are played and rendered, is Foveated Rendering.

While the technique, which uses an eye tracker integrated into a VR headset to reduce the amount of rendering that needs to be performed by greatly

reducing the image quality in the players peripheral vision zone, isn't a new concept, breakthroughs in Artificial Intelligence (AI) have improved it to the point where companies are now able to deliver 1080p and 4K experiences with up to 90 percent less network bandwidth. And that's game changing - literally.

Furthermore, while the technique for now is mainly constrained to VR environments there's no reason why in the future companies couldn't develop eye tracking software that uses the cameras in your devices to create the same effect outside of VR.



THE IMPORTANCE OF CODECS.

Codecs are the oxygen of the streaming media market, so no codecs, no streaming media. From shooting video to editing and encoding streaming media files for broadcast codecs are involved every step of the way, but as the way we create and distribute content changes the codec market will change too and, as always, standards will come and go.

Furthermore, as the type of information that needs to be encoded in the game streams also evolves, for example in response to the emergence of new end user devices and technologies, such as digital sense systems, haptics, and even one day neuro-tech, game developers will increasingly need to incorporate an increasingly wide range of sensory information into their game streams so that, for example, gamers can physically feel the hit from a laser rifle round and a punch to the gut. And while this is something that today very few game studios are exploring inevitably over time more and more studios will investigate new ways to make the gaming experience even more engaging and intense.

Also, as the Creative Machines I discussed earlier change the way game studios generate content, enabling us to move from using predominantly human based production techniques and mixed human-machine production techniques to fully autonomous ones that leverage different technology combinations, such as 5G, AI, and GPU's in time we'll also see a fundamental shift in the protocols of entertainment delivery.

In short, in time, when all of these things are taken into account, it's highly likely that we will move from the closed streaming video codecs that many of today's cloud game streaming services use today, such as VP9 and traditional video codecs, and embrace the more open object based protocols associated with games and the production of synthetic content in all its forms.

Imagine, therefore, being able to play a game like Halo in virtual reality while wearing a haptic suit - not only will these streaming platforms need to encode audio, objects, and sound, but they will also have to try to encode and transmit sensory information that ties into the

game, so when someone gets hit the information to re-create that hit using the gamers haptic suit also needs to be included.

CODECS WITH A TWIST

While codecs have gotten more advanced over the years one of the frontiers that they haven't quite crossed yet is the ability to leverage the first generation of neural network chips that will increasingly be embedded into all of our devices and Artificial Intelligence (AI) in general, which includes the same Generative Adversarial Networks (GANs) that we're using to create increasingly advanced and sophisticated synthetic content as discussed in earlier chapters, to perform some of the downscaling and upscaling duties.

As a result, over time, again, it is inevitable that companies will increasingly rely on AI to do alot of the heavy lifting and we're already beginning to see the start of this new trend as companies understanding of neural networks especially improves and

Notes:

as the amount of computing resources they consume falls.

So far this has culminated in the production next generation adaptive AI codecs such as Samsung's AIScaleNet that can, for example, adapt to changes in bandwidth and optimise the streaming quality of Over The Top (OTT) services by compressing and downscaling 8K content to 4K quality, transmitting it, and then using AI to upscale the content back to 8K quality.

Furthermore, when you consider the tremendous gains that other industries have seen when using AI, such as the automotive, financial services, and healthcare industries, this will just be the beginning so it's not going to be too much of a leap to say that AI will be a game changer in this field aswell.

Notes:

FUTURE ARENAS AND EXPERIENCES



WHEN WE discuss the future of arenas, the places where people as well as future machines and digital entities go to play and compete with one another, even our human imagination isn't going to be the limit because increasingly these arenas, whether they're physical, mixed, or virtual, are going to be ideated not just by humans, but also by increasingly imaginative and innovative Creative Machines - the same machines we're seeing emerge today and that I'll discuss in more detail in the following chapter.

Furthermore, as increasingly powerful exponential manufacturing technologies continue to arrive our ability not to just create and imagine digital arenas will improve over time, but our ability to quickly build, adapt and even morph physical arenas, whatever their form and wherever they are, will also exponentially improve over time as well.

In short, if it can be imagined, we'll be able to bring it to life, and that includes everything from the places and entities and even entirely new digital and biological lifeforms. And when these two worlds, the physical and digital worlds merge, which they inevitably will, then not only will our experiences become more intense but our relationship with reality will change too.

All this then leaves us with just one question: What happens when our imagination is no longer the limit?



Virtual Arenas delivering Cyber-Physical experiences



Mixed Reality Arenas delivering Cyber-Physical experiences



Physical Arenas delivering Physical only experiences



High Intensity



Medium Intensity



Low Intensity

BREAKING REALITY ...

Increasingly our concept of reality is going to change as our physical and digital worlds continue to merge in new and extraordinary ways ...

... AND SWITCHING UP INTENSITY

... But that's not the only thing that is going to change in the future. The intensity of the experiences in all of these different arenas and realities will also increase until they are on a par with one another.

EVERY MIND AND WORLD IS YOUR PLAYGROUND.

HAVE YOU ever heard of the saying "The world is your playground?" Well, as we race headlong into the future we're going to have to amend that saying to include the digital realm, and switch out the singular word "world" and replace it with "worlds," because where we're headed "Every world is your playground." And that's awesome. Furthermore though, as we see neural interfaces rise to prominence in the future we can also see a day when someone else's mind could be your playground too - just strap yourself into your fully immersive gaming rig and jack into the playscape of someone else's consciousness.

Today we already know that everything and anything can be an arena, from our back yards to our kitchen tables, from the largest stadiums in the world to everything above and below the open skies, with the primary differences between them all being the space, spectator capacity, and experiences they deliver.

As the technologies we develop continue to advance and mature though the

sheer variety of arenas that we can create and experience will explode, and physical, mixed reality, and virtual reality arenas, and any combinations thereof, will become our new playgrounds. Furthermore, technology advances will eliminate the boundaries between them all, and we'll be able to push the limits to unexpected new highs.

FUTURE ARENAS

Ostensibly, just as there is today, in the future there will be three main types of arenas, namely physical reality, mixed reality, and virtual reality arenas, and in time they will all increasingly merge together in new and unexpected ways - all of which I'll discuss in more detail later in this chapter.

Built in new ways

However, while the arenas and the experiences they drive are evolving so are the tools and technologies that we'll use to ideate and create them, and in many ways these are just as

extraordinary and fascinating as the arenas themselves. So much so that I've given this particular topic its very own chapter.

FUTURE EXPERIENCES

When we discuss future experiences there are two primary perspectives that I'll discuss in detail - the perspective of the athlete, and the perspectives of the spectators.

Needless to say technology is going to have a big impact on both of these areas, but as different technologies arrive and mature the other impact that they're going to have is on the user experience within each type of arena, after all everyone has their own tempo and every moment has its own intensity, and of course both of these vary depending on the activity we're participating in and our mood and motivations at the time, but what if you could have just as an intense experience in the virtual arenas as you could in physical ones? And this is an area I'll also go into in more detail in this chapter.



BYTES STUCK BEHIND THE GLASS.

IN TIME it is inevitable that the digital and physical worlds will merge until they reach a point where it's difficult to distinguish which is which, and in time the arenas that this cyber-physical union will create will go way beyond our wildest dreams and help us create and experience arenas that are as amazing as they are immersive.

For now though all of this future promise starts here, in an entirely different type of arena in the world of gaming and E-Sports where the arenas are still locked behind glass.

As we continue to see gaming's rise to prominence on the global stage, whether it's thanks to the emergence of cloud game streaming services or E-Sports, or both, and the introduction of new technologies, such as haptics, neural interfaces, and VR, to name but three of many, it's inevitable that these arenas and experiences will one day be able to leap out from behind the glass and that we'll be able to experience and interact with them in new ways.

In the next couple of decades the arenas

that gaming and E-Sports fans experience today will start their own transformations in earnest.

First we'll see the continued ingress into the space of haptics and wireless virtual reality rigs which will be the way that many gamers get a taste of their first "real" cyber-physical experiences.

In time, as virtual gaming inevitably takes off, we'll then begin to see the introduction of what I'll call "Constructs" - purpose built physical arenas that are augmented with the latest technologies, from electro-magnetic floors that let gamers run marathons while staying in the same spot, through to AI and machine vision systems that help the organisers replace traditional motion capture rigs and let them seamlessly manage the spaces and gamers as they do battle with one another in virtual worlds.

Then, finally, in the latter half of this timeline we'll see the cyber-physical experience ratchet up and become even more intense as increasingly powerful neural interface systems, that not only let

2020 - 2040

gamers explore and interact with their virtual worlds using just their minds, but that also let future Creative Machines, like the ones I've discussed, read their thoughts and construct new procedural game content on the fly, take a more prominent role - at which point we then move onto the next timeline and go even deeper into the rabbit hole.



MIX IT UP AND GHOST ME.

WHEN WE look into the future of Augmented Reality and Mixed Reality arenas and their role in this world I have to say they kind of feel like a halfway house - especially when you compare them against the types of physical, cyber-physical, and virtual arenas we'll be able to experience in the future.

While Augmented Reality, for example, undoubtedly has its merits the biggest draw back of the technology today is that it's predominantly screen based so while that's all well and good for playing single-player games and turning your kitchen table into a Angry Birds arena, when it comes to playing single-player or multi-player sports, let alone sports that require you to have one or more free hands and a free range of movement it's very difficult to see just how AR in its current form will scale and become invisible enough of a technology to provide users with the sporting experiences they crave.

After all, could you imagine playing even a simple game of virtual catch or tennis in your living room with a friend while

waving your smartphone or tablet about? No, neither can I, and as interesting and as promising as AR is these two examples almost perfectly explain why the current technology is flawed.

However, fear not, all is not lost. Once we flip out AR on a smartphone for AR in a pair of smart contact lenses, or those rather cringe worthy albeit interesting glasses, well that screen locked experience disappears and suddenly AR is set free and the arenas gets a whole lot more interesting and interactive. And as an example, now imagine playing in the park hurling virtual fireballs at each other and suddenly you'll appreciate the completely different type of experience between the two examples. And now that AR is free our imaginations can get busy again.

As our minds wander now thinking about the new arenas we can create, and then eventually sports which I will come to in a later chapter, the world is literally our playground, and when we bring Mixed Reality into the game not only do we have a way to interact with these arenas using AR, but we also have a way to

2020 - 2040

morph and bend them and everything in them to our wills. At which point the only question that's left is what arena do you want to create?

Moving one step on though, just as the digital world is increasingly able to merge with and invade the physical world, the physical world can also invade our digital worlds and this is where the incredibly interesting concept of Ghosting now comes in that allows you to virtually compete and play off against anyone anywhere in anything.

Imagine, for example, being able to sprint against a virtual Usain Bolt from 2008 as he smashes it at the Beijing Olympic Games, swim against Michael Phelps, or being able to pit yourself against an entire sports team in real time during live events in a cyber-physical battle royale. And if you want to do more than just take part and feel the

Notes:

experiences, well just strap on that haptic second skin.

While creating and initiating this kind of cyber-physical play off requires a mix of equipment, for example HD cameras at one or both ends if you want a two way competition, Video to Digital AI generators, codecs, and powerful networks that can stream those digitised experiences from one user to another to and from your smart contact lenses the fact of the matter is that those technologies are already here.

All this leaves then you with one major problem during this timeline - you're going to either have to stop bragging about how much faster or better you would be if that was you in the game, or get your game face on and prove it in the virtual arena.

Furthermore, as we consider how these new forms of arenas could impact and merge with the E-Sports market, well, who's not to say you aren't as good as you say and that a scout won't sign you up? In which case don't forget to cut the futurist a cheque. Thanks.

Notes:

THE SIMULATION BECOMES REAL.

THE ONE thing that true virtual reality enthusiasts crave is immersion, full immersion, and the thrill of being mentally and physically teleported to any time and place they can imagine. And where we're going in the long run they're not only going to get their wish, they're going to get much more than they bargained for.

Let's face it most of today's VR experiences are lame, sorry, I mean tame, by future standards. The equipment is bulky, the resolutions are low, there are gaps in the experiences, and unless you've got a sweet rig then you're stuck in one place and then latency creeps in and ruins your timing. And as for content, well, several studios shut down recently. So you have to admit that while the concept of VR might be intriguing the experience isn't as compelling as it should be. Yet.

But by now you should be getting the picture. This isn't what the future looks like, not by a long shot, so let's dive into that big virtual ocean that is the future and explore the art of the possible.

While getting to the point where users believe that the simulation is the real world as they do in the Matrix will take time, and in the interim most of the effort in this field is going to be focused on miniaturising the tech involved, such as the bulky VR headgear and the bulky haptic clothing, and experimenting with creating new products that enhance the user experience and increasingly tie all of our five senses, not just the sense of sight, sound, and touch, into these worlds. But as developers set forth to conquer all of these challenges elsewhere others will increasingly focus on an entirely different problem - that of motion.

After all, being fully immersed in a virtual environment is one thing, but if you can't walk or run around it as you would in the real world then, well, you're still getting only part of the experience even if you can smell the cordite in the air and feel the sweat on your brow and taste the copper in your mouth. And this is where the floors we stand on, whether it be in an arena or at home are an issue - they're just too static. As a result several companies are now developing electro-

2020 - 2040

magnetic floors that, in essence, move under you and move with you, following the world and your movements wherever it is your feet want to take you.

It goes without saying that all the developments across all these different areas will play a vital role in helping us create increasingly realistic and intense virtual environments that feel increasingly real and immersive to us. And as all these technologies mature, as the clunky headsets first become wireless, then miniaturise, then become brain controlled, and then towards the end of this timeline disappear altogether as we see the emergence of screenless display systems that beam VR content directly into your eyes, and as the bulky haptic suits become lycra thin then nano thin and essentially become nothing more than a second skin, we'll be well on our way to creating the ultimate magic act and making the illusion real.

Notes:

Over time as all these developments become increasingly mature and commercialised we'll also see giant leaps in the way content is created and delivered.

Today's relatively manual content creation techniques and workflows will become increasingly automated, and then we'll see the Creative Machines I've discussed assume more of the load, generating 4K, then 8K, then 11K, then 16K synthetic content faster than we could have using traditional techniques, and then as these Creative Machines start moving into the cloud, improve and mature, increasingly we'll reach the point where they can read a user's behavioural and biometric cues and data to generate procedural content on the fly that's tailored to their individual point in time preferences.

Bringing all of this together it's increasingly easy to see therefore how we progress from today's rather tame virtual experiences to a world where the headsets are gone, where haptic suits become as invisible as our own skin, and where we can experience and interact with high definition worlds using nothing but our thoughts and in very much the same way we interact with the real world. And then of course, as we push

the timeline further out, we go deeper down the rabbit hole.

As I close this timeline out there is a wrinkle though, and it's an interesting one. So far I've discussed improving the quality of the virtual world and our experience with it, but what if we could mix things up and turn our real world, our actual real world, into a virtual experience, a virtual experience where, as you walk down the street, for example, an AI renders a completely new virtual world over the top of the real one, turns office blocks into trees or dinosaurs, or turns cars into giant candies?

While this kind of technology still leverages the same VR technologies I've just talked about hopefully you can see the differences in the approach, and when it comes to building new types of arenas, well, frankly, how many options do you want!?

Notes:

THE FUTURE OF GAME PRODUCTION



THERE ARE three types of games experiences today, namely physical, virtual, and mixed, and the same will be true in the future. But in the future the way these games are produced, distributed, consumed, and played, will all change dramatically and be much more intense.

Just as it is inevitable that cloud game streaming, as I've discussed, will become the dominant way we consume games in the future inevitably game production itself will undergo the same shift, but with a couple of extra plot twists.

Firstly, the tools that game studios use today will increasingly be moved into the cloud, but once there in time they will at first be augmented by new Artificial Intelligence (AI) toolsets that help studios produce higher quality content more efficiently, and then increasingly the entire game production process will be automated and managed by Creative Machines that will devise, create, and then procedurally generate their own synthetic content and synthetic games on the fly in response to gamers real time feedback including the analysis of live camera feeds, user engagement and response times, and much more.



Notes:

ONE OF, if not the greatest challenge, that researchers developing synthetic content generators face is that of being able to get their creations to create content that passes the famous Uncanny Valley test - the hypothesised relationship between the degree of an object's resemblance to a human being and the emotional response it generates. In short, the machine's ability to generate content that is good enough into fooling people that it's the real deal.

ACHIEVING IMPERFECTION

Ironically, however, in the past most researchers in the space did their best to develop creative machines that generated flawless output, but as they quickly found out, even though the outputs were world class the critics were still able to quickly identify the fakes because they were, frankly, too perfect. And life, especially humans, aren't perfect.

As a result one of the greatest turning points in generating passable content, besides from the amazing advances in the technologies, tools and techniques, was a shift in researchers attitude and their sudden determination to introduce subtle imperfections into their work.

One of the best examples of imperfection at work is human speech. For example, while early synthetic speech models were designed to be perfect in both their flow, pronunciation, tone, rhythm, and vocabulary, the fact of the matter is that no human alive today speaks perfectly. Our speech is imperfect and our conversations are filled with stutters, slightly random tone and rhythm changes, pauses, and Umm's and Ahh's.

As a case in point when Google first demonstrated Duplex, now part of Google Assistant, on stage at their I/O conference it was the introduction of these imperfections that amazed and fooled the people listening in the audience into thinking they were listening to a real person booking appointments. And synthetic voice isn't the only type of content that is now benefiting from this lack of perfection, we're seeing it emerge in imagery and other content.


All that said though there are naturally times when you want your content to lack imperfections, such as music performances, and other types of content, and now that researchers have figured out how to include imperfections in their work their next challenge is to be able to turn it on and off, like a switch, at will, and embed enough intelligence into their creative machines so that those

Notes:

machines understand when, and when not to go off the grid.

Notes:

SYNTHETIC AUDIO



EXPLORE MORE. Click or scan me to learn more about this emerging tech.



Notes:

OVER THE years there have been many attempts to accurately synthesise the nuances of different sounds as well as the human voice, with plenty of enterprises finding it difficult enough to create high fidelity creative machines capable of realistically synthesising just a single voice - let alone systems that are accurate and responsive enough to fool people into thinking they are in a genuine moment listening to a wide range of real audio streams.

However, as this particular field becomes table stakes rather than just a nice to have, and as we continue to see the rise and adoption of Behavioural Computing platforms, including Digital Personal Assistants, enterprises around the world have significantly stepped up their efforts to conquer it. And, supported by increasingly capable and powerful exponential technologies they're now making significant headway on a wide range of fronts.


One of the most significant breakthroughs in this space in recent years has been enterprises experimentation with Deep Learning and Wavenets which, after a significant amount of testing and development, have now largely helped researchers in the field conquer one of the fields greatest challenges - dynamically synthesising raw audio's

extremely fast tick rate which, for human speech, for example, is over 16,000 samples per second and where each sample is influenced by all of the previous ones. And now that some of these foundational technologies and tools continue are maturing it becomes easier to adapt them to synthesise an increasingly wide range of sounds, from speech to music and beyond.

However, as we look further out, it could be argued that this challenge will pale in comparison to creating the world's first truly intelligent so called Conversational AI's that not only synthesise realistic human-like audio but will also be able to have expansive and wide ranging conversations with us.

FUTURE HISTORY ...

A VISUAL HISTORY OF THE FUTURE



2017 saw the first basic but passable synthetic voice generator emerge.



2018 saw WaveNet and new Deep Learning models emerge that were capable of synthesising the human voice with uncanny accuracy, and new basic multi-discipline synthetic music generators and high fidelity Text to Audio generators emerge.



2019 saw the live demo of the first synthetic human voice generator that passed uncanny valley, more synthetic music and voice generators emerge, that near uncanny valley, and the first ever commercial record deal for an AI.

Commercialisation begins.



In 2020 we will see the emergence of the first Conversational AI platforms complete with synthetic voices that are able to pass uncanny valley. However, the overall development direction of the technology will be driven by tech, not content, companies. And open source models will accelerate the development of the overall field.



By 2025 this technology will be able to perfectly copy and mimic any individual and be capable of synthesising complex speech complete with accents and emotions.

It will be moderately commercialised and regulators will have started stepping up their scrutiny of the field.

By 2030 this technology will be mature and fully integrated into studio workflows as well as consumer applications. It will be fully commercialised.



By 2035 the technology will be fully adopted and ubiquitous, and its impact on society will be the source of great debate.

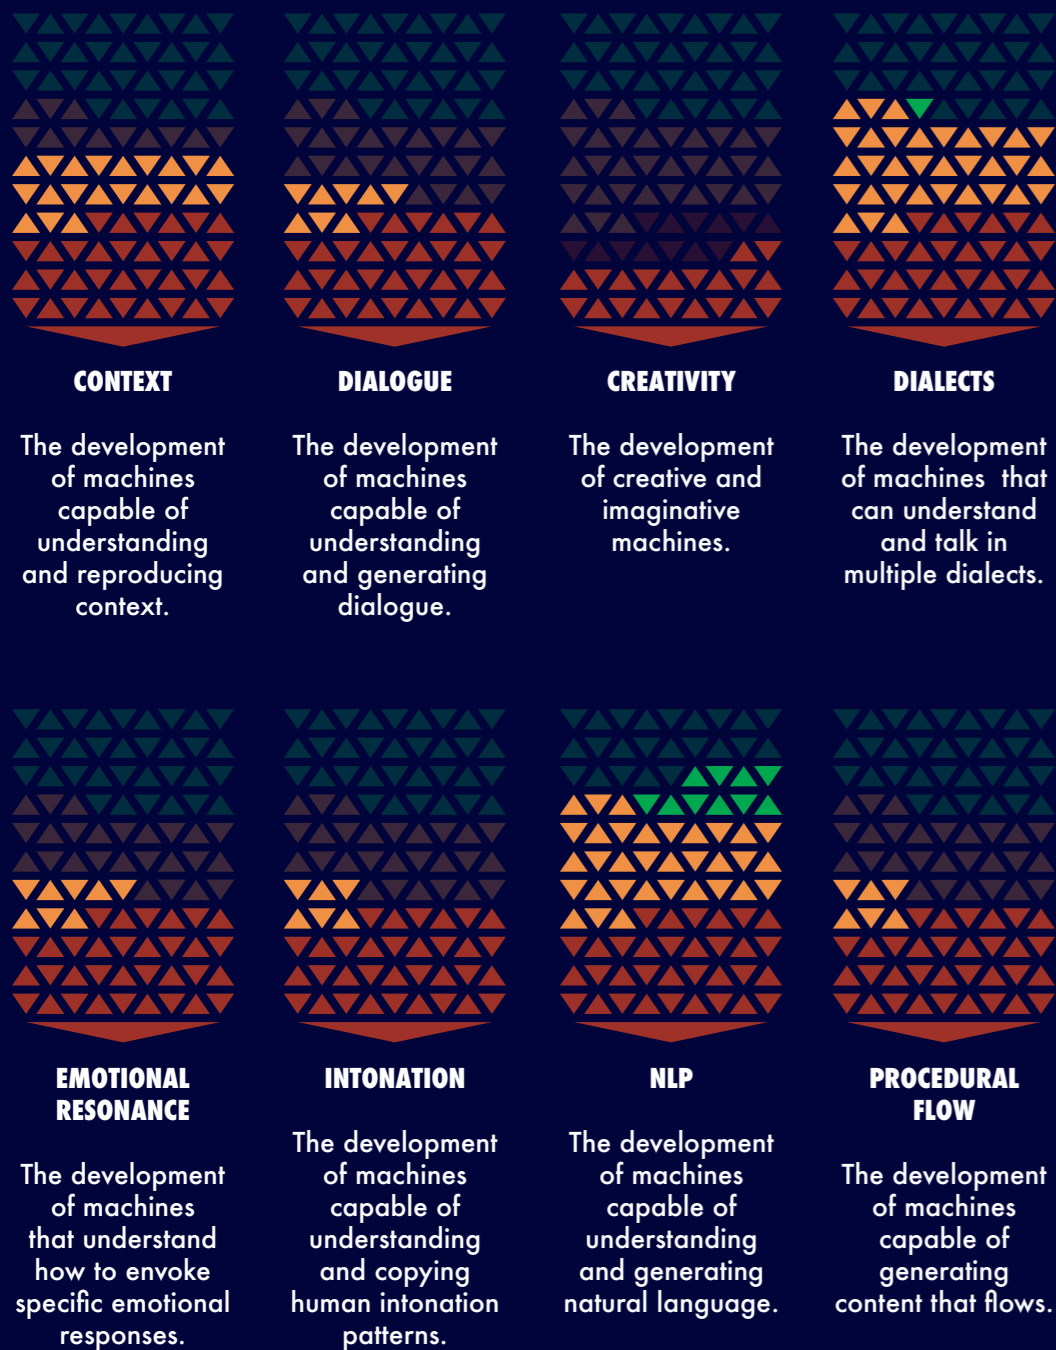


Figure 4. Current state of synthetic audio development by AI discipline.

Source: Multiple Sources

THE WORK STILL TO DO

While many of us might think that creating a machine that's capable of spitting out decent sound, in whatever form it takes, is easy, you can clearly see that it isn't.

Furthermore, reaching the point where researchers are finally able to create a single universal and generalised synthetic audio generator that can be used across a wide variety of use cases relies on developments not in any one single area, but, as you can see from the adjacent discipline areas, many.

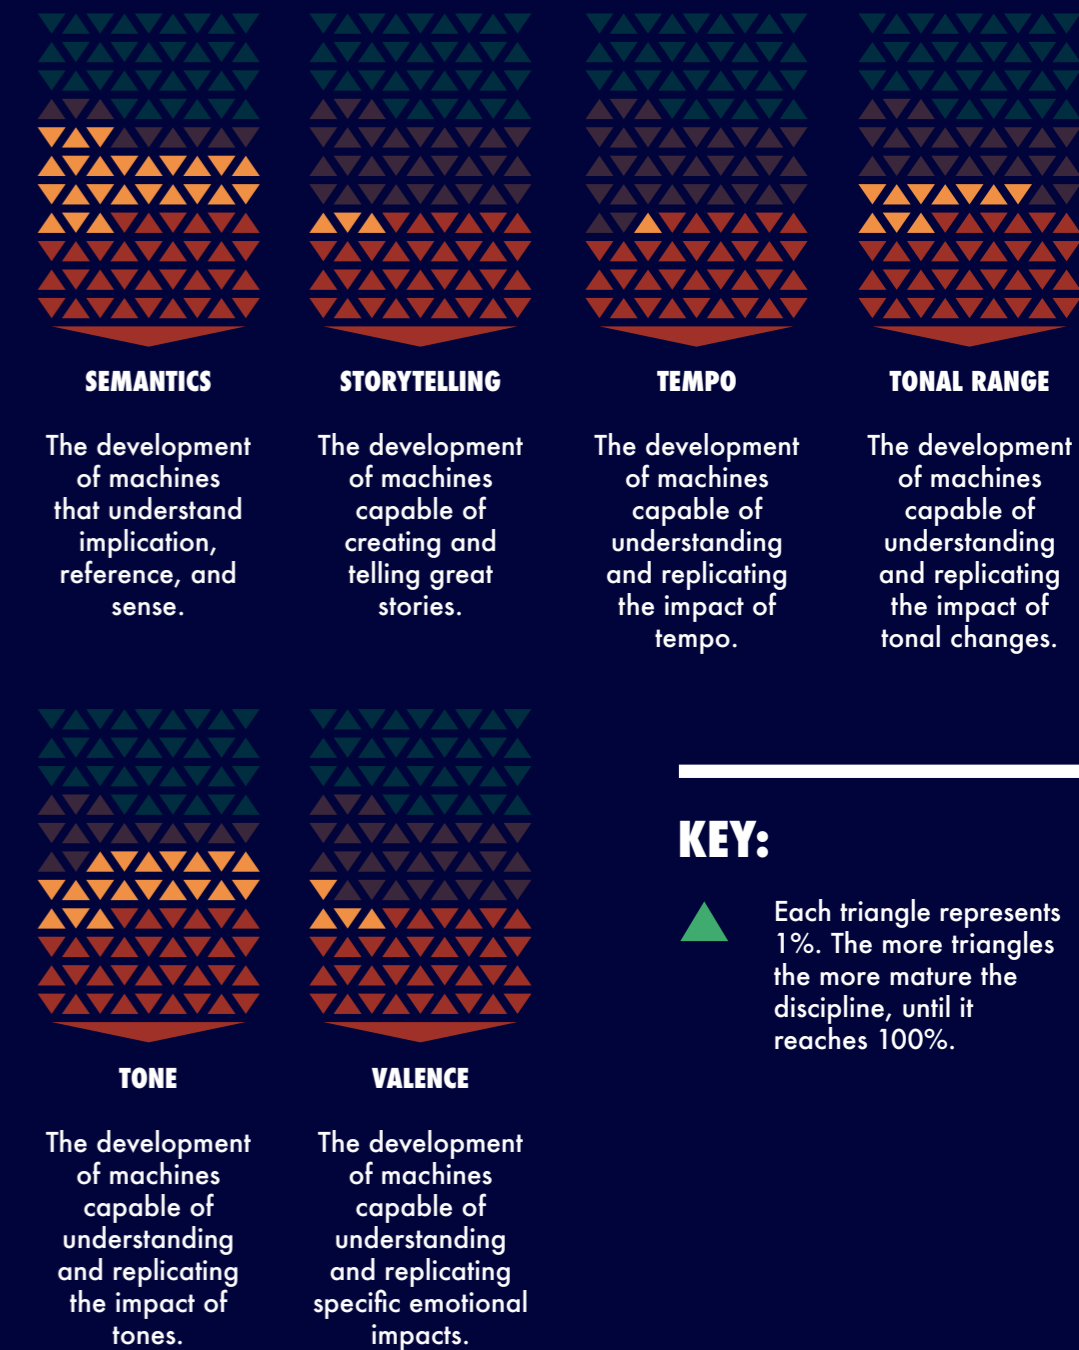
For example, while there are an increasing number of creative machines capable of generating synthetic music, with a couple of them even being signed by a major record label, trying to create music that people are willing to pay for, music that strikes an emotional chord and pulls at the heart strings, relies on a wide range of development areas - including, but not limited to the machine's ability to understand context, emotion, and semantic meaning, and then, at the least, translate those into compositions that have an appropriate flow, rhythm, and tonal qualities.

Even though there have been significant breakthroughs in the field over the past

year, let alone years, as I've discussed, there is still clearly a lot to be done before we get to the point where we have the ultimate creative machine in this field - a machine, for example, where you can simply push a metaphorical button and have it churn out chart topping hit after chart topping hit. Which then brings me on to another point.

As we look at all of the synthetic audio platforms being developed today they are all being developed in what's called a "narrow" way. In other words, a debating machine will one day be good enough to hold its own against the world's top human debating teams, but ask that same machine to create a podcast or book a calendar appointment and it'll fall over faster than a one legged flamingo in a hurricane.

However, as time progresses all of the capabilities of these narrow AI systems will be combined beneath one umbrella - an Artificial General Intelligence architecture that means that eventually one machine will be able to do it all, generate chart topping hits and generate catchy marketing content, and help you translate mandarin Chinese into Gaelic while at the same time it's advising you on which mortgage to buy.



KEY:

▲ Each triangle represents 1%. The more triangles the more mature the discipline, until it reaches 100%.

Notes:

TREND IMPACT AND TRAJECTORY

This trend will have a wide range of impacts including on the following market categories:

- Behavioural computing
- Debating
- Digital assistants
- Marketing content
- Music
- Podcasts
- Robo-Advisory
- Soundtracks
- Spoken word
- Universal translators

The impact of synthetic audio on global culture, industry, and society, needless to say, will be wide ranging, leading to everything from the development of new interesting use cases such as being able to advise and converse with your co-workers and family from beyond the grave, all the way through to job losses across a wide range of career categories, from call center operators, sales people, and voice over artists to musicians.

Meanwhile from a trajectory perspective as with all of the synthetic content generators we're discussing researchers will continue to develop new machine

models, iron out the bugs and the kinks, and eventually those machines will leave the labs, be adopted by enterprises and integrated with other products and services, and be adopted by consumers - while the regulators play catch up.

TECHNOLOGIES BEHIND THE TREND

Over the past few years the number of enterprises developing creative machines capable of generating synthetic audio, that includes everything from music to the multi-lingual spoken word, has increased steadily, and as the development of these machines advances so too does the range of audio content they can generate. However, the technologies and tools researchers use, and the techniques they use to develop and train their models depends on the particular field they're interested in.

The primary areas of research in this field today include:

- Conversational AI
- Music generation
- Natural language translation
- Speech synthesis
- Text to Music generation
- Voice mimicry

Notes:

1. Audio Mimicry

In a world where we can increasingly clone everything from livestock and pets it's only natural that somewhere there are researchers that want to clone people's voices - with all of its obvious pros and cons. Over the past couple of years there has been a tremendous amount of focused development in this space, much of it focused on one goal - to use AI to create the most realistic artificial voices in the world, whether those voices are generated from scratch, as I'll discuss in the Spoken Word section below, or, in this case, to mimic the voices of real human beings. And the latter is hard, but for the enterprises that can pull it off they'll have a system that can get anyone to say anything.

Every individual's voice has its own patterns and nuances, and sampling someone's voice can now be done in as little as a minute from public video footage, after which a mixture of different deep learning algorithms disassemble the audio into its constituent components and analyse it based on accent, cadence, pitch, speed and more. And as these algorithms sort they learn, building up their own models.

In many cases developing these AI's, most of which are built in PyTorch and

TensorFlow frameworks and trained on petaflop scale computers in the cloud, and work by comparing every aspect of a recording to every other voice recording in their databases in order to build a digital profile with a few million data points required to build each voice - voices that are now starting to pass the uncanny valley test and even fool people's own families.

2. Music Generators

Generative Adversarial Networks (GANs) are considered a state of the art method for generating everything from synthetic imagery to video. However, while they are a must have in those fields researchers focusing on generating other types of content, such as music, have struggled to apply them, and up until recently their only recourse was to try to adapt the many Autoregressive (AR) AI models such as WaveNets and Transformers that dominate the field because of their ability to predict single sample sounds - something that gives them a distinct edge when it comes to synthesising voices, but something's that makes them painfully serial and slow when generating music. The result of all this is, of course, that traditionally researchers in the field ended up having to augment their research with techniques

Notes:

including distillation or specialized kernels.

Now though researchers, fed up of this painfully slow approach are turning their attention to developing new types of GANs that are capable of generating entire music sequences in parallel, and synthesising audio significantly faster than real time - with some speeds exceeding 50,000 times faster than standard WaveNets.

Unlike traditional WaveNet autoencoders that rely on using a time-distributed latent code these new GANs generate the entire audio clip from a single latent vector, or trait, which makes it a lot easier for them to disentangle global features such as pitch and timbre from all of the other noise. Then, using datasets of musical instrument notes, such as the NSynth dataset, researchers have been able to independently control pitch and timbre using Progressive GAN architectures and upsample the tracks by using convolution generators and combining them with other single vector sounds in order to generate the full music track.

While the results that some of the researchers are pushing out now are impressive, with some music models even being signed by companies including

Time Warner there is still a lot of work to do with some researchers thinking that their work might be further improved still by the development of multi-scale GAN conditioning models that can handle variable length outputs, and by replacing upsampling convolution generators with more flexible differentiable synthesisers.

3. Spoken Word

Trying to create a synthetic speech platform that passes the uncanny valley test is complex, if for no other reason that every one second of human speech contains over 16,000 so called waveforms, and when it comes to developing state of the art platforms that can generate speech from scratch, complete with accents and emotions, one of the leading models in this field is the WaveNet model first developed by Google DeepMind whose second iteration, just a year later was over 1,000 times faster than the original and capable of generating far higher quality audio. What set this second model asides from the first was the development of a new Probability Density Distillation (PDD) technique that allowed it to work in a massively parallel computing environment.

The original WaveNet model used

Notes:

autoregressive connections to synthesise the waveform one sample at a time, with each new sample conditioned on the previous samples, and while this produced high-quality audio with up to 24,000 samples per second its sequential generation is too slow for production environments.

To get around this problem the researchers needed a solution that could generate long sequences of samples all at once with no loss of quality so they developed the PDD technique, where a fully trained WaveNet model taught a smaller "student" model that was better suited to modern parallel computational hardware. This student model is a smaller dilated Convolutional Neural Network (CNN) that was similar to the original WaveNet, but, crucially, the generation of each sample doesn't depend on any of the previously generated samples meaning the researchers could generate the first and last word, and everything in between, at the same time.

During training, the student network starts off in a random state. It is fed random white noise as an input and is tasked with producing a continuous audio waveform as output. The generated waveform is then fed to the trained WaveNet model, which scores each sample, giving the student a signal to

understand how far away it is from the teacher network's desired output, and over time, the student network can be tuned - via backpropagation - to learn what sounds it should produce. Put another way, both the teacher and the student output a probability distribution for the value of each audio sample, and the goal of the training is to minimise the so called Kullback Leibler Divergence between the teacher's distribution and the student's distribution.

Ironically, this new technique has parallels to the set up for Generative Adversarial Networks (GANs), that are used to create other forms of synthetic content, with the student playing the role of generator and the teacher as the discriminator. However, unlike GANs, the student's aim is not to "fool" the teacher but to cooperate and try to match the teacher's performance.

SUMMARY

The human voice is one of the most recognisable sounds on the planet but it is so rich with information that it's incredibly difficult to reproduce, but as researchers continue to push the boundaries and get closer to cracking it it opens up a whole new world of opportunities.



EXPLORE MORE. Click or scan me to learn more about this emerging tech.



Notes:

WE ARE bombarded by imagery every day, whether it's in adverts and online, in comic books and publications, and a myriad of other places, but when it comes to our ability to detect and sort the synthetic imagery from real imagery some categories of synthetic images are easier for us to detect than others.

AS THE quality of synthetic imagery improves researchers will next focus their efforts on creating the technologies and techniques that help control and modify the final results.

The reason for this, of course, is obvious - the more real imagery we are bombarded with throughout our daily lives the easier it is for us to spot the fakes, with prime examples being the human face and human form.

As a result the enterprises that are working to develop creative machines capable of generating imagery that passes the uncanny valley test face a barrage of critics - the public. But as the technology, techniques, and tools improve rapidly it's clear for everyone to see that the platforms and the outputs they generate have improved dramatically in just the past five years

alone, and we are only going to accelerate from here.

As a consequence it won't be too long now before the vast majority of imagery that these systems generate, whether it is simple 2D cartoon like imagery or more complex and contextual 3D imagery, will surpass anything that human creatives can produce.


FUTURE HISTORY ...

A VISUAL HISTORY OF THE FUTURE



2014


2018 saw an increase in GPU computing power, the development of more powerful GAN models, and a dramatic increase in the volume of training data, as a result synthetic imagery now starts surpassing uncanny valley.



2019


In 2020 researchers attention will turn to creating the tools needed to control and modify the outputs of the models. Also, as the corpus of available data used to train models increases the range of imagery the models are able to produce will broaden dramatically.

It will start commercialising.



2025

By 2030 the field will be fully mature, and the only areas left to conquer will be those relating to specialist content types. Most researchers will have now moved their attention to other fields.




2035

2017 saw researchers focus in more on developing synthetic image generators, but the quality was low.



2018


2019 saw researchers increasingly apply their expertise across multiple image domains and develop new 2D to 3D Domain Transfer techniques, as a result more imagery passes uncanny valley, and researchers in the field now start turning their attention to the next problem - creating synthetic video content.



2020

By 2025 the output will be fully controllable and researchers will have turned their attention to developing 3D and dynamic content generators, and integrating the tech with other complimentary tech fields.

The tech will be fully commercialised and will be increasingly ubiquitous.



2030

By 2030 the technology is fully matured, commercialised, and ubiquitous across multiple platforms.

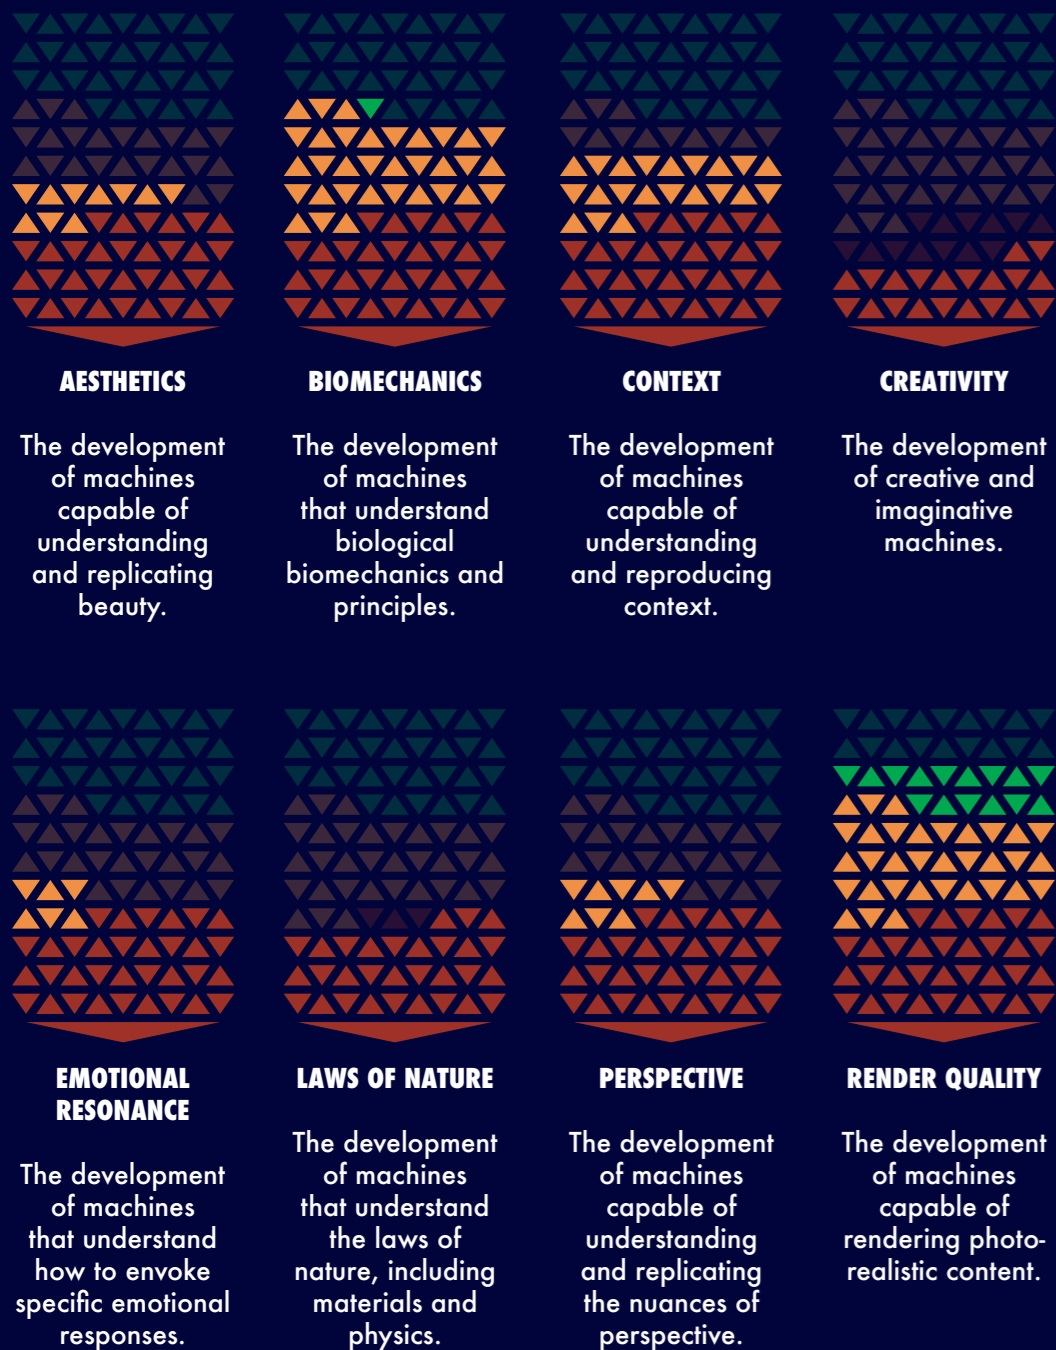


Figure 5. Current state of synthetic imagery development by AI discipline.

Source: Multiple Sources

THE WORK STILL TO DO

Of all the different synthetic content types arguably the easiest type of synthetic content for creative machines to learn how to generate is synthetic imagery, because unlike some of the other synthetic content fields, such as synthetic video, generating good enough synthetic imagery requires a much smaller number of AI disciplines to combine and come together before researchers are able to finally produce a single universal and generalised synthetic imagery generator.

Furthermore, to make things easier there is already a huge corpus of training content for them to learn from, and an increasingly vast pool of computing power that they can draw on.

All that said though there is no getting away from the fact that images convey a large amount of visual information, and because sight is, arguably, our primary sense, if every little detail in the images isn't generated perfectly then people notice - and people can be harsh critics.

As a result there is, arguably, a lot more pressure on the researchers in this field to go past just "good enough" and create state of the art machines that fool even the harshest critics - and those machines are already emerging.

Notes:

TREND IMPACT AND TRAJECTORY

This trend will have a wide range of impacts including on the following market categories:

- Animations
- Avatars
- Artwork
- Stock imagery

The impact of synthetic imagery will mainly be felt on creatives working in the art and photographic industries whose livelihoods today mainly rely on their ability to create and sell a diverse array of content. However, the same technology that takes away will also give non-creatives the ability to create new imagery in new and exciting ways and let them retain creative control while the machines bother with the generation.

Meanwhile from a trajectory perspective as with all of the synthetic content generators we're discussing researchers will continue to develop new machine models, iron out the bugs and the kinks, and eventually those machines will leave the labs, be adopted by enterprises and integrated with other products and services, and be adopted by consumers - while the regulators play catch up.

TECHNOLOGIES BEHIND THE TREND

Overall it is almost obvious to say that this particular field is dominated by Generative Adversarial Networks (GANs), however, while the results that early GANs generated were impressive enough to get people to pause for breath the content they generated was often very difficult to control and modify post production. As a result, and as researchers start creating the next generation of GANs adaptability and usability are now two of the key areas of focus.

The primary areas of research in this field today include:

- Image to Image Domain Transfer
- Image to Image generation
- Image generation and manipulation
- Speech to Image generation
- Text to Image generation

1. Cartoon GANs

Everybody loves animations and cartoons which is why, as you'd expect, some researchers are focused on developing creative machines capable of generating animation content.

Unlike generating traditional imagery

Notes:

though this form of content, such as Japanese Anime, has distinct aesthetics, and today relying on traditional manual transformation techniques to generate real world scenes require considerable expertise and expense as artists must painstakingly draw lines and shade colours by hand in order to create high quality scenes and reproductions.

In order to try to solve this problem in the past researchers in this space turned to Non Photo-realistic Rendering (NPR) as well as more traditional Convolutional Neural Networks (CNN) in order to try to develop working solutions, but both of these techniques proved either too time consuming, and, or, impractical because they didn't produce satisfactory cartoon-like results. After all, different cartoon styles have unique characteristics involving high-level simplification and abstraction, and cartoon images especially tend to have clear edges, smooth colour shading and relatively simple textures, which presented challenges for the Texture-Descriptor based loss functions used in existing methods.

After working through all of the available options most researchers, it appears, are now focusing their energies developing Generative Adversarial Network (GAN) frameworks composed of two CNNs

that enable style translation between two unpaired datasets - a Generator for mapping input images to the cartoon manifold, and a Discriminator for judging whether the image is from the target manifold or synthetic, with residual blocks being introduced to simplify the training process.

And in order to avoid slow convergence and obtain the high quality stylization, dedicated semantic content loss and edge-promoting adversarial loss functions they need they've integrated an initialisation phase into this new, so called "cartoonisation" architecture.

An example of this is shown in a GAN named CartoonGAN that shows the importance of each component - the initialisation phase performs a fast convergence to reconstruct the target manifold; sparse regularisation copes with style differences between cartoon images and real-world photos while retaining original contents, and the adversarial loss function creates the clear edges.

Using this technique both real world photos and cartoon images are used for the initial model training phase while the test data contains only real world pictures.

Notes:

When compared to recently proposed other CNN based image transformation frameworks, such as CycleGAN or Gatys et al's Neural Style Transfer (NST) method, the new technique more successfully reproduces clear edges and smooth shading while accurately retaining the input photo's original content.

However, while the developments so far have been impressive NST only uses a single stylisation reference image for model training which means it cannot deeply learn a particular anime style, especially when there are significant content differences between the stylization reference image and the input images. But, as ever, with more development and training the results will improve.

2. Imaginative Machines

Creating AI's that are imaginative is an incredibly interesting research area, full of both opportunities and mine fields, that has a wide range of applications - especially in the generation of synthetic imagery and video. And, as difficult and as onerous as it might sound to try to create a machine with its own imagination, fortunately the research teams involved in these types of projects

have their own creativity and ingenuity to draw on.

For example, if you're handed a note that asks you to draw a picture of a bird with a yellow body, black wings and a short beak, then chances are you'll start with a rough outline of a bird, then glance back at the note, see the yellow part and reach for a yellow pen to fill in the body, read the note again and reach for a black pen to draw the wings and, after a final check, shorten the beak and define it with a reflective glint. Then, just for completeness, you might sketch a tree branch for the bird to sit on. And now this ability, once just the domain of humans, can now be performed by an AI, and it's already probably better at drawing than you are.

The technology, which the researchers behind the most interesting project in the field, simply call "The drawing bot," can generate images of everything from ordinary pastoral scenes, such as grazing livestock, to the absurd, such as a floating double decker bus. But by far the most interesting aspect of what this particular AI is generating is the fact that each image contains details that are absent from the original text descriptions, indicating that this AI does indeed have its own imagination.

Notes:

At the core of the model is a technology known as a Generative Adversarial Network (GAN). The network consists of two machine learning models, one that generates images from text the descriptions and another, known as a discriminator, that uses text descriptions to judge the authenticity of generated images. The generator then attempts to get fake pictures past the discriminator, and working together, the discriminator rejects the images that don't pass muster and continuously pushes the generator toward perfection.

During their work the researchers trained the drawing bot on datasets that contain paired images and captions, which allowed the model to learn how to match words to the visual representation of those words. The GAN, for example, learns to generate an image of a bird when a caption says bird and, likewise, learns what a picture of a bird should look like.

While GANs work well when generating imagery from simple text descriptions such as a blue bird or an evergreen tree, the researchers found that the model quality stagnated with more complex text descriptions such as a bird with a green crown, yellow wings and a red belly, and that's because the entire sentence serves as a single input to the

generator, and the detailed information of the description is lost. As a result, they found that the original images the model created were often a blurry greenish-yellowish-reddish bird instead a close, sharp match with the original description, and this is where they, again, took inspiration from humans.

As humans draw, we repeatedly refer to the text and pay close attention to the words that describe the region of the image we are drawing, so to capture this human trait, the researchers created what they call an Attentional GAN (AttnGAN) that mathematically represents the human concept of attention. It does this by breaking up the input text into individual words and matching those words to specific regions of the image, and once they wrapped this development into their models they found that the synthetic imagery improved dramatically.

3. Steerable GANs

An open secret in traditional machine learning circles is that many AI models work beautifully when used to generate imagery based on standard benchmarks, but that they lack the flexibility and usability that users expect once they leave the lab and start being used in production to create tailored content.

Notes:

This is something that researchers have attributed to the fact that the models used to create certain imagery are often trained on biased, narrow datasets, for example, of cats and dogs that are all centred in the middle of images, that lock their models into generating imagery with very specific styles and viewpoints that users then find difficult to modify to their individual briefs and requirements when they decide to use them to generate content.

More modern generative models too are no exception, and while there have been significant advancements in using GANs to generate different types of dynamic imagery the fact remains that once generated those images are, again, difficult to modify. And this is where so called "Steerable GANs" now come into play that allow users to easily tailor and transform these generated images, such as by adjusting camera angles, colour, and lighting, in any way they like, and they work by letting users drag mouse icons around a 2D latent space, much in the same way we use colour pickers today, to change a wide variety of details in real time.

Researchers in this specific field go on to hypothesise that the degree of distributional shift that can then be applied to images post production is

related to the breadth of the training data distribution, and as we use more diverse data sets to train our models the adaptability and flexibility of this technique will improve.

4. Style GANs

Since Generative Adversarial Networks (GANs) were first developed in 2014 by Google Researcher Ian Goodfellow the technology has been widely adopted by researchers interested in using the technology to create synthetic imagery and domain transfer.

After some early failures GANs have made huge breakthroughs recently and can now produce highly convincing fake images of animals, landscapes, human faces, and much more. But while researchers know what GANs can do a lack of transparency in their "Black Box" inner workings means that most GAN results are still achieved mainly through trial and error, which, needless to say offers only limited control over the synthesised imagery.

In order to overcome this problem more modern techniques, from companies such as DeepMind and Nvidia, incorporate so called Style Transfer Techniques (STT) to create new generator architectures that

Notes:

are capable of learning and separating out the different aspects of an image unsupervised; and that enable intuitive, scale-specific control of the resulting synthesis.

And here's an example of how they work: Given an input facial image, the style-based generator can learn its distribution and apply its characteristics on an entirely new and novel synthesised image. While previous GANs couldn't control what specific image features they, or their users rather, wanted to synthesise, using a technique called Style Mixing (SM) these new generators can control the effect of a particular style, for example, high level facial attributes such as identity, pose, and shape, without changing any of the other generated features. Not only does this enable better control over specific features, such as eyes and hair styles, and stylistic content, but it now gives users the ability to more accurately create the images they desire.

Stochastic Variation (SV) is another key property that researchers have introduced that allows GANs to realise the randomisation of certain image attributes, such as the placement of facial hair, freckles, pores, and stubble density, in imagery of human faces.

As a next step researchers are now

planning on "quantifying interpolation quality and disentanglement" and propose a variety of new more automated techniques that include Perceptual Path Length (PPL) and Linear Separability (LS) that can be applied to future generator architectures, and as a result of all of these changes researchers have seen impressive results using these modified GANs to generate images of bedrooms, cars, and cats, as well as anime maestri and real world scenery.

SUMMARY

The ability to produce a diverse range of high quality imagery on demand will be very appealing to those enterprises and individuals that today have to pay large sums of money to access and publish copyrighted original works, however, the main impact of the field will be the sudden explosion in new content.

Notes:



EXPLORE MORE. Click or scan me to learn more about this emerging tech.



Notes:

WHILE WE could be forgiven for just thinking about synthetic content in terms of creative machines ability to generate and copy more “traditional” media types, such as imagery and video, it would be remiss of us to not include the new types of media that are emerging, and their both obvious and less obvious uses.

As everything within our world becomes increasingly driven and tied to the digital world the fact is that the overwhelming majority of the content we consume today only exists on in what I call surface media, for example, on the web, on screens and smartphones, but over time more of the content we consume will be more immersive, whether it’s consumed during sessions held in Virtual Reality environments, or, in time, via neural interfaces where content is streamed straight to and from our minds - something that, as science fiction as that might sound, has already been demonstrated with human volunteers multiple times.

However, when we start looking at alternative use cases for simulated and immersive content there are a wide number of interesting and promising growth areas. On the one hand researchers are now turning their attention to finding new ways

to easily and quickly generating immersive procedural content for gaming applications and VR environments, but increasingly one of the hottest alternative uses cases in town right now is the ability for organisations to virtually test and train new AI models and products within simulated environments, as well as innovate and iterate them. And while there are now many examples of where these so called simulation engines have been used some of the most intriguing involve using them to cram hundreds of man years of learning into minutes in order to rapidly train neural network models and robots, and using them, in combination with other technologies, tools and techniques, to accelerate product innovation - from hardware through to software - by multiples of tens of thousands.



FUTURE HISTORY ...

A VISUAL HISTORY OF THE FUTURE





2017 saw researchers develop new Video to Video Domain Transfer models capable of digitising and altering streaming video in real time, that dramatically cuts down the time to “digitise” real video.

2018 saw a variety of developments across multiple AI disciplines and the combination of AI’s conversant in alternative domains. These combinations were then used to create realistic simulation environments to train everything from AI’s, self-driving vehicles, and robots, millions of times faster than previously possible.




2019 saw the development of the first AI capable of synthesising realistic game environments using voice commands, and the continued development of more advanced Video to X simulation models. Researchers push harder to auto-generate photo-realistic VR environments using these techniques.

In 2020 researchers will focus on developing faster and leaner models that deliver higher quality output, and have a better understanding of the laws of nature that govern our world. Researchers will also develop new techniques and tools to control and modify the output, and begin integrating it with other systems.



By 2025 it will be increasingly difficult to distinguish the simulated worlds produced by these models from the real thing, and the output will be much more controllable. However, computer processing and network speeds could still be an adoption bottleneck as the technology starts to commercialise.

By 2030 the simulated worlds that these systems generate will be fast and easy to create, and fully modifiable, albeit that they will likely still need a level of expertise to tweak.



Meanwhile attention will have shifted to creating procedural content generators, and integration with other technologies will accelerate.

By 2035 the technology is fully matured, commercialised, and ubiquitous across multiple platforms.



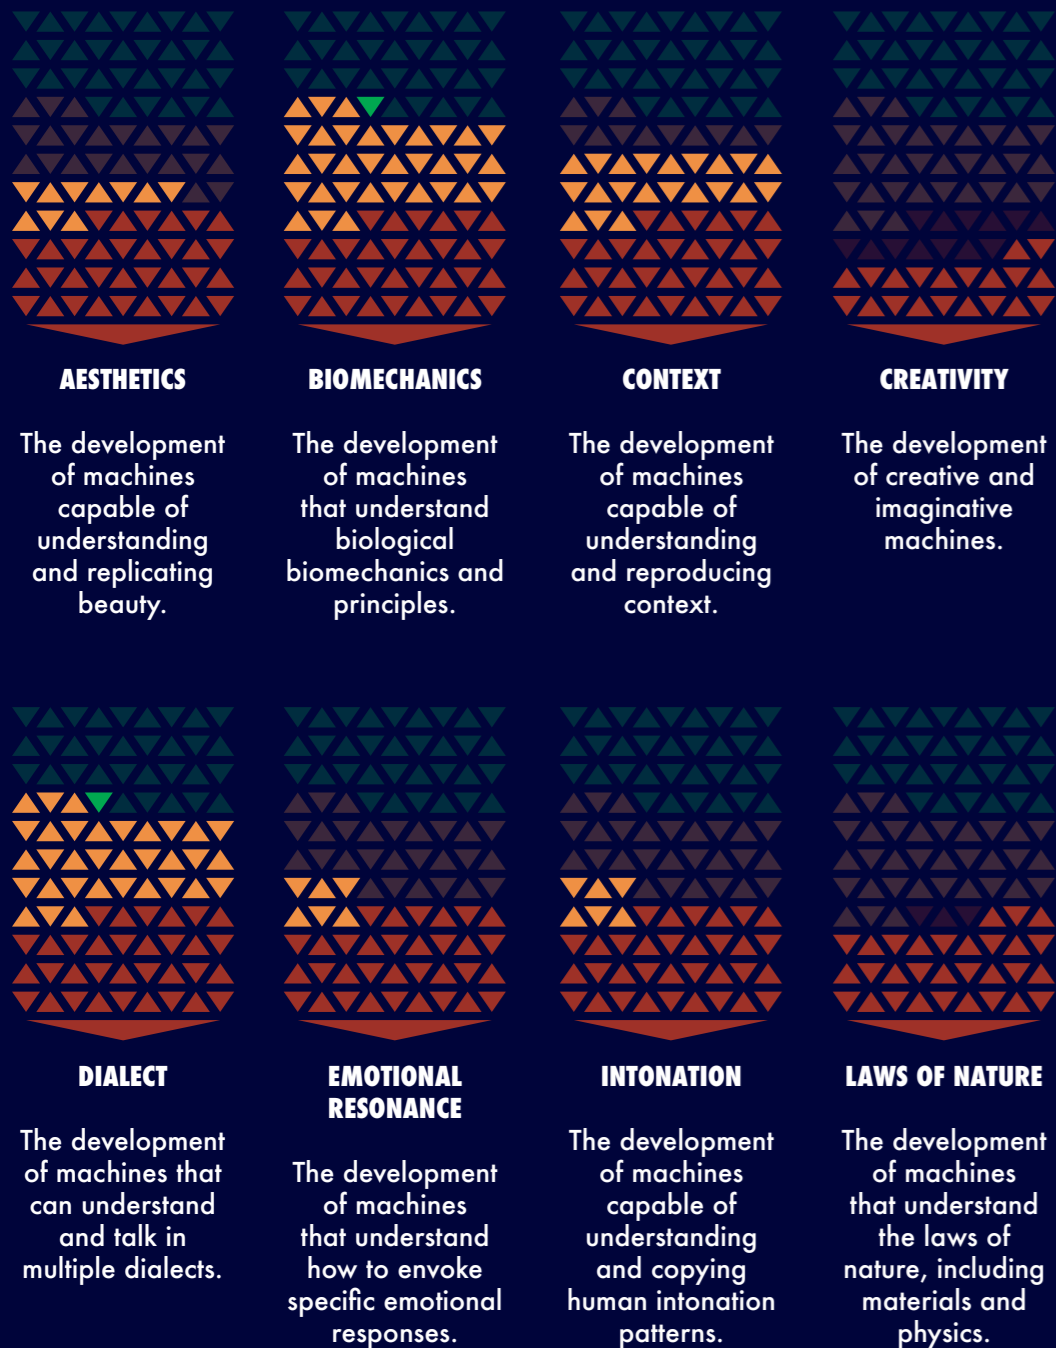


Figure 6. Current state of synthetic simulations development by AI discipline.

Source: Multiple Sources

THE WORK STILL TO DO

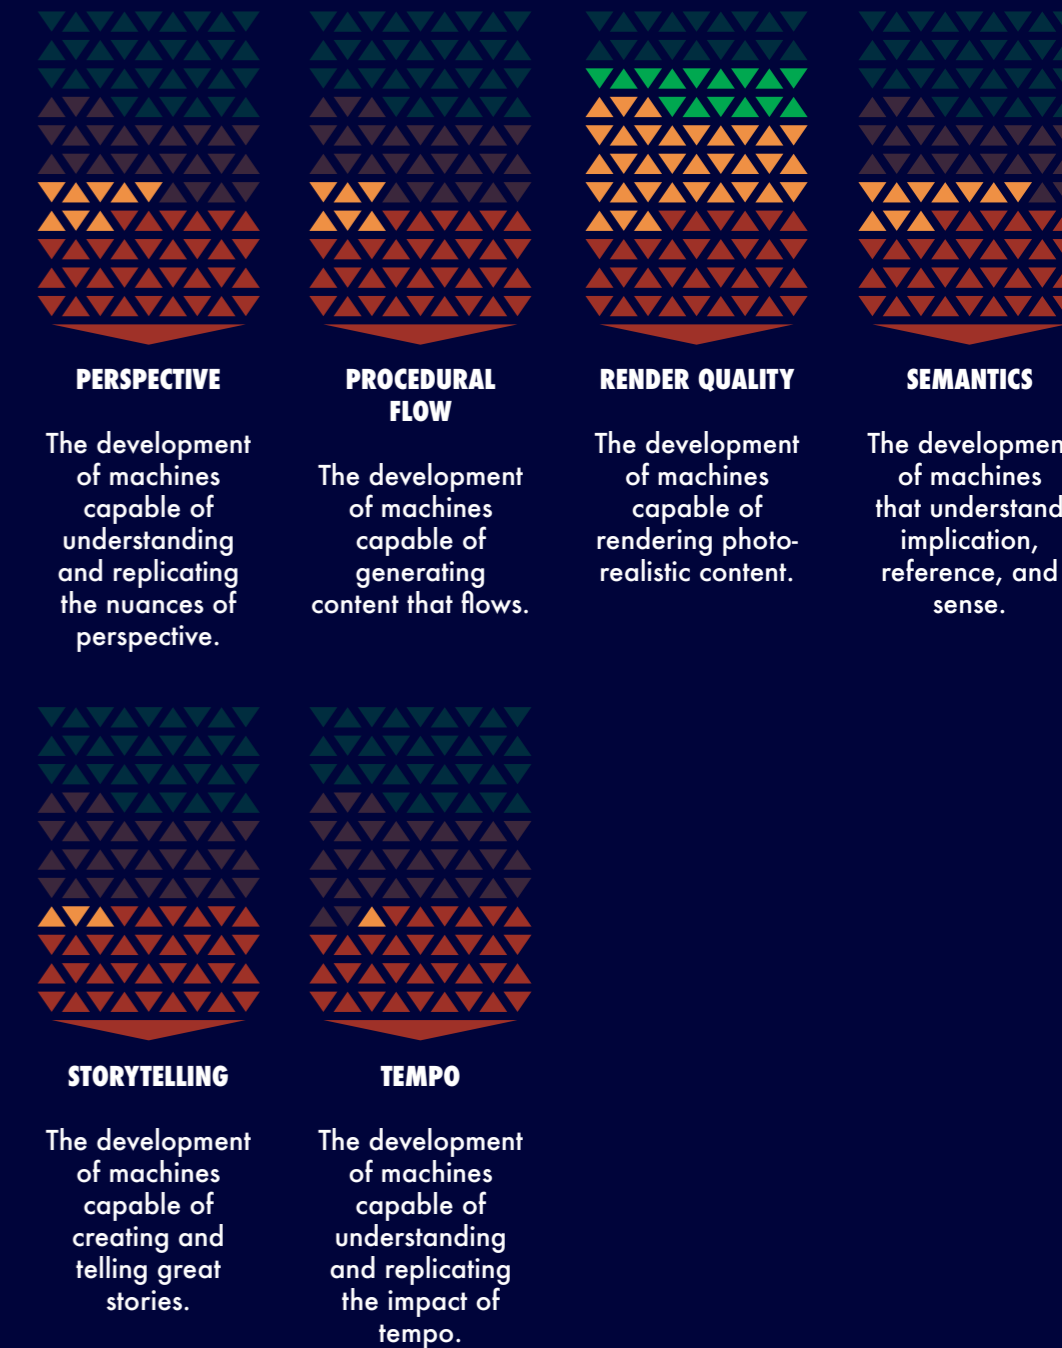
Synthetic simulations are, in my mind at least, one of the most fascinating areas of creative machine development today, if not only for the reason that on the one hand these are the platforms that will help us create truly expansive and immersive virtual reality environments, as well as realistic simulation environments that can be used to accelerate the development and training of future generations of AI's and robots, accelerate innovation and product development by thousands fold, and be used as the medium in which future movies, such as the recent Lion King, and other types of content are "filmed" and produced.

This is truly a revolutionary field - and for all the right reasons, but because of that it is also highly reliant on not one but many different AI disciplines coming together before researchers are finally able to create a universal and generalised synthetic simulation generator.

Furthermore the level of difficulty is increased several fold because, in short, the end goal for many researchers in this field is to create truly expansive and immersive digital environments and worlds that mirror the real world - real

life, but in digital form. And better.

As a consequence, for this field to be considered fully mature there are a huge number of different disciplinary areas that need to mature and then be combined, and this is a formidable challenge. In time though, as each individual challenge is overcome the creative machines in this space will become some of the most game changing platforms on the planet.



Notes:

TWO WORLDS COLLIDING

Having looked into all the different synthetic content fields in depth it's this field, the Synthetic Simulations field, that's perhaps the most interesting and understated especially when we consider that eventually the technologies, techniques, and tools that we use to create these realistic simulated environments will be the same technologies, techniques, and tools that we use to create more of tomorrow's synthetic content, most notably dynamic content such as videos and movies.

As our ability to create, and then importantly control and modify, these synthetic simulations improves it's easy to see how with some tweaking the technologies, techniques, and tools used to create today's AAA computer games could be used to create tomorrow's ultra-realistic movies, and so on.

A great example of this is Disney's Lion King 2019 re-creation where, in a world first, the company behind the film built and filmed the entire movie, everything from the characters to the environment, in Virtual Reality using the Unreal engine - an engine that up until that point had mainly been used to develop games.

As a result of this it is inevitable that

enterprises in this space, for example the gaming companies, will eventually break out of their existing industry niches and challenge not just the big studios but also every other content creator out there.

In short the digital world of gaming and simulations will collide with the world of mainstream mass market digital creators and studios.

Furthermore, when we consider the implications of such a coming together on the user experience this means that user experiences can be more intense because more information, such as haptic information, can be encoded into these environments easier than they can using other more "traditional" synthetic video technologies, techniques, and tools.

Additionally, not only will this give creators more fine grained control over each individual experience as users flip between devices but it will also allow enterprises to adjust their business models in new ways and offer tame, moderate, and intense experience packages. And much more.

TREND IMPACT AND TRAJECTORY

This trend will have a wide range of impacts including on the following market

Notes:

categories:

- Gaming
- Simulation engines
- Video production
- VR environments

The impact of synthetic simulations will be nothing short of revolutionary and will impact every aspect of global culture, industry and society, and they will, arguably, be one of the biggest drivers in accelerating the global rate of change.

One of the most exciting opportunities enabled by this field, for example, is to accelerate the rate of product development and innovation by multiples of thousands by, for example, using simulations to accelerate the rate of Artificial Intelligence (AI) model development, and by training drones, robots and other products millions of times faster than we can using today's traditional techniques.

Asides from this, however, as the quality of synthetic simulations improves we could also see a day when movies are constructed and filmed in immersive Virtual Reality (VR) worlds, and that is before we discuss the convergence of the digital and physical worlds and the opportunities that that brings.

Meanwhile from a trajectory perspective as with all of the synthetic content generators we're discussing researchers will continue to develop new machine models, iron out the bugs and the kinks, and eventually those machines will leave the labs, be adopted by enterprises and integrated with other products and services, and be adopted by consumers - while the regulators play catch up.

TECHNOLOGIES BEHIND THE TREND

While this field is relatively new it's one that is accelerating rapidly as the amount of computing power enterprises and researchers have access to scales almost exponentially, so, as you'd expect as different teams of researchers experiment with new techniques it's also a field that is accelerating rapidly.

The primary areas of research in this field today include:

- Text to Video generation
- Video to Synthetic Video generation
- Video to Video Domain Transfer
- Video to VR generation

So with that in mind now let's take a look at some of those technologies, tools, and

Notes:

techniques.

1. Unreal Engines

One of the most promising techniques at the moment involves combining streaming video input with a Video to Video synthesis AI architecture, and the Unreal Engine to create so called realistic "Semantic maps" of environments and scenes which is aided by the fact that the AI assigns every pixel on the screen with an identifier label, which then helps with post processing, and, if required gives consumers a simple way to edit and change different aspects and attributes of the environment. So, for example, in an urban environment some pixels will be labelled as "Car" and others as "Building," and so on. These clumps of pixels are then also given clearly defined edges which helps the Unreal Engine to produce a type of sketch of the environment that then gets fed into a neural network that renders the scene.

Despite the fact the field is advancing quickly though it's still likely that it will be a couple of years before these digitally rendered worlds will be ready to be incorporated into Virtual Reality worlds, and there are still many details and kinks to be ironed out before the renders become photo-realistic

SUMMARY

Perhaps of all the technology fields to watch and keep an eye on this is the one, however, your view of this particular field and its relevance to you will evidently be coloured by the sector you reside in.

Notes:



EXPLORE MORE. Click or scan me to learn more about this emerging tech.



Notes:

ENTERPRISES AROUND the world have been trying to use Artificial Intelligence (AI) in all its various forms to analyse text for arguably a couple of decades now, whether it's for national security purposes or academic ones. But the phenomenon of using AI to generate text, for a variety of reasons, is a relatively new phenomenon that has its roots in the world of intelligence gathering where national security agencies such as the NSA, faced by the huge exascale volumes of intelligence data, wanted to find new ways to condense and summarise it all into easily digestible narratives that were easier for their human analysts to digest.

The global national security community, however, aren't the only ones who face a data deluge, and over time many of the concepts developed and honed in their field were translated across to other sectors and, where necessary tailored accordingly. As a result the first commercial platforms to emerge were news generators, so called Robo-Journalists capable of ingesting large volumes of data, making sense of it all, and then condensing it down into easy to digest news articles.

FUTURE HISTORY ...

A VISUAL HISTORY OF THE FUTURE




2017 saw the continued development of synthetic text generators capable of summarising volumes of textual information, contextual and semantic reasoning platforms, and the increased use of these platforms to create professional, commercial Robo-journalists, among other things.

2018 saw a variety of developments in the field, but in the main it was pretty much business as usual as companies refined the accuracy of their models, and continued to develop more robust systems.





2019 saw the development of the world's most dangerous synthetic text generator, the world's first AI generated book and copywriters, and the first attempts to use synthetic text generators to create scripts for production adverts, movies and TV series.

In 2020 the existing text generator models will be trained using an increasingly large corpus of data, and researchers will start to develop new techniques and tools to control, modify and increase the accuracy of the output.



By 2025 the technology will be almost fully mature and controllable. The only fields left to conquer will be specialist fields and multi-lingual challenges, and researchers will increasingly turn their attention to integrating these systems with others.



By 2035 the technology is fully matured, commercialised, and ubiquitous across multiple languages and platforms.

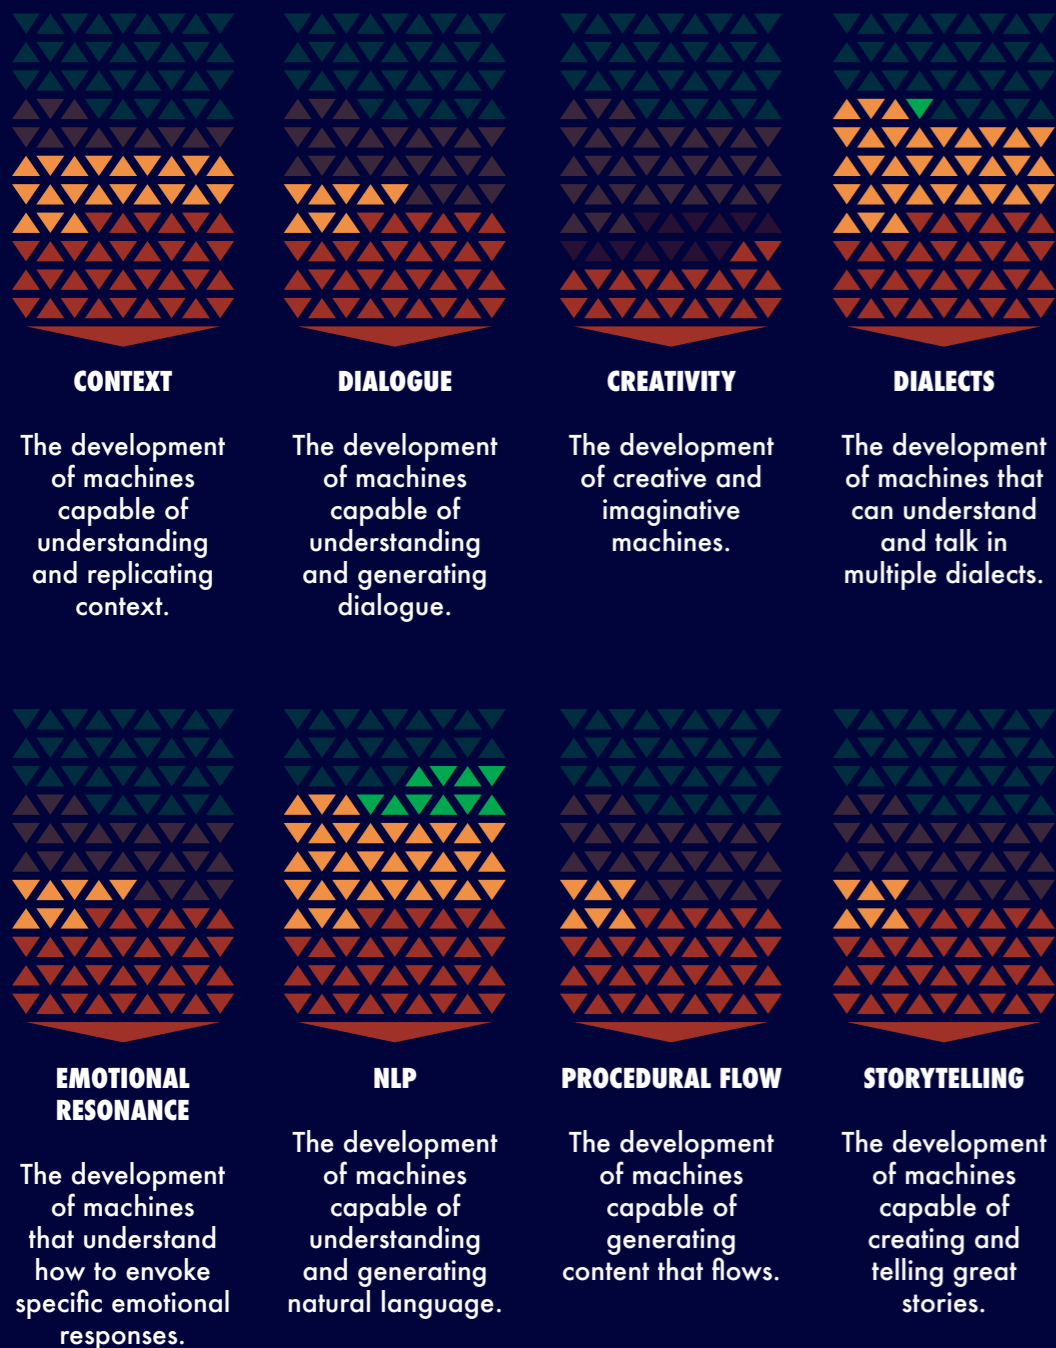


Figure 7. Current state of synthetic text development by AI discipline.

Source: Multiple Sources

THE WORK STILL TO DO

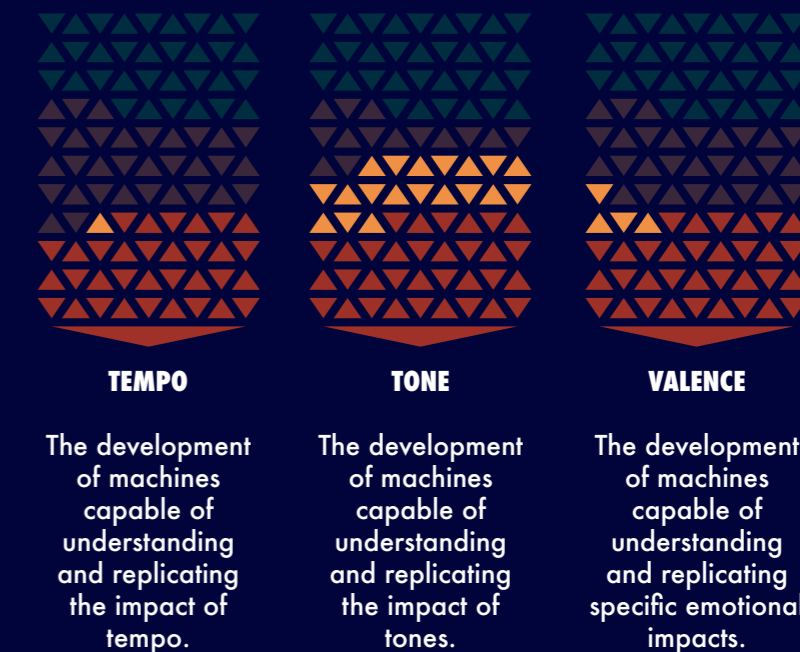
Unlike many of their other synthetic content cousins, there are already a number of good synthetic text generators in the wild and in commercial use in fields as diverse as copy writing and journalism, and the rate of development in this field is accelerating fast, so much so that we have already seen these platforms write the first robo-novels and even rudimentary movie scripts.

That said though, like many of the technologies in this megatrend, it's not a hill they have to conquer it's a mountain, and as a result there are a number of different AI discipline areas that need to combine together before we have a truly universal and generalised synthetic text generator.

As a result of their early successes many researchers in the field are now turning their attention to developing creative machines capable of generating more long form content - a feat that's much more difficult to achieve because, again, like the other synthetic content fields, these machines have to draw on a variety of other disciplines that include everything from learning that stories have a beginning, a middle, and an end, all the way through to being able to understand context, and generate text

capable of soliciting the right emotional responses.

As these machines command of the natural language improve though in time researchers in the field will be able to devote more of their time training their models to generate the right rhythm and flow for pieces, find new ways to help them unleash their creativity and imagination, and create stories to be proud of.



Notes:

TREND IMPACT AND TRAJECTORY

This trend will have a wide range of impacts including on the following market categories:

- Articles
- Copy writing
- Narratives
- Novels
- Reports
- Research papers
- Scripts
- Summaries

THE TECHNOLOGY BEHIND THE TREND

Recently a number of enterprises, such as OpenAI, a non-profit AI research firm in the USA, have released a range of new AI language models that are capable of generating convincing passages of prose. So convincing, in fact, that not only have some of these caught the media's attention and imagination, but in some cases the researchers have refrained from open sourcing the models in the hope of stalling their eventual and, unfortunately, inevitable weaponisation – especially when it comes to using them to mass produce fake news. Putting this into context it's not for no reason that the team of researchers behind one of the most advanced models called their

creation "potentially most dangerous AI out there," because it was that good at creating fake content.

While these impressive results are a remarkable leap beyond what traditional language models have been able to achieve up until now the technique involved in creating them isn't all that new. One of the latest breakthroughs, for example, was driven by feeding the model an ever increasing amount of training data - a technique that's also behind many other recent advancements in the field which now include many of the models being used by everyone from the Wall Street Journal, to the researchers who are using them to create the first ever AI generated books – the first of which, a research book that summaries all of the latest Lithium Ion battery research, has already been released.

"The latest [OpenAI] model is surprising people in terms of what they can create with more data and bigger models," says Percy Liang, a leading computer science professor at Stanford University.

Once trained the passages of text many current models produce are good enough to masquerade as something written by humans, however, this ability to write should not be confused with a genuine

Notes:

understanding of language, which is the ultimate goal of the subfield of AI known as Natural Language Processing, or NLP for short.

In fact, getting machines to the point where they have a genuine understanding of natural language is a challenge that has so far largely eluded NLP researchers, and it's a goal that many experts think could take years, even decades, to achieve and involve techniques that haven't been developed yet.

That said though there have been some noteworthy developments in the field with the development of new models that are being trained to generate new scripts for a wide array of popular series including Game of Thrones.

At the moment four different philosophies of language that currently drive the development of NLP techniques, so let's look into them all.

The primary areas of research in this field today include:

- Copy writing
- Long form text generation
- Short form text generation
- Script generation
- Summarising

1. Distributional Semantics

Linguistic philosophy. Words derive meaning from how they are used. For example, the words "cat" and "dog" are related in meaning because they are used more or less the same way. You can feed and pet a cat, and you feed and pet a dog. You can't, however, feed and pet an orange.

How this translates to NLP. Algorithms based on distributional semantics have been largely responsible for many of the recent breakthroughs in NLP, and, like many fields, it's a field that's accelerating. Researchers use machine learning to process text, finding patterns by essentially counting how often and how closely words are used in relation to one another. The resultant models can then use those patterns to construct complete sentences or paragraphs, and power things like auto-complete or other predictive text systems.

In recent years, some researchers have also begun experimenting with looking at the distributions of random character sequences rather than words, so models can more flexibly handle acronyms, punctuation, slang, and other things that don't appear in the dictionary, as well as languages that don't have clear delineations between words.

Notes:

Pros and Cons. These algorithms are flexible and scalable because they can be applied within any context and learn from unlabelled data, but on the negative side the models they produce don't actually understand the sentences they construct - at the end of the day, they're writing prose using word mere associations.

2. Frame Semantics

Linguistic philosophy. Language is used to describe actions and events, so sentences can be subdivided into subjects, verbs, and modifiers, such as who, what, where, and when.

How this translates to NLP. Algorithms based on frame semantics use a set of rules or lots of labelled training data to learn to deconstruct sentences. This makes them particularly good at parsing simple commands, and thus very useful for creating chatbots or voice assistants. If, for example, you asked Alexa to "Find a restaurant with four stars for tomorrow," such an algorithm would have to figure out how to execute the sentence by breaking it down into the action "find", the what "restaurant with four stars," and the when "tomorrow."

Pros and Cons. Unlike distributional

semantic algorithms though, which don't understand the text they learn from, frame semantic algorithms are able to distinguish between the different pieces of information in a sentence, and these can be used to answer questions like "When is this event taking place?" On the negative side however currently these algorithms can only handle very simple sentences and therefore fail to capture nuance because they require a lot of context specific training. Also as a consequence they're relatively brittle and not particularly flexible.

3. Model Theoretical Semantics

Linguistic philosophy. Language is used to communicate human knowledge.

How that translates to NLP. Model theoretical semantics is based on an old idea in AI that all of human knowledge can be encoded, or modelled, in a series of logical rules. So, if you know that birds can fly, and hawks are birds, then you can deduce that hawks can fly. This approach is no longer in vogue, however, because researchers soon realised there were too many exceptions to each rule, for example, penguins are birds but they can't fly, which would result in an error.

Notes:

But algorithms based on model theoretical semantics are still useful for extracting information from models of knowledge where the information is always consistent and logical, such as databases. Like frame semantics algorithms, they parse sentences by deconstructing them into parts, but whereas frame semantics defines those parts as the who, what, where, and when, model theoretical semantics defines them as the logical rules encoding knowledge.

For example, consider the question "What is the largest city in China by population?" A model theoretical algorithm would break it down into a series of self-contained queries such as "What are all the cities in the world?" "Which ones are in China?" "What are the cities' populations?" "Which population is the largest?" It would then be able to traverse the model of knowledge to get you your final answer.

Pros and Cons. These algorithms give machines the ability to answer complex and nuanced questions, but on the negative side they require a model of knowledge, which is time consuming to build, and are not flexible across different contexts.

4. Grounded Semantics

Linguistic philosophy. Language derives meaning from lived experience. In other words, humans created language to achieve their goals, so language must be understood within the context of our goal oriented world.

How this translates to NLP. This is the most modern approach and the one that many experts in the field think holds the most promise. It tries to mimic how humans pick up language over the course of their lifetimes, so the machines start with a blank state and learn to associate words with the correct meanings through conversation and interaction.

In a simple example, if you wanted to teach a computer how to move objects around in a virtual world, you would give it a command like "Move the red block to the left" and then show it what you meant, and then over time the machine would learn to understand and execute the commands without the need for human interaction.

Pros and Cons. In theory, these algorithms should be very flexible and get the closest to a genuine understanding of language, but on the negative side teaching is very time intensive - and not all words and phrases

Notes:

are as easy to illustrate as “Move the red block.”

In the short term, most experts think that the field of NLP will see much more progress from exploiting existing techniques, particularly those based on distributional semantics, but in the longer term they all believe they have limits.

“There’s probably a qualitative gap between the way that humans understand language and perceive the world and our current models,” says Liang.

Closing that gap would probably require a new way of thinking, he adds, as well as much more time.

SUMMARY

Like every creative machine field it’s fair to say that generating synthetic text, for whatever purpose, is a field that on the one hand is coming under increasing scrutiny because of the impact that the ability to generate synthetic text will have on society and especially trust, and on the other hand over the past year research in the field has progressed rapidly with an increasingly large number of AI models being produced that many experts have deemed too “dangerous” to be released into the wild.

However, as the number of models produced, and the volume of research and papers grow, it is inevitable that the technology will escape into the wild and be democratised for better and worse.

Notes:

THE FUTURE OF GAME CONSUMPTION



TODAY GAMES are played and consumed in a variety of different formats and ways. There are, of course, games that are still played physically such as the ancient games of Chess and Majong, and there are, of course, games that are played digitally, such as Forza and StarCraft, and obviously the biggest impact of the digitisation of games has been to free them up from being played in just one way on one thing, so as technology continues to develop and advance, and as the number of devices and gadgets we have continue to explode, so too do the variety of ways we can play games.

experience when you compare it to, say, playing it in virtual reality with a haptic suit.

DEVICES ARE AND ARE NOT IRRELEVANT

Generally the theory goes that one of the greatest impacts that cloud game streaming will have, from a user perspective at least, will be to make the end device irrelevant because you'll be able to play your games on any device anywhere.

However, while that viewpoint is correct from a technological perspective it's wrong from a user perspective because different devices can deliver very different user experiences. After all, playing Halo, for example, on a smartphone will still be a very different



THE FUTURE GAME EXPERIENCE

YOU MIGHT have heard of the phrase “In the game.” As we charge headlong into the future this could very well change from being a figurative phrase to a literal one as technology increasingly makes it possible for us to not just create realistic simulations, but also live in them - much in the same way that we live and interact with the real world today. So, let’s plug into the Matrix and dive in.

Today there are three types of games experiences, namely traditional, where games are played on screens, mixed reality that merges the digital and real worlds, and then virtual reality where gamers are immersed in completely virtual worlds. And the same will be true in the future.

Today though the experiences that these different formats generate primarily tie into two of our five senses, namely sight and sound. But as we race into the future they will increasingly be able to tie into and influence all five senses, including smell, taste, and touch, as well as a sixth sense which is arguably our “Soul,” and I’ll discuss how we do that later in this chapter.

In the future though game experiences will also become much more intense, more personal, and more real. Note

that I said real, and not realistic. This distinction is important because realistic simply suggests that it’s difficult to distinguish a simulated environment from a real one. But if a gamer can be fooled into thinking that a simulated environment is real though then their entire concept of reality can break down and ultimately their own reality changes - and that takes us down the equivalent of Alice in wonderland’s rabbit hole.

All this said though it will still be down to gamers to decide for themselves just how real and immersive they want their future game experiences to be, and new technologies will simply be the door through which they access them.

THE FLAT SCREEN

The vast majority of games today are played very much in the same way that we played computer games back in the 1980’s - via a screen.

Sure, the content, quality, and the resolution of the games have changed dramatically during that period as we’ve seen dramatic advances in computer power, and yes, today we have a choice of more screens than we’ve ever had before, but fundamentally it’s the same. And screens are going nowhere any time

Notes:

soon - even as we see the emergence of new forms of retinal display systems that beam images directly onto the retina, and continuous improvements in alternative screen types such as AR glasses, E-Inks, metalenses, and even pico projectors.

PERSONAL SOUND

The arrival of new metamaterials are now opening the door to an entirely new world of sound, one where sound is personal, and one where only you hear the sounds that the game developers want you to hear while your friends sitting next to you hear a completely different set of sounds - and all without the need to isolate yourself from each other using old fashioned headphones.

In short, as the technology develops we can now see a time when headphones are old tech, and when we finally are able to ditch them gamers will be able to loose yet another piece of technology that will only help make the game experience even more realistic than before.

MIXED REALITY

As our gadgets and devices, which are at

the edge of the network, continue to get more powerful so too will the augmented and mixed reality experiences they can be used to create, and game studios will increasingly take advantage of this to create not just new games, but also new sports.

VIRTUAL REALITY

Today we think of virtual reality environments as fully simulated digital environments that we access using clunky headsets. We dive in, and then we dive out when we're done, and have to massage our necks. But, as headsets improve in time and as the resolution and fidelity of the worlds gamers can access using them improves, first from 4K 60fps to eventually 12K 120fps, then in time today's relatively basic simulations will become increasingly photo-realistic and life-like.

New technologies, such as metalenses, will also help us get rid of these clunky headsets and replace them with, in the first instance, nano-thin glasses, and then eventually smart contact lenses which are perhaps the ultimate gamers accessory.

BEYOND VIRTUAL REALITY

Notes:

There is another VR technology emerging though, and it's an odd one frankly. While it uses the same headsets we've all come to know and hate, and in time it could also use all of the other technologies we use to deliver VR experiences this new form of experience, which doesn't have a formal name yet takes a leaf out of AR and mixed reality's play books except for the fact that instead of overlaying a few digital elements over your real world it replaces your entire real world with simulated elements instead.

For example, on your way to work, if you're wearing a headset, it will change your entire real world environment into a digital simulation and then replace every real world element with simulations of your choosing - turn the people into aliens, the sidewalk into a beach, and the skyscrapers into cucumbers. The sky is literally the limit.

NEURO JACKING

Arguably though, despite all these other advances and options, the ultimate gamer experience has got to be one that allows gamers to jack their brains directly into the game, and in doing so bypassing all of the other fancy technology and gadgets that are simply there to clutter

up their shelves.

As we continue to see significant advances in neuroscience, and the development of both invasive and non-invasive Brain Machine Interfaces (BMI), that are now at a point that let us upload knowledge to people's minds, and stream their thoughts to television, this reality is closer than some people might think. Albeit that it's still decades away from being commercialised if only for the reason that it'll freak the hell out of regulators.

FULL IMMERSION

Being able to jack your brain into a game might be all very well and good though, but as advanced a technology as that is, and as advanced an experience that it'd deliver, it's still highly likely that there would be gaps in the user experience. For example, your sense of touch will likely not be stimulated enough or in the right ways, plus those feelings would be simulated and not really real, and as for being able to physically navigate and explore game environments, well, while you can trick your mind into thinking your moving and navigate worlds simply using your mind, the fact is that you'd likely still be sat on your sofa like a couch potato gathering

Notes:

dust.

So, for those people who are looking for the ultimate full immersion experience we have plenty of options.

When it comes to satisfying their sense of sight we can turn to VR and neuro jacking to help us create and immerse them in photo-realistic environments and simulations, and when it comes to sound we can use new personalised sound technologies to create ultra-rich soundscapes and soundtracks.

With regards to the sense of smell and taste, well, again, there are an increasing number of what are being referred to as Digital Sense technologies emerging in the world's labs that, while there might be questions about their design and execution, are already being integrated with VR headsets to fool gamers into thinking they're smelling and tasting real environments and real foods.

And as for touch, well, here we have haptic technology to thank for bringing the sense of touch to life, and making everything from a simulated hug or punch feel real with real world consequences that, in the case of the latter, if not careful can even cause bodily bruising and damage. As we continue to see advances in the field

of computing, electronics, fabrics, and sensors, increasingly haptic technologies are being embedded into everything, from gloves and vests through to entire immersive suits.

But, as technology helps gamers realise the ultimate immersive experience one challenge that we have yet to talk about is that of movement, real movement, and the ability to physically explore and navigate worlds in the way you'd move around in and explore the real world. Here we have the advent of electromagnetic floors to thank because they achieve just that - they let gamers get off of their couches and physically walk or run through simulated environments as they would in the real world.

Notes:

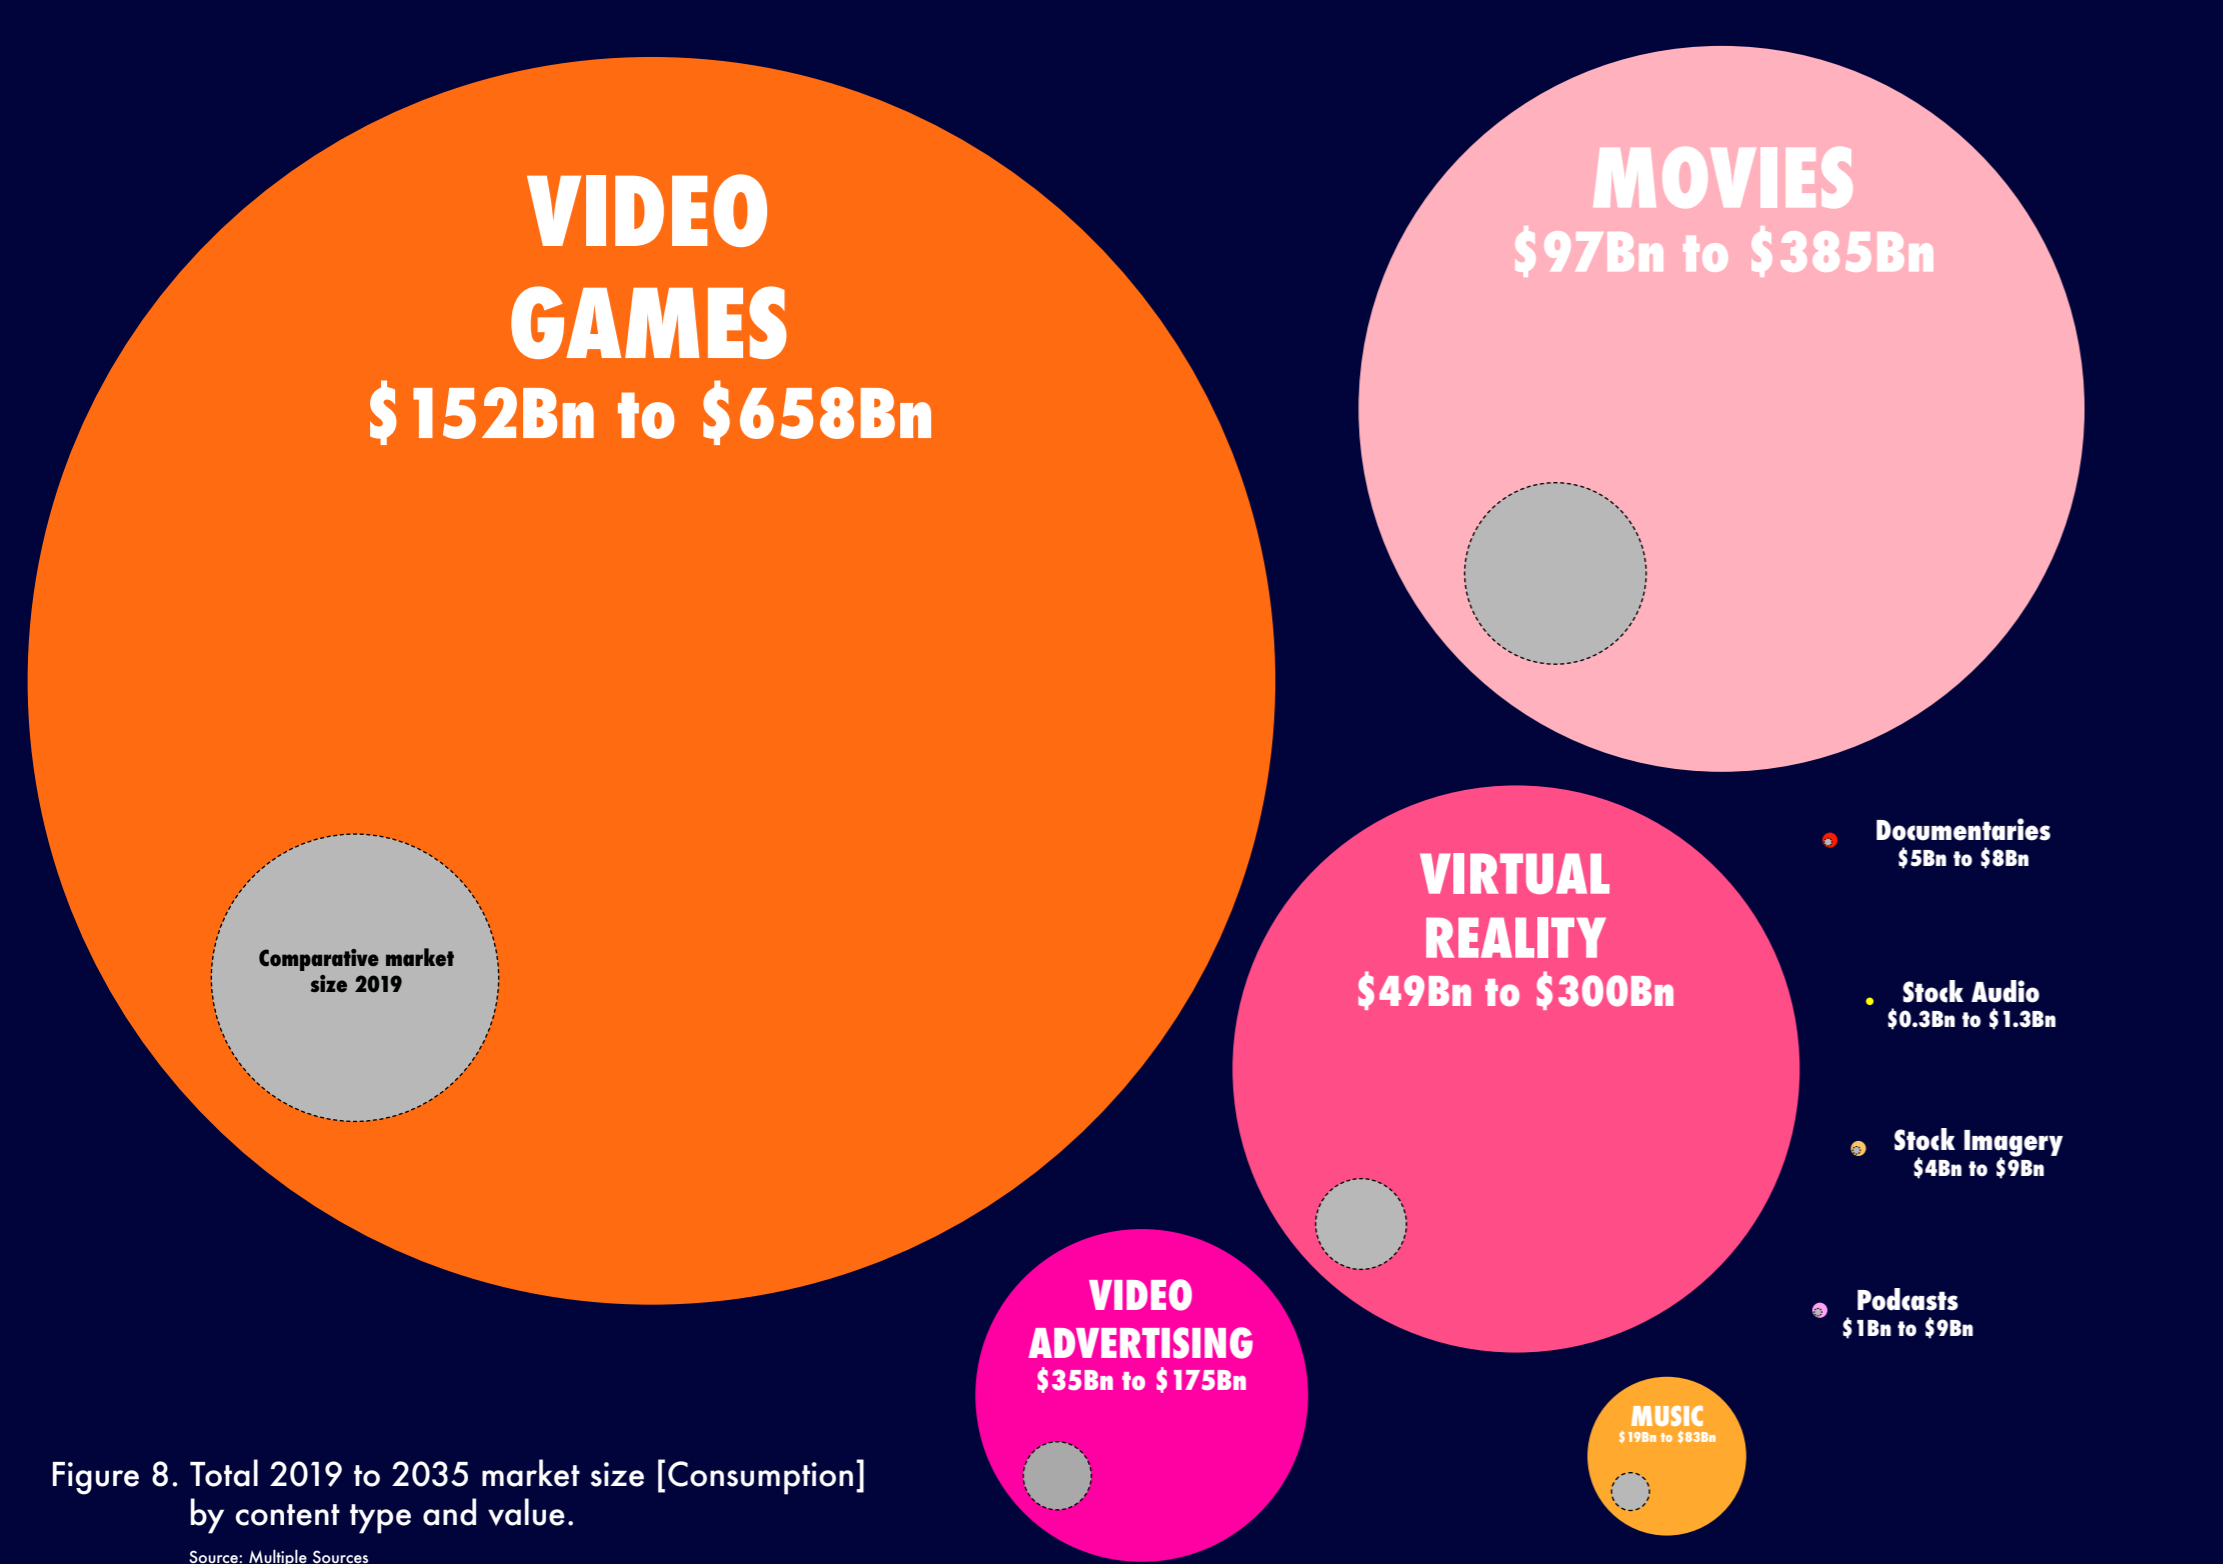
A hand is shown interacting with a futuristic digital interface. The interface displays various data visualizations, including a line graph with multiple colored lines (green, red, blue) and a grid. The background is dark with blue and purple lighting, suggesting a high-tech or gaming environment. The hand is positioned as if touching or pointing at the screen.

FUTURE MARKET SIZING AND PLAYERS

AS YOU would expect the global gaming market today is huge with an estimated 2 billion players, and as we continue to see improvements in technology that give more people more access to quality game experiences it's inevitable that that number is going to grow significantly over the coming decades.

Today the global gaming industry is estimated to be worth over \$152 Billion, and that's before we see some of the new distribution and business models bite - let alone some of the new technologies we've discussed such as AR, MR, and VR.

In this section I'll take a closer look at some of the new market entrants and players.



ECOSYSTEM: >> START

THE CLOUD GAMING ecosystem is still relatively nascent, even by the technology industry's standards, and is only a few years old at best. As a result the startups and enterprises that are trying to establish themselves in the market are relatively few with most enterprises keeping an eye on the space and waiting for the sparks to take hold before they make the leap and enter it.

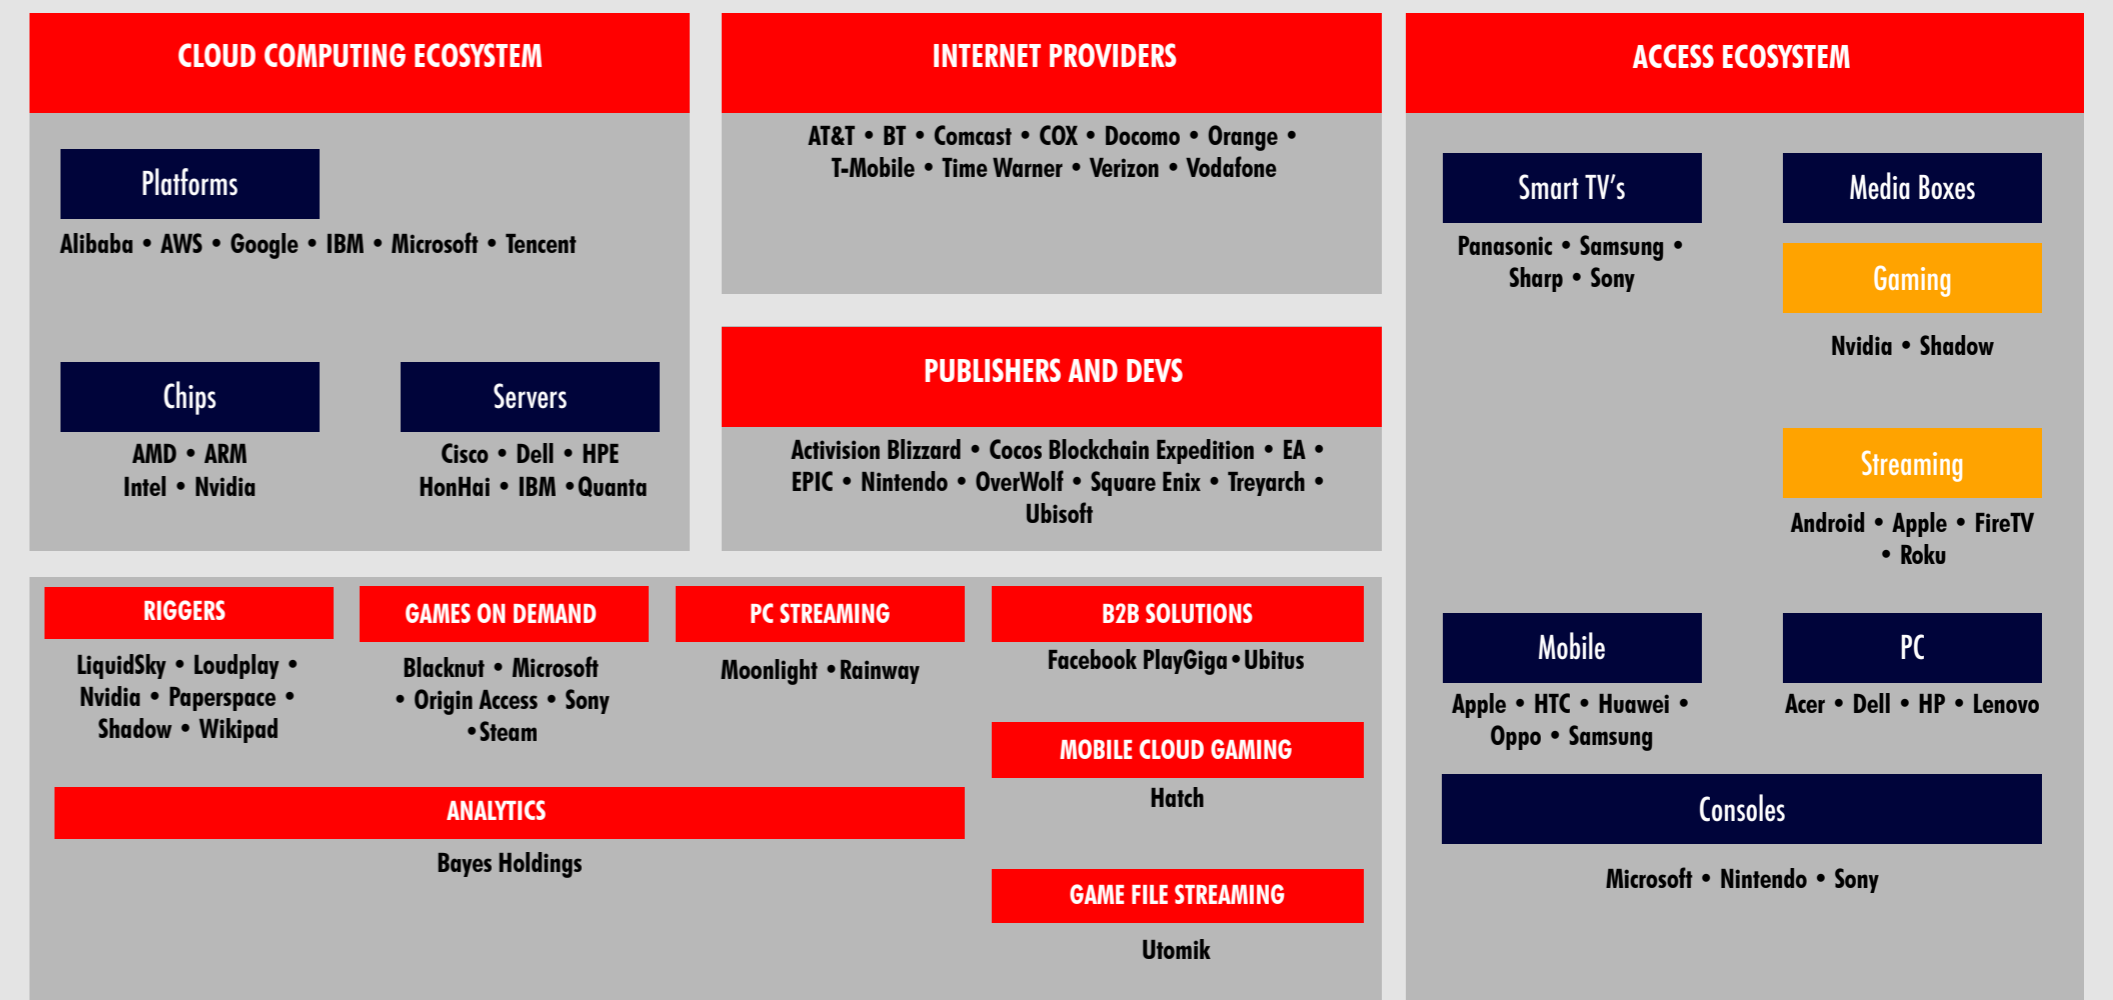


Figure 8. Total 2019 to 2035 market size [Consumption] by content type and value.

Source: Multiple Sources

Notes:

When we take a deeper look into the ecosystem and the way it will develop over the coming years it is highly likely that it will mature along similar lines to today's business environment where the dominant cloud hyperscale companies, such as AWS, Google, and Microsoft, will increasingly want to own the full stack and become vertically integrated cloud game streaming companies who not only run and operate the hardware at the bottom of the pyramid, but also own the developer studios, integration layers, and all the associated ecosystem of services on top of it all.

Given this environment though it is also likely that third party companies who are not owned by these hyperscalers, such as Activision, EA, EPIC, Nvidia, Sony, and friends will be forced to either attempt to build their own cloud game platforms, likely on top of the hyperscalers infrastructure, before building out their own proprietary development stacks and ecosystems in an attempt to defend and extend their existing market positions.

If this primary strategy doesn't work though and these companies start ceding market share to the hyperscalers then you can expect to see them form some unusual alliances not unlike the one we've already seen form between

rivals Microsoft and Sony as they try to consolidate their market position and try to out manoeuvre companies like Google.

In addition to this these companies and strategies though, at the other end of the spectrum, I'd expect companies like Samsung, who own a significant portion of the stack, especially when it comes to the end user interface by way of Smart TV's and smartphones, will also likely build out their own cloud game platforms as they try to disintermediate the other companies who will inevitably try to reduce them to a simple commodity hardware play.

So, with that out of the way now let's have a look at some of the more interesting third party companies, remembering that this information is an overview and any actual interest in the companies should be followed up with appropriate due diligence:

BLACKNUT

Blacknut is a subscription based cloud gaming company that operates an app first model that gives their users access to over 300 games, however while users can stream games to their set top boxes and PCs they won't be able to stream

Notes:

to smartphones unless they're on a 5G network.

COCOS BLOCKCHAIN EXPEDITION

Based out of Beijing, China, Cocos is a decentralised gaming and digital assets development platform that's built using their proprietary engine, and with over 1.1 million developers million developers using the platform they have one of the largest game developer communities in the world.

HATCH

Based out of Finland Hatch is a subscription based cloud gaming company with over 150 premium games , leagues, social scoreboards, and a child friendly business model. Unlike Blacknut they can stream games to mobile, and have significant partnerships with a number of telecoms providers and consumer electronics companies.

OVERWOLF

Overwolf is an API based game development platform that gives devs all the tools they need to build high performance high quality games, and

they have a strong ecosystem and strong stats that include over 7 million daily engagements, over 5 million monthly gamers and over 115 million game sessions a month.

SHADOW

Shadow is a subscription based cloud gaming platform that lets users install their own games on remote servers and stream a fully featured 4K 60Hz PC gaming experience to all of their devices.

CONCLUSION

PEOPLE SAY change is a constant, but in today's technology fuelled world this simple phrase can be a deceiving, and often comforting, misnomer because change isn't constant, it's exponential.

Today those changes are rippling through our global culture, industry and society faster now than they have at any point in human history, and the future belongs to the brave and the bold, not the meek and the mild.

Notes:

THIS IS NOT THE END.
EXPLORE MORE.



Copyright © Matthew Griffin, 311i Ltd. All Rights Reserved.

Produced in the United Kingdom.

This document is current as of the initial date of publication and may be changed at any time. Not all offerings are available in every country in which 311i operates.

The information in this document is provided "As Is" without any warranty, express or implied, including without any warranties of merchantability, fitness for a particular purpose and without any warranty or condition of non-infringement. 311i products are warranted according to the terms and conditions of the agreements under which they are provided.

This report is intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. 311i shall not be responsible for any loss whatsoever sustained by any organisation or person who relies on this publication.

The data used in this report may be derived from third-party sources and the 311i does not independently verify, validate or audit such data. The results from the use of such data are provided on an "as is" basis and the 311i makes no representations or warranties, express or implied.

UK311-20200109-DOC01