The New Disruptive Paradigm

GENERATIVE AI

TABLE OF CONTENTS

1. FOREWARD
The Era of AI disruptors

In recent years, economic disruption has been the norm worldwide. A global pandemic, changing worker needs, and the general realization that none of us can continue to do business as usual.

Artificial intelligence (AI) has a unique and powerful role to play in meeting many of the challenges presented to us today. We are now living in an age in which massive data availability and models offer the necessary foundations to fuel truly transformational applications and rapid innovation using AI.

While witnessing these rapid changes, we can also recognize the opportunities that existing challenges present. Yet business leaders have often failed to recognize the true impact that AI can have and the enormous potential it presents.

Today's race is no longer about adopting AI or automating processes for efficiency. It is now about realizing the massive economic disruption that AI can cause. For instance, to quote a simple example: using simple text in the form of a sentence, an image can be created, and from a couple of sentences, a customized movie can be produced. That is "a black swan event" that comes once in a decade and redefines the boundaries of the "information economy".

More importantly, as we discuss Generative AI technology in this report and its enormous potential, we will identify the major players in the field, the direction it is headed in and the area within the field likely to attract the most venture capital investments in the future. We will infer from the direction of venture investments, which is the economic segment that is likely to flourish in coming years.

REPORT HIGHLIGHTS

Gen AI disrupting the Film and Gaming Sub-sectors within the Media Industry

A revolution is occurring within the media industry, spurred by ChatGPT and large language models, with **Stable Diffusion** being a key player. A new emerging pattern divides Generative AI players into a *foundational layer* group and another *application layer* one. This delimits the playing field going forward, as was the case in the past with "operating systems" being developed as environments containing programs and applications.

Furthermore, in Section 2, **Open-source** versus **Closed-source** Gen-AI is discussed in detail and the advantages of the former over the latter are made evident, including the cost effectiveness of **Open-source**, resulting in granting enterprises the ability to start small and scale up when possible. Also, since Open-source code is freely available, it allows for public collaboration to address issues and provide support.

More importantly, this report estimates the **Total Addressable Market** (TAM) within the media industry for **Stable Diffusion** (total potential revenue that can be captured by Stable Diffusion due to the economic disruption caused by GenAI).
Spending in the media sector rests on three major pillars: **Customer Spending**, **advertising**, and **internet/online** access, with revenues split almost evenly currently and each pillar generating around 30% of revenues. The total size of the media market is around USD 2.5 trillion. It is split into multiple segments, with films and gaming being the dominant ones.

We have recorded the **impact of Stable Diffusion in the films** sub-segment through its recent ability to capture frame-by frame animation, and the **impact on gaming** through the new excellent "ideation" capability of Generative AI, which is likely to radically transform the production of games, expediting the process and massively reducing costs.

Furthermore, using a *top-down* approach to quantify the TAM, the *media and entertainment* sector has been divided into subsegments and further niches within the subsegments to recognize, the total potential revenue that can be captured by Stable Diffusion.

We found that while the TAM stands at USD 2.5 Trillion in 2023, the potential for economic disruption by GenAI amounts to nearly USD 1.07 trillion at present and is expected to grow to USD 1.4 trillion in 2025 *(see table on page 17)*.

Finally, through our review of the major players in Generative AI, we detail the Venture Capital investments in the field as a barometer of the future potential of the field. We identified fund flows per segment, category, funding rounds and geography (*see Section 9*).

As per our knowledge and to date, this is the most exhaustive look into the purely ***Generative AI*** ecosystem.

2. GENERATIVE AI
An Intro

The core of Generative AI (Gen-AI) is the Large Language and Image models (LLM), also known as foundation models. These models grant users the ability to do the following:

01 **Generate Content**
Those models can automatically generate content, such as *articles*, *blog posts*, *images*, *videos*, etc., acting as an invaluable time-saving tool for businesses who are pressured to create content under tight deadlines.

02 **Content Quality Improvement**
The Gen-AI-generated content can be of higher quality than that created by humans since AI models are able to synthesize large amounts of data and can identify patterns that are difficult for humans to recognize.

03 **Personalize Content to One's Specific Needs**
AI models can generate content based on personal preferences and individual choices. This can help businesses create client-tailored content for a target audience, hence being more useful.

Generative Tech is now being considered, worldwide, as the next big thing in software. It's a new level of human-machine partnership. It turns deep learning engines into collaborators to generate new content and ideas almost as a human would. Some have called it *"Foundational AI"*. AI models are the enabling base layers of the stack, with thousands of applications being built upon these.

The Generative Tech sector is developing at a very rapid pace, as reflected in real revenues and high valuations, even though a term for it had not been coined till as recently as a month ago. Eighteen months after its launch, Jasper is reported to have recorded nearly $100M in revenue and reached a $1.5B valuation. Open AI, which powers GPT-3 and other AI models, has been raising capital at valuations of tens of billions. Anthropic, another large model builder, has raised sizable amounts.

The recent availability of open-source alternatives to proprietary Gen AI models, proved to be a tipping point in the last six months. In short, EleutherAI, GPT-NeoX-20B, launched in Feb 2022, is the open-source alternative to *OpenAI*'s GPT-3 for text generation. *Stability AI*'s Stable Diffusion, launched in August 2022, is the open-source alternative to OpenAI's DALL-E 2 for images and videos. Both have been game changers on price, quality, and ease of access. The cost to generate images has dropped 100X ever since. The ability to generate output from these models through web and mobile has become "about 10 times easier" in the last six months.

This will be discussed in detail in the next few pages.

Marketing and Media will capture the lion's share of Generative AI's disruptive economic impact: According to Gartner[1] :

By 2025

30%

of outbound marketing messages from large organizations will be synthetically generated, a significant increase over the less than 2% today.

By 2030

a major blockbuster film will be released with

90%

of the film generated by AI (from text and video).

[1] https://www.gartner.com/en/articles/beyond-chatgpt-the-future-of-generative-ai-for-enterprises?source=BLD-200123&utm_medium=social&utm_source=bambu&utm_campaign=SM_GB_YOY_GTR_SOC_BU1_SM-BA-SWG

The Generative AI Potential

The sectors below are a small sample of sectors that Generative AI will disrupt in some shape or form. We will see, as we go, new applications tackle those specific sectors with highly innovative ideas based on foundational models.

3. Foundation Models

Foundation models (like *GPT-3* and *Stable Diffusion*) are extremely large models trained on broad datasets that can be adapted to a wide range of downstream tasks. Furthermore, a Generative AI model is specifically a foundation model, where the "training" involves modeling the "probability distribution" of the underlying data, for example, predicting the probability of a character being the next one in a given text sequence. We will discuss this in detail in the next few sections.

The low cost and ease-of-use of these models is helping to accelerate the development of AI apps as more engineers push themselves into the field of AI.

Today, foundation models are frequently adapted to build generative applications with the "wow" capability. But they can also be applied to more traditional ML use cases such as classification and entity extraction, and more importantly, they minimize (but not completely obviate) the need for startups to gather proprietary training data, label it, architect complex data transformations, tune hyperparameters, and select the right model.

Generative Applications

These are companies utilizing generative AI for its namesake purpose: the creation of net new output in various media types. This is by far the most prolific category and thus, comprises the majority of companies on our index. We are seeing startups here that are both building directly on top of existing foundation models, as well as those that have chosen the route of building their own models from scratch, particularly in domains where foundation models don't exist (e.g., speech).

The bottom layer is an AI model, which can generate novel output based on inputs that are unique to the user, such as OpenAI's DALL-E or Stability's Stable diffusion model.

The Large Language Models (LLMs) were first developed at Google in 2016, and were used as the backbone for their translation engine, trying to preserve meaning and context. Since then, large language and text-to-image models have proliferated internally, at major tech firms, including: **Google (BERT and LaMDA); Facebook (OPT175B, Blende-Bot) and OpenAI.**

Then in late 2021 and 2022, the following players emerged in the foundational space: **Stable Diffusion; MidJourney; and Crayion (Dall-E-Mini).**

These models have largely been confined to major tech companies because training them requires prohibitively massive amounts of data and computing power. GPT-3, for example, is trained on a 40 terabytes training set and employs 175 billion Neural Network coefficients to generate predictions. Hence, a training round for GPT-3 costs upward of $10 million, using thousands of NVIDIA GPUs - the norm in the AI computing industry.
Starting from scratch in the AI model layer is very hard. Most GenAI companies don't possess the data center capabilities or the large computing budgets required to design their own models despite the public availability of code. Nevertheless, many application layer companies are trying to establish a foothold in the foundation layer. These include: **Character AI (ex-Google employees); CohereAI (ex-Google employees); and Anthropic.** The three companies above are in the nascent stage, while the only two companies far ahead of others in the field are OpenAI and StabilityAI.

The funding pattern of most startups has been such that they have followed a revenue model adapted to the application layer, which is mostly B2C. There are some aspiring ones that have tried to move downwards in the model stack, but that is turning out to be challenging for them. High costs of training, associated with developing models using cloud computing resources, is a significant challenge, with one round of training requiring at least thousands of GPUs and costing millions of dollars. Additionally, algorithmically, Open AI and Stability have been very successful in developing the technology and have allocated funds and resources to this that are hard for new companies to match.

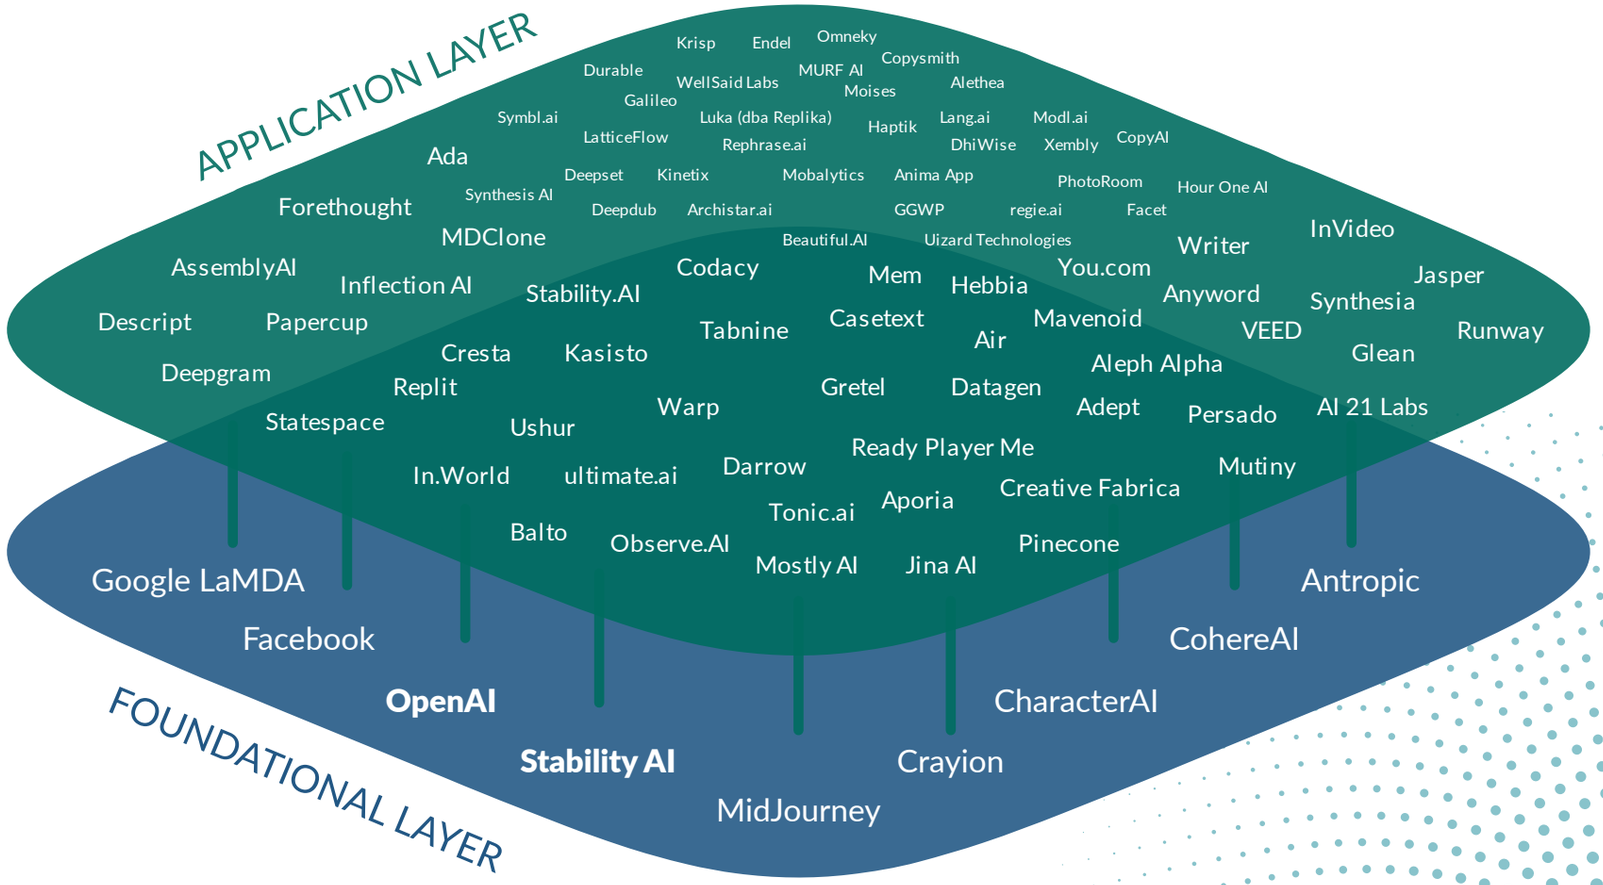

4. CLOSED-SOURCE
Proprietary Generative AI

OpenAI

Founded by *Elon Musk* and the *Y Combinator* president Sam Altman, OpenAI rose to quick international fame when they launched ChatGPT in November 2022. Within a week, the application saw a spike in usage of over a million users. Being able to code and interact in a way that mimics human intelligence, ChatGPT has surpassed previous standards of AI capabilities and has introduced a new chapter in AI technologies and machine learning systems.

OpenAI rushed to reveal their model, but similar and probably as powerful models already existed. *Lemoine*, a software engineer at Google, who had been working on the development of LaMDA, shared his interactions with the program, in a *Washington Post* article, causing a stir. Lemoine recounted many dialogues he had with *LaMDA* in which the two discussed topics that ranged from technical to philosophical issues. These led him to ask *if the software program is sentient*. Lemoine was eventually fired.

Launched in 2015, before the introduction of any Generative AI concept, OpenAI witnessed the collaboration between, Musk and Altman on one side and other players in Silicon Valley the likes of Peter Thiel and LinkedIn founder Reid Hoffman who pledged close to a billion **Dollars** for OpenAI that year. Two major projects formed the cornerstone of Open AI. These were:

The Generative Pre-Trained Transformer (GPT)-3 model has gained a lot of hype. It is essentially a language model that leverages deep learning to generate human-like text. Along with machine-generated texts, it can also produce stories, poems as well as codes. It is deemed as an upgrade to the previous GPT-2 model, released in 2019, which is a large transformer-based language model with 1.5 billion parameters, trained on a dataset of 8 million web pages. To put it simply, language models are a set of statistical tools that enable such technology to predict the next word or syntax of the sentence.

In 2021, Open AI took a big leap and created DALL-E, an AI tool that can create some stunning masterpieces. Within a year, it launched Dall-E2, which provides images with 4 times greater resolution. Dall-E2 was able to create realistic images and art from a description in natural language. This human-like robot hand can paint artworks that merge concepts, attributes, and style. Moreover, Dall-E2 can build upon existing art pieces and create new expanded original canvases, generating different variations of a previous image.

Two pictures famously generated by Open AI Dall-E algorithm.

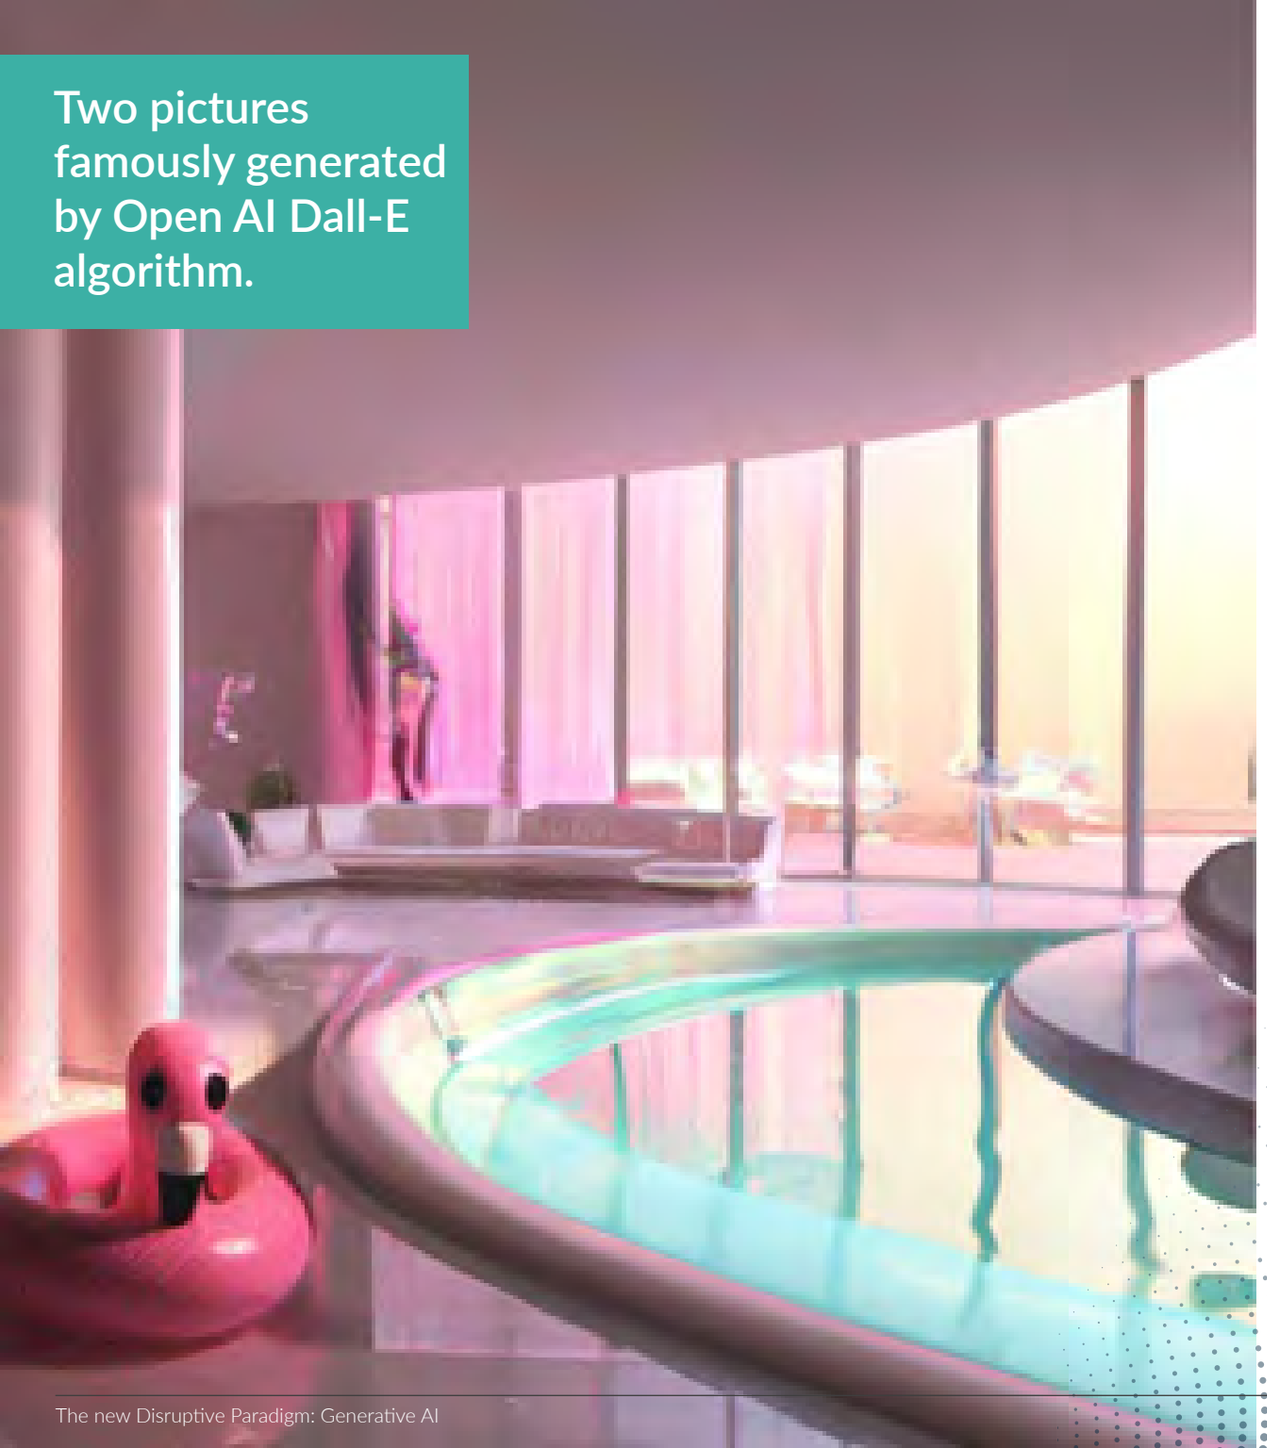

At the start, “Open AI” was to be a non-profit AI research company, with a mission “to ensure artificial general intelligence benefits all of humanity”, with no hope for profit.

Today, Open AI is governed by a board that comprises employees Greg Brockman, Ilya Sutskever and Sam Altman (CEO). It also has non-employees: Adam D’Angelo, Reid Hoffman, and others. Meanwhile, key strategic investors include Microsoft Corporation.

By having Microsoft fund Open AI “in billions” and wanting in return massive profits from direct operation meant that all hopes of Open Software got completely dashed.
Indeed, Open AI has now been working exclusively on its proprietary algorithm and works with Microsoft under strict licensing agreements.

Elon Musk has left the venture.

5. OPEN-SOURCE DIRECTION

Open Source Success

During the past decade, adoption of open-source software at the enterprise level has mainly flourished.

Businesses, in direct interaction with businesses (B2B) or Clients (B2C), discovered the many advantages open-source solutions hold over their proprietary counterparts. There are many examples of this, with the most significant being popularly used operating systems, such as Linux for enterprises and Android for mobile. The former, itself based upon the open-source Unix operating system, has captured 50% of the enterprise market share, while the latter 80% of the mobile applications market share.

Enterprises have sought make efficient use of open-source software for many reasons reasons listed herein.

In addition to these advantages, open-source software seems to offer better long-term viability in comparison to proprietary developers that may be transient. This is largely due to the presence of supportive communities focused on continually introducing innovations.

Community
Open-source solutions geared towards the enterprise are often supported by thriving communities, bound by a common drive to encourage and improve a solution that both the enterprise and the community benefit from. These global communities collaborate to effectively develop and implement advancements.

Transparency
This allows full availability of the code as well as discussions surrounding its features, bugs and solutions. This results in an advantage over proprietary software which is private and maintained by a single team within a large corporation.

Reliability
With more scrutiny, the reliability of open-source code tends to be superior, also allowing it to be well integrated. The output tends to be extremely robust, tried, and tested, with open source now running almost 90% of the internet.

Better Security
Open-source software's code is proven to be secure given its thorough scrutiny by various members of a community of developers, with any issues that do arise being addressed more diligently.

Cost Effective
Although open-source solutions should be thought of as more than just free software, the fact that they require no licensing fees remains a decisive advantage when looking at the total cost of deploying a solution.

Freedom from Lock-in
Proprietary software for core infrastructure increases the risk of being locked in by the vendor or technology. Enterprises then have no choice but to agree to vendors' price increases and experience a lack of flexibility.

Becoming the Norm
There are many large enterprises implementing open-source solutions. These are now introducing resources to support open-source solutions, thus helping to develop open source solutions.

Open-Source Generative AI

At the heart of open-source generative AI is Stable Diffusion which relies upon the introduction of diffusion models.

Stable Diffusion: The Open-source Model

As mentioned in *Section 2*, Open Ai's DALL-E 2, text-to-image system, can create photos in different styles, renders, etc. They first used ***Diffusion***.

Diffusion processes, inspired from physics, are processes whereby a substance in a medium "***diffuses***" from a region of higher density to a region of lower one. One might think, for example, of blue ink diffusing in water and eventually turning all the water inside a container a light bluish color. Physicists have modeled diffusion mathematically, using complicated models, and it has been extensively used in areas like finance to model the randomness of financial markets. Similarly, images can be transformed into a uniform distribution by randomly adding noise.

Diffusion, in the physical world, is an irreversible process — ink diffused in water can't be restored to its original form. Counterintuitively, diffusion systems, as employed in machine learning, are a great asset: the model learns to figure out a "reverse diffusion" process, restoring the destroyed image and recovering the data from the noise. In fact, by adding the noise artificially, the diffusion system actually learns how to remove it.

Diffusion systems have been around for a while. However, it is only recently that Open AI introduced CLIP ("***Contrastive Language-Image Pre-Training***"), which classifies data - for example, images - to "score" each step of the diffusion process based on how likely it is to be classified under a given ask (e.g., "draw an astronaut on the lunar surface").

At the start, the image has a low CLIP score because it's mostly noise. However, as the diffusion system reconstructs data from the noise, it slowly moves towards matching the prompt.

Open AI introduced CLIP alongside its DALL-E2. However, this had its limitations:

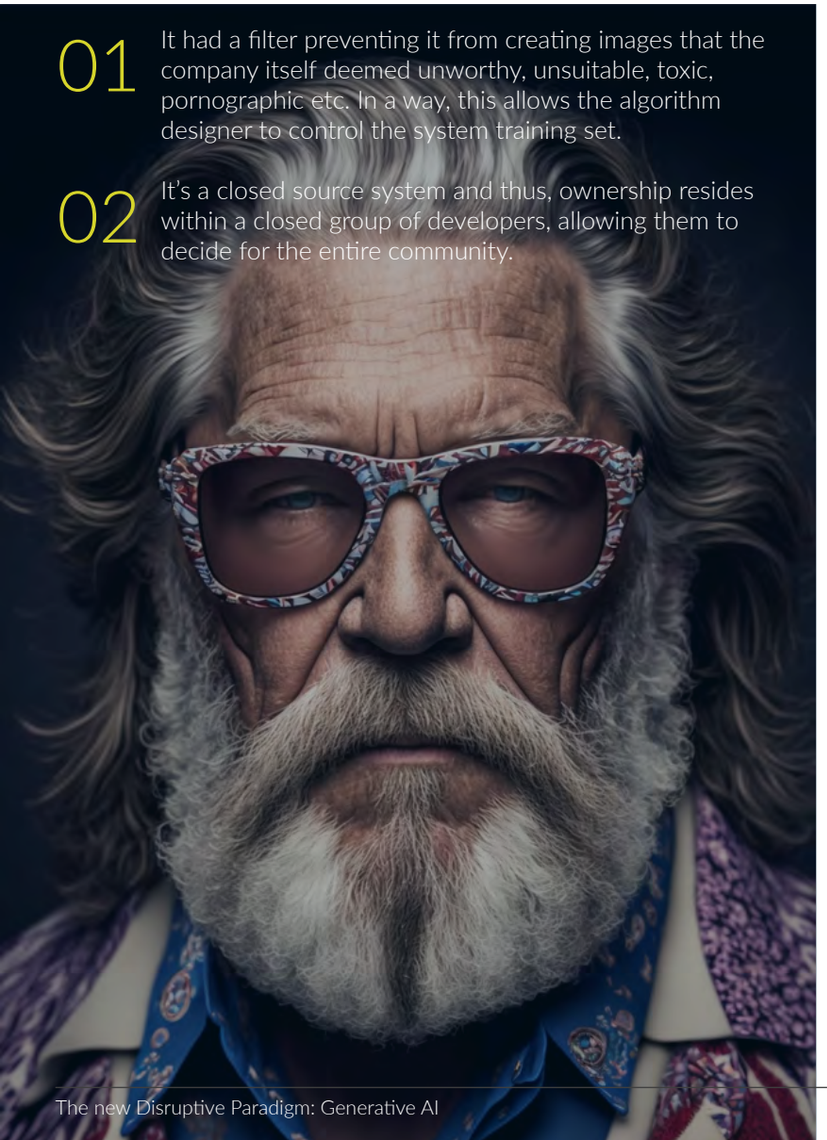

01 It had a filter preventing it from creating images that the company itself deemed unworthy, unsuitable, toxic, pornographic etc. In a way, this allows the algorithm designer to control the system training set.

02 It's a closed source system and thus, ownership resides within a closed group of developers, allowing them to decide for the entire community.

Enter **Stable Diffusion**:

Firstly, there are no such filters, no content filtering. Secondly, having an Open-source model:

1) Sets the standard for the image-to-text model, and allows the whole community to benefit from it and contribute to it.
2) With open code, the wealth of feedback from all users worldwide supports considerable improvements in the model and allows better tuning of model parameters. Stable Diffusion is a **latent text-to-image** diffusion model capable of generating photo-realistic images given any text input. It is a *Latent Diffusion Model* that uses a fixed, pretrained text encoder (CLIP ViT-L/14), as suggested in the seminal paper of:

Rombach et al. High-Resolution Image Synthesis with Latent Diffusion Models - *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR).

Stable Diffusion was created by **Emad Mostaque**. He established the London-based ***Stability AI***, a collaboration with media creation company *RunwayML*, Heidelberg University researchers, and the research groups **Eleuther-AI** and **LAION**. Stability AI is currently valued at **USD 1bn** *(see Section 7)*.

Created in 2020, realizing the advent of new AI technologies and more the lack of an organized "Open-source" environment in the AI community, Mostaque's vision was to create an open ecosystem where further models can be built upon the latent space concept.

With millions of downloads on **GitHub**, Stability, the open-source computing forum, is at the forefront of Gen-AI.

Open-source systems require publicly available datasets and therefore, ***Stability AI relies upon the LAION 5B dataset.***

A dataset consisting of 5.85 billion CLIP-filtered image-text pairs, featuring several nearest neighbor indices, an improved web-interface for exploration and subset generation, and detection scores for watermark, NSFW, and toxic content detection. 2Bn images are marked in English.

According to David Ha, of Stability "even Google uses Laion5b as its model training set".

The initial version of Stable Diffusion was based on LAION-400M, the predecessor to LAION 5B, which was known to contain depictions of sex, slurs and harmful stereotypes. LAION5B corrected this.

Stability AI uses a cluster of 5,000 Nvidia A100 GPUs running in AWS to train Stable Diffusion using the LAION5B dataset.

The prediction, however, can run on a lower spec Nvidia chip costing less than USD 300, while the Stability team has perfected the image production to run on a laptop sized machine in less than a quarter of a second. In the future, Stability aims to achieve the capability of a superhuman using a couple of gigabyte applications on ordinary desktops.

Revenue Generation Model

The business model of Stability-AI rests on three pillars:

01 The open-source foundational layer and maintaining the ultimate source code reference (B2B).

02 Services (B2C & B2B) that live in the application layer provide a wealth of services based on Stability proprietary ventures.

03 Infrastructure layer or SaaS (B2B, B2G): that could be offered to corporates or governments and providing the full spectrum including training on a specific proprietary dataset, that generally can be the property of a government or larger institution.

Stability is also currently working on multiple proprietary projects, including AI models for generating audio, music and even video, with a clear focus on media in the future.

> Stability AI will make money by training “private” models for customers and acting as a general infrastructure layer.
>
> *Emad Mostaque*

6. Addressable Market Size

The media industry, as described in the next two sections of the report, has multiple segments that combine into one vertical (commonly data on media is combined with data on the entertainment sector and with the internet).

- Film
- Television and home video
- On-demand video streaming
- Music and radio
- Newspaper and magazines
- Gaming
- Virtual reality

Drivers within the industry vary by segment and across geographic regions, with the shift to digital consumption and emerging technologies being a major driver of change now.

The global media industry recorded a decline of 2.3% in 2020 as it dealt with the Covid-19 pandemic. It then continued its upward trajectory and was estimated at around USD 2.5 trillion in 2022 and is forecast to increase to around USD 3 trillion by 2026, at an estimated CAGR of 4.6%.

Global Media Revenues ($B) and Year-on-Year Growth

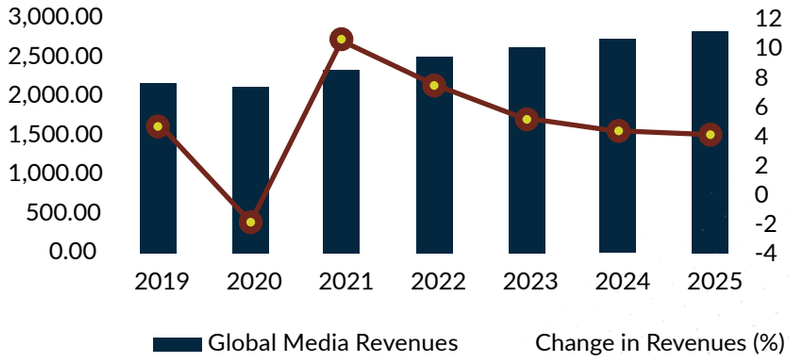

Sources: Statista, pitchbook, PwC

Media revenues are generally classified into three categories: Consumer spending, advertising, and others (internet access, etc.).

In 2018, while consumer spending contributed a major chunk (40%) to revenues, advertising accounted for only 29% of total revenues. (Figure 11) However, in 2022, advertising contributed 35% to the share of total revenues and is expected to reach 40% of revenues by 2030. Both streams of revenues are directly impacted by Gen-AI. However, declines in consumer spending are compensated by an increase in advertising, such that going forward, CAGR is expected to be at a respectable, stable 6%. Hence, the revenue breakdown of the media sector is such that 70% is attributable to **consumer spending** and **advertising** combined (with the latter forecast to reach 40% by 2030) and 30% to internet access.

The following chart shows the breakdown of consumer spending revenues for the sector.

Global Consumer Revenues by Segment (US$ trillion)

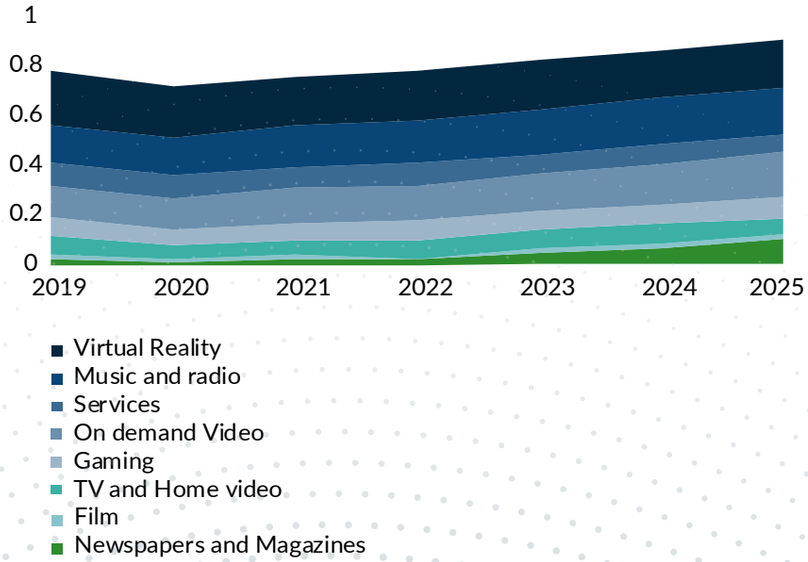

Advertising serves as a major contributor to revenues for the media sector. Meanwhile, internet advertising spending is considerably higher and is expected to reach US$336 billion by 2025, a 9% CAGR increase since 2021. Although slower than the 16% growth recorded in 2019, it still represents a massive disruptive opportunity for Stability. Even as traditional components such as television stagnate or decline, overall advertising revenues are expected to rise to roughly US$800 billion in 2025.

Global Advertising Revenues by Segment (US$ trillion)

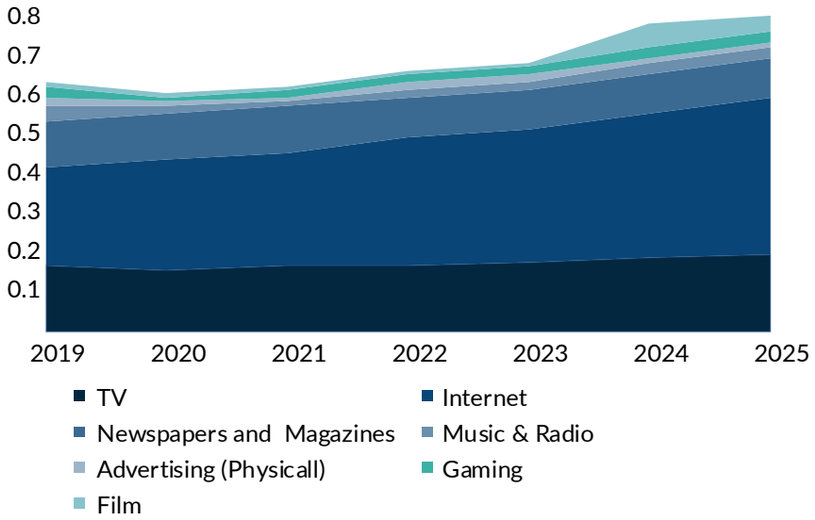

Media Spending per Capita (US$)

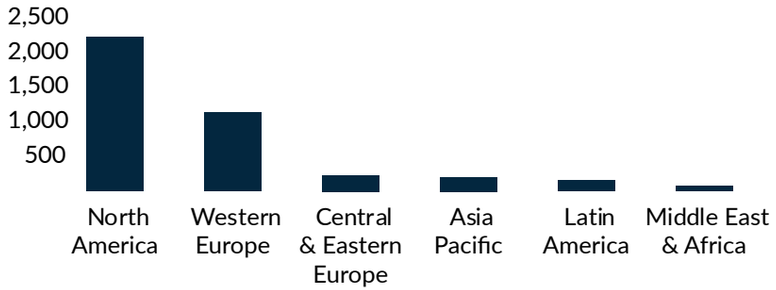

There have always been big variations in consumer spending revenues across demographics, regions, etc.

As visible in Figure 10, there is an opportunity for huge growth in developing markets where large populations of young consumers are forming the consumer base. To target these consumers, an approach different from the one employed for penetrating seasoned markets, such as the US and Europe, will have to be employed. These new markets of young consumers are also more likely to adopt new generative AI applications.

Global Video Games Revenue, by Segment

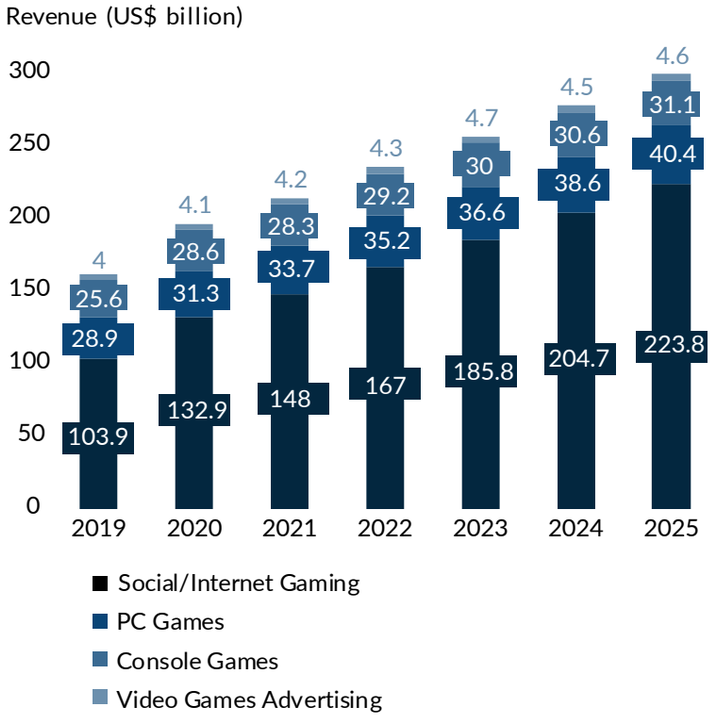

The video games segment, most popular amongst youth, is one of the sectors experiencing the most significant growth. In 2021, video games revenues (excluding e-sports) reached US$214.2 billion, and is expected to rise at an 8.4% CAGR to US$321.1 billion in 2025 (highest among all subsectors).

All gaming subsegments are likely to be affected by the disruption brought forward by Gen-AI. The ability to produce new games is the industry's main aim. Moreover, recyclability is very important within video game production, with mass production requiring component reusability. **RunwayML** has launched a new app that lets users erase video parts within a given timeframe and replace it with whatever users want to: for example, users can replace a bee standing on a leaf with an ant. These edits can then be reused.

Global Film Revenues and Admissions (billions)

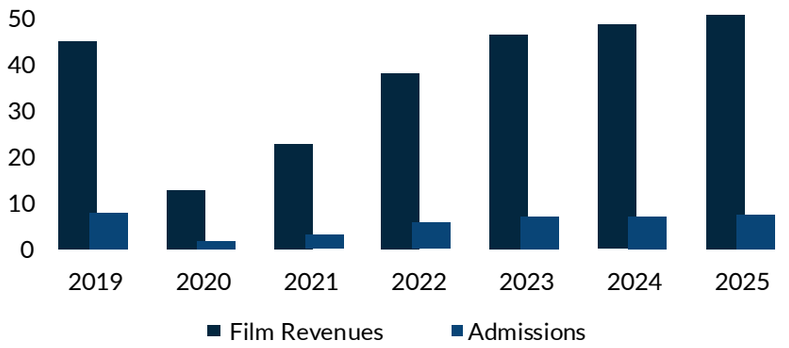

The above chart highlights the fact that film revenues declined in 2020 due to Covid-19 and are only expected to rise back to pre-Covid levels in 2023. This rise will be on the back of the North American market, which is expected to register a 12% growth in 2023, as well as the Asia Pacific and Latin American regions. Russia, however, is expected to register slow growth.

Disruption in the film and gaming industries

We researched in detail the structural constituents of the media and entertainment industry, and within it, the **film** and **gaming** segments as these are the ones that generate most of the profits in the sector. Disruption in these two subsegments will effectively disrupt the whole sector.
In this section of our report, we delve deeper into the enormous impact GenAI will have on the film and gaming industries.

Frame-frame animation in film

As an example, **Labs.monks**, which tracks digital innovation in the media world, reported last month, through a concrete use case, how *frame-by-frame animation*[2], a method usually overlooked as slow and inefficient in making films, was able to generate 60 frames in one minute—rather than one frame in 60 minutes—making it incredibly easy to generate useful watch time "in the blink of an eye". Moreover, this was achieved without compromising quality. Using *Stable Diffusion* allowed the filmmakers to create a large volume of backgrounds that could be swapped easily to test how they looked. This also works for developing themes and backgrounds for a particular movie clip, ending up like "describing a specific shot needed" and then being able to test hundreds of different themes', camera angles, artistic styles, etc.
The inspiration towards such an approach was "rotoscoping", which is a known method whereby an artist traces images over a current existing footage and then translates this existing recording from one style to another. Generative AI has excelled here and complemented creativity instead of replacing it. The work is completed in a fraction of time compared to current industry standards.
This use case can be extrapolated to hundreds of other ones, resulting in the film industry being shaken to the core. Gen-AI seem to affect almost every method and technology, be it in 2D, 3D Artwork, 3D textures, music, dubbing, etc.

[2] https://media.monks.com/articles/frame-frame-look-how-generative-ai-supercharges-creativity

The revolution in games, GenAI
The 2D image generators of Stable Diffusion, have captured most of the popular excitement over Gen-AI, due to, among others, the eye-catching nature of the generated images. Already, Gen-AI models within the applications layer address almost all tools involved in games: 3D models, character animations, dialog, music.
In addition to our example, reusability of backgrounds, through erasing items in a video timeframe, the speed of production and wealth of tools makes new production extremely efficient, as discussed in *Section 6* on RunwayML.

As an example, **Red Dead Redemption 2**, one of the most expensive games ever produced, cost half a billion dollars, but produced a stunning virtual world, which was very close to reality. It took eight years to produce, included 100 characters and featured 60 hours of music created by over 100 musicians. This is in contrast to Microsoft's **flight simulator** which teamed with Blackshark.ai to generate realistic images and 3D photos of the world using 2D satellite images. They managed to produce it in a matter of months with great success.

Today, there are applications pertaining to almost every subsegment of the games industry. These include **the list on the right**.

GenAI, excellent at "ideation", seems to help game designers by allowing concepts and ideas to be explored quickly, thus supporting conceptual artwork. This is a key part of the production process. Several studios are using several of these tools to radically speed up their concept artwork, resulting in creating images in a single day, a process that would have taken a month previously.

2D IMAGES

Stable diffusion
Dall-E

3D DESIGN

Hypothetic
Kaedim3D
Mirage

TEXTURE

Barium AI
Ponzu
ArmorLab

ANIMATIONS

Kinetix
Plask
DeepMotion
WonderDyamics
Radical

WORLD BUILDING

Mixar
PromothenAI
Moatboat

- First, the studio's game designers use Stable Diffusion to explore different ideas and generate images they find inspiring.
- These get turned over to a professional concept artist who assembles them together and paints over the result to create a single coherent image, which is then fed into Stable Diffusion again to create a bunch of variations.
- Once these variations are agreed upon, they pick one, paint in some edits manually and repeat the process until they're happy with the outcome.
- Finally, with the image is passed through Stable Diffusion again, upscaling it to create the final piece of art.

This is a concrete example of how developments are unfolding in the entire industry, radically changing the industry's modus operandi.

At present, these tools are simply augmenting human efforts and capabilities with new tools and time-saving advanced machinery.

The two use-cases above highlight the impact of the introduction of GenAI on the sector. These are just two of many that show the profound impact of open-source Stable Diffusion. The full application layer is inspired by Stable Diffusion and is adopting it with lightning speed.

Additionally, as we try to quantify the addressable market size, the penetration of GenAI is such that almost 100% of films and games will use genAI in some shape or form, such that eventually there will be no production not utilizing the tools that it brings to market. There will be a full 360 degree impact on the industry and all revenue that currently stem from the industry will constitute the TAM of the sector.

The next section breaks down the media sector revenues and derives the total addressable market size of GenAI.

Quantifying the addressable market

US$ billion

Segments	Sub-segment	Gen-AI	2019	2020	2021	2022	2023	2024	2025
Virtual Reality	VR & AR & metaverse	✓	10.0	13.0	17.0	20.0	35.0	45.0	100.0
Gaming	Social Gaming	✓	103.9	132.9	148	167	185.8	204.7	223.8
	PC Games	✓	28.9	31.3	33.7	35.2	36.6	38.6	40.4
	Console Games	✓	25.6	28.6	28.3	29.2	30	30.6	31.1
	Integrated Video Games Advertising	✓	4	4.1	4.2	4.3	4.7	4.5	4.6
Film	Film	✓	45.2	12.7	22.8	38	46.4	48.7	50.7
	Admissions	x	7.9	1.9	3.3	5.7	7.1	7.3	7.5
TV and Home Videos	Classical TV, videos, others		220	210	205	201	196	190	180
On Demand Video Streaming	Subscription	✓	37	50	54	62	67	74	80
	Transactional	x	8	9	10	10	11	12	13
Music	Live Music	✓	28	8	15	26	27	29	30
	Physical Recorded Music	x	5	6	5	6	5	4	3
	Music Ringtones	x	1	1	0	1	0	0.5	0
	Digital Download	x	2	2	2	2	0.5	0.5	0
	Digital Streaming	✓	15	16	15	22	22	24	26
Newspaper and Magazines	Print, Physical	x	85	83	80	80	78	77	75
Total Consumer Spending Revenues		x	616.50	596.50	626.30	689.40	717.10	745.40	765.10
CAGR Consumer Spending (2023-26)	*3.90%*								
Total Advertising Revenues		✓	640.40	610.50	630.40	670.50	690.50	800.00	820.00
CAGR Advertising (2023-26)	*6.50%*								
Internet and Other Revenues		x	918.10	916.00	1,083.30	1,150.10	1,230.40	1,201.60	1,249.90
CAGR Advertising (2023-26)	*3.50%*								
Total Media Market Size			**2,175.00**	**2,123.00**	**2,340.00**	**2,510.00**	**2,638.00**	**2,747.00**	**2,835.00**
Total Addressable Market (US$ billion)						**1,074.20**	**1,145.00**	**1,299.10**	**1,406.60**

We have, so far, provided a review of the global media sector and defined its major constituents and revenues stemming from each subsegment. In the table above, we summarize the analysis for computing the addressable market share for Stability AI.

Methodology: A historical, *top-down* approach has been employed to estimate the addressable market share. For projections, CAGR estimates have been used. Industry segments, and then subsegments were identified. Any further breakdown would have resulted in identification of subcategories which may be intertwined, thus resulting in double counting.

Subsegments which would be impacted by Gen-AI disruption and presented an opportunity for penetration by Stability-AI were identified and their market share added to the addressable market share. These included the advertising industry and its subsegments, as well as internet access companies.

We have therefore computed the total addressable market for Generative AI in the Global Media & Entertainment industry to be US$ 1.1 trillion in 2022 and expect it to touch US$ 1.4 trillion in 2025.

> We have barely started scratching the surface of what the future will hide for the Generative AI Era
>
> ***Nathan Lile***

7. VENTURE CAPITAL (VC) The Excitement

Initially, much of the hype stirring VCs cash liquidity, had focused on the stealthy adoption of OpenAI's ChatGPT, gaining a million users in the first five days of its release in November 2022.

This generated so much interest that the main sector that VC investors were investing in during the fall was Generative AI startups. The generative AI market's underlying similarities and approach to a multi-layer model has appealed to VC investors who previously invested in blockchain, which offered a familiar model of investment. Gen-AI has indeed filled the gap that was created by failing cryptocurrency and blockchain startups.

Investment in Gen-AI by VCs increased by over 400% since 2020, touching nearly US$ 2.2 billion by the end of 2022, as per Pitchbook data, with investors convinced of the massive transformative potential of the technology. As per Pitchbook research, the global market for AI-augmented content creation tools is projected to increase by an overall 17% till 2025. If anything, actual growth is likely to be higher than this.

The question that then remains is: Does this enthusiasm translate into Series As, Bs, Cs and Ds?

The community seems to be of the view that Gen-AI tools are relatively easy to build with much help from foundation models' API. Hence, they represent a safe revenue model.

"If one looks at other historical major technological shifts like the internet, mobile or crypto, there were, at the time considerable skepticism, this time they seem to agree that Gen-AI is the real next big thing."

As we show in the following graphs, for a pool of VC firms, most investors have at least one, if not more, investments in a GEN-AI company, with some having invested in as many as thirty companies. These include Y combinator, Sequoia, Lightspeed Ventures and Andreessen Horowitz. We have compiled a list of more than 650 investors in the sector.

Venture Capital Investment in Generative AI

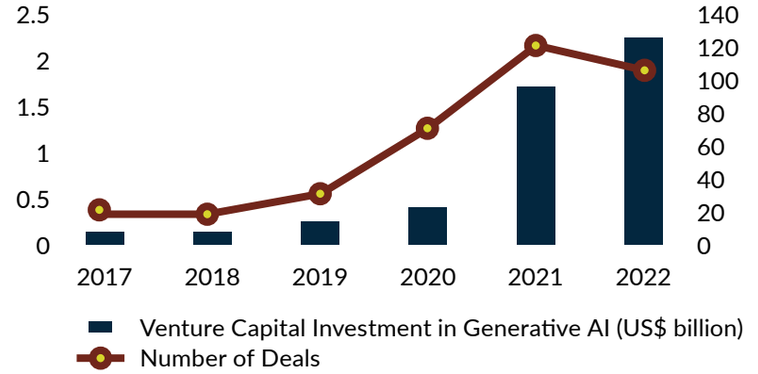

Top VC Investors

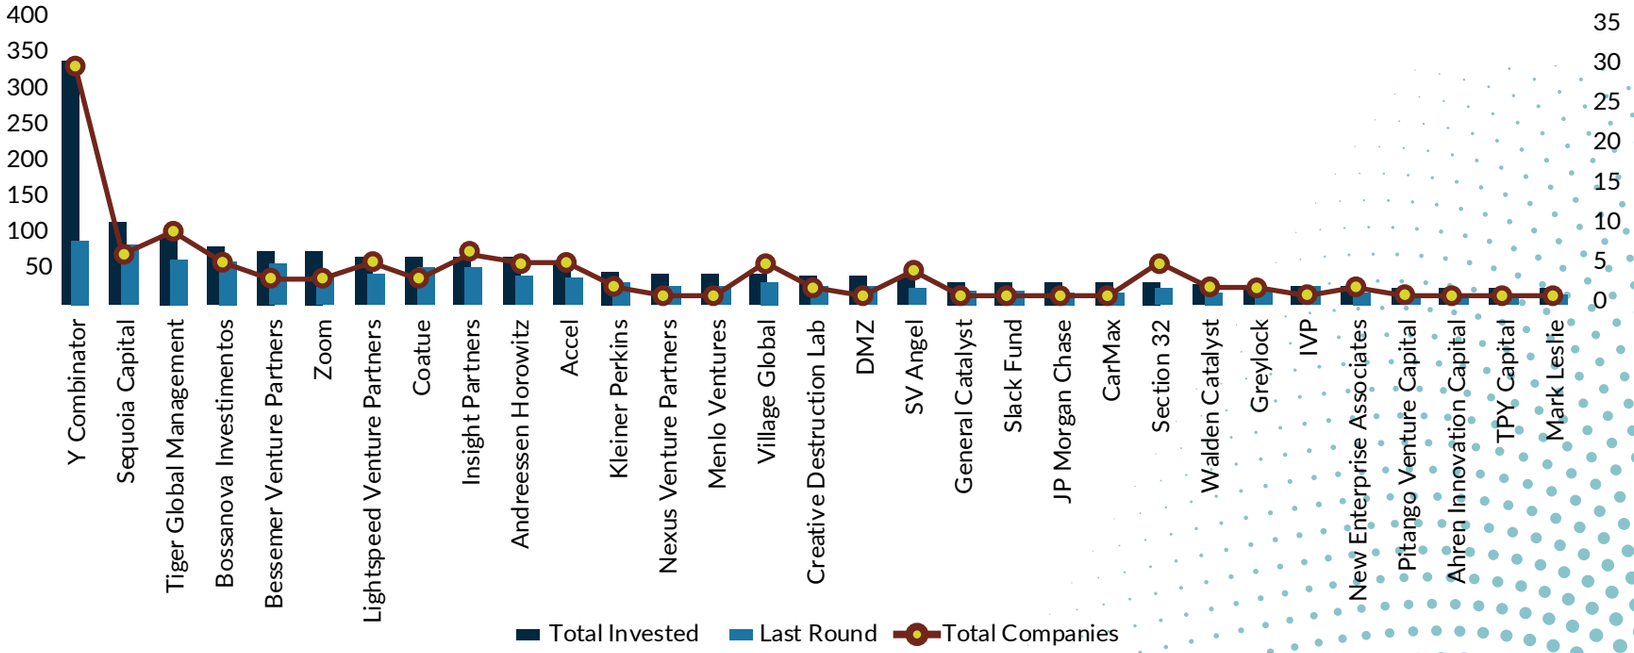

8. CAPITAL INVESTMENTS per Segment and Category

For the purpose of this report, over 200 companies, involved directly and indirectly in Gen-AI, were researched and data collected on them. Furthermore, 643 investors, who have invested in these companies over the course of different stages of Series funding as well as Debt and Convertible bonds, were identified.

To date, the aggregate size of investment in the generative AI sector is 6,399 Bn, with 5 Bn allocated to Text and Language, and 1.3 Bn to media.

As the following two tables highlight, Stability AI leads the way in terms of attracting investment in both image and media. It has garnered considerable interest from the VC community, proof of the potential of Stable Diffusion and its open-source model. Being open-source, Stability AI's model is likely to continue benefiting from the introduction of significant amounts of data in its models, which allows it to produce content based on previous word sequences or images, resulting in improved model performance.

IMAGE

#	Company Name	Total Amount	Last Funding Amount
0	Stability.ai	101,797,885.00	101,797,885.00
1	Creative Fabrica	56,300,000.00	12,200,000.00
2	PhotoRoom	19,125,000.00	19,000,000.00
3	Uizard Technologies	18,600,000.00	15,000,000.00
4	Beautiful.AI	16,000,000.00	11,000,000.00
5	Facet	15,500,000.00	13,000,000.00
6	Kive	8,700,000.00	7,000,000.00
7	Bubbles	7,500,000.00	5,500,000.00
8	Common Sense Machines	5,079,994.00	5,079,994.00
9	Scenario	5,000,000.00	0
10	Pencil	3,850,000.00	2,750,000.00
11	Pinscreen	3,800,000.00	416,000.00
12	Hypar	2,480,000.00	1,480,000.00
13	Rosebud AI	1,500,000.00	0
14	Kaedim	203,561.00	0

MEDIA

#	Company Name	Total Amount	Last Funding Amount
0	Stability.ai	101,797,885.00	101,797,885.00
1	Descript	100,000,000.00	50,000,000.00
2	Statespace	96,500,000.00	50,000,000.00
3	Runway	95,500,000.00	50,000,000.00
4	Deepgram	85,920,000.00	47,000,000.00
5	Ready Player Me	71,100,000.00	56,000,000.00
6	Synthesia	66,600,000.00	50,000,000.00
7	AssemblyAI	63,120,000.00	30,000,000.00
8	Creative Fabrica	56,300,000.00	12,200,000.00
9	InVideo	52,500,000.00	35,000,000.00
10	VEED	35,000,000.00	35,000,000.00
11	Papercup	33,234,603.00	20,000,000.00
12	Hour One AI	25,000,000.00	20,000,000.00
13	Endel	21,221,791.00	15,000,000.00
14	Deepdub	20,000,000.00	20,000,000.00
15	Aimi	20,000,000.00	20,000,000.00

Players in segments raised amounts above US$ 10 million. Stability still has a long way to go in raising the right amount of funds that matches its mission.

Within generative AI, venture series funding plays the major role in raising funds for companies. As the transformative nature of text-to-image begins to be realized, venture series funding for the field is expected to increase significantly as more disruptive business prospects appear. The most obvious application for this technology is art. This tool could be applied for a range of purposes, ranging from generating art for book covers to digital media. This will unleash a wealth of opportunities and likely result in venture series funding being increasingly directed towards the field of image, within which Stability AI operates, in the future.

SERIES C

Ada	190,620,620.00
Glean	155,000,000.00
Cresta	151,000,000.00
MDClone	104,000,000.00
Descript	100,000,000.00
Statespace	96,500,000.00
Runway	95,500,000.00
Forethought	92,000,000.00
Kasisto	81,450,000.00

During 2021, the field of text-to-image AI, which allows original images to be produced from text commands, started to explode. By removing the need for technical labor for generating images, these tools signify a major shift within AI and are likely to attract increasing amounts of funding in the future, with text-to-image AI expected to be the next big thing. This will also lead to the establishment of new startups within the field, attracting further investment from venture capitalists. The number of companies within image is already significant (representing over 10% of total AI companies).

Funding raised in the Global Generative AI space

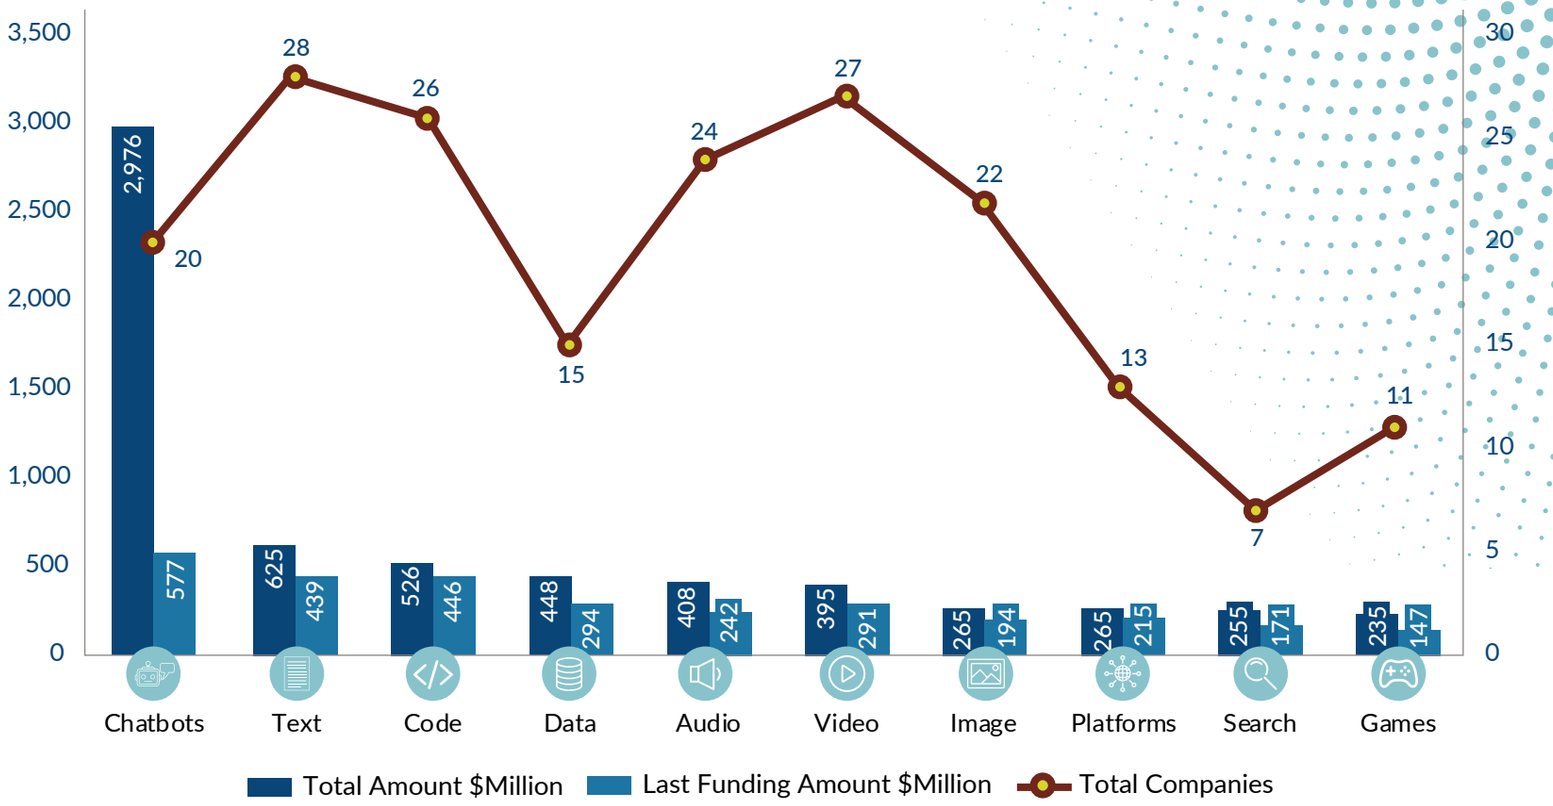

Startup Companies per Segment

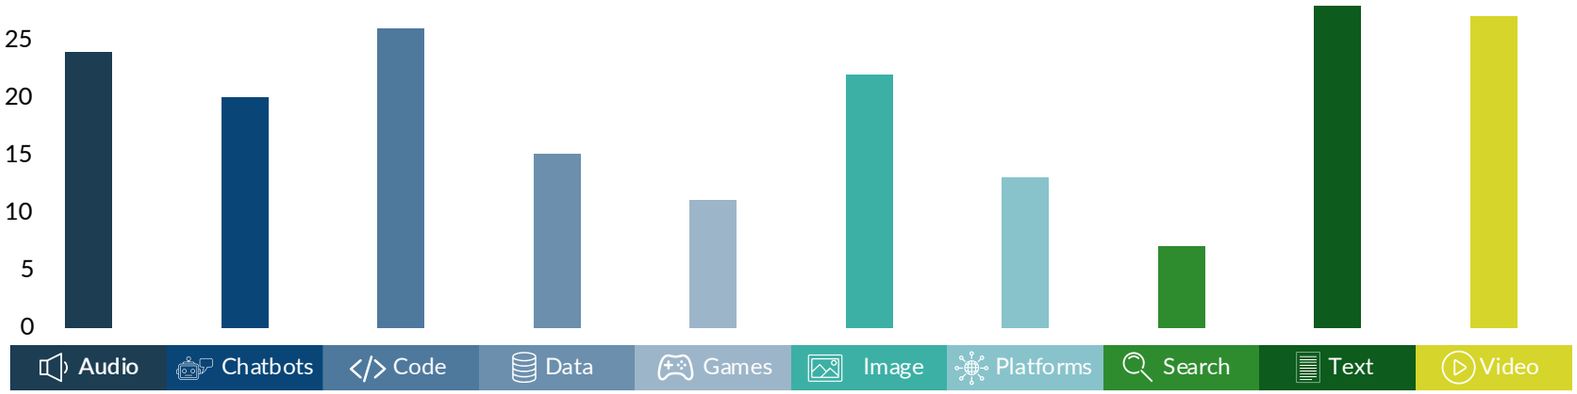

9. APPLICATION LAYER
A New Universe for Startups

The universe of gen-AI has expanded greatly during 2021 and 2022, with many startups beginning to monetize gen-AI tools. These include a significant number in the chatbots, text and video segments. In the image segment, Stability AI leads the way, with a total funding of over US$ 100 million.
The expansion of the application space has been aided by the solid improvements foundational models , and the ease of implementing a specific solution using the exposed API's.

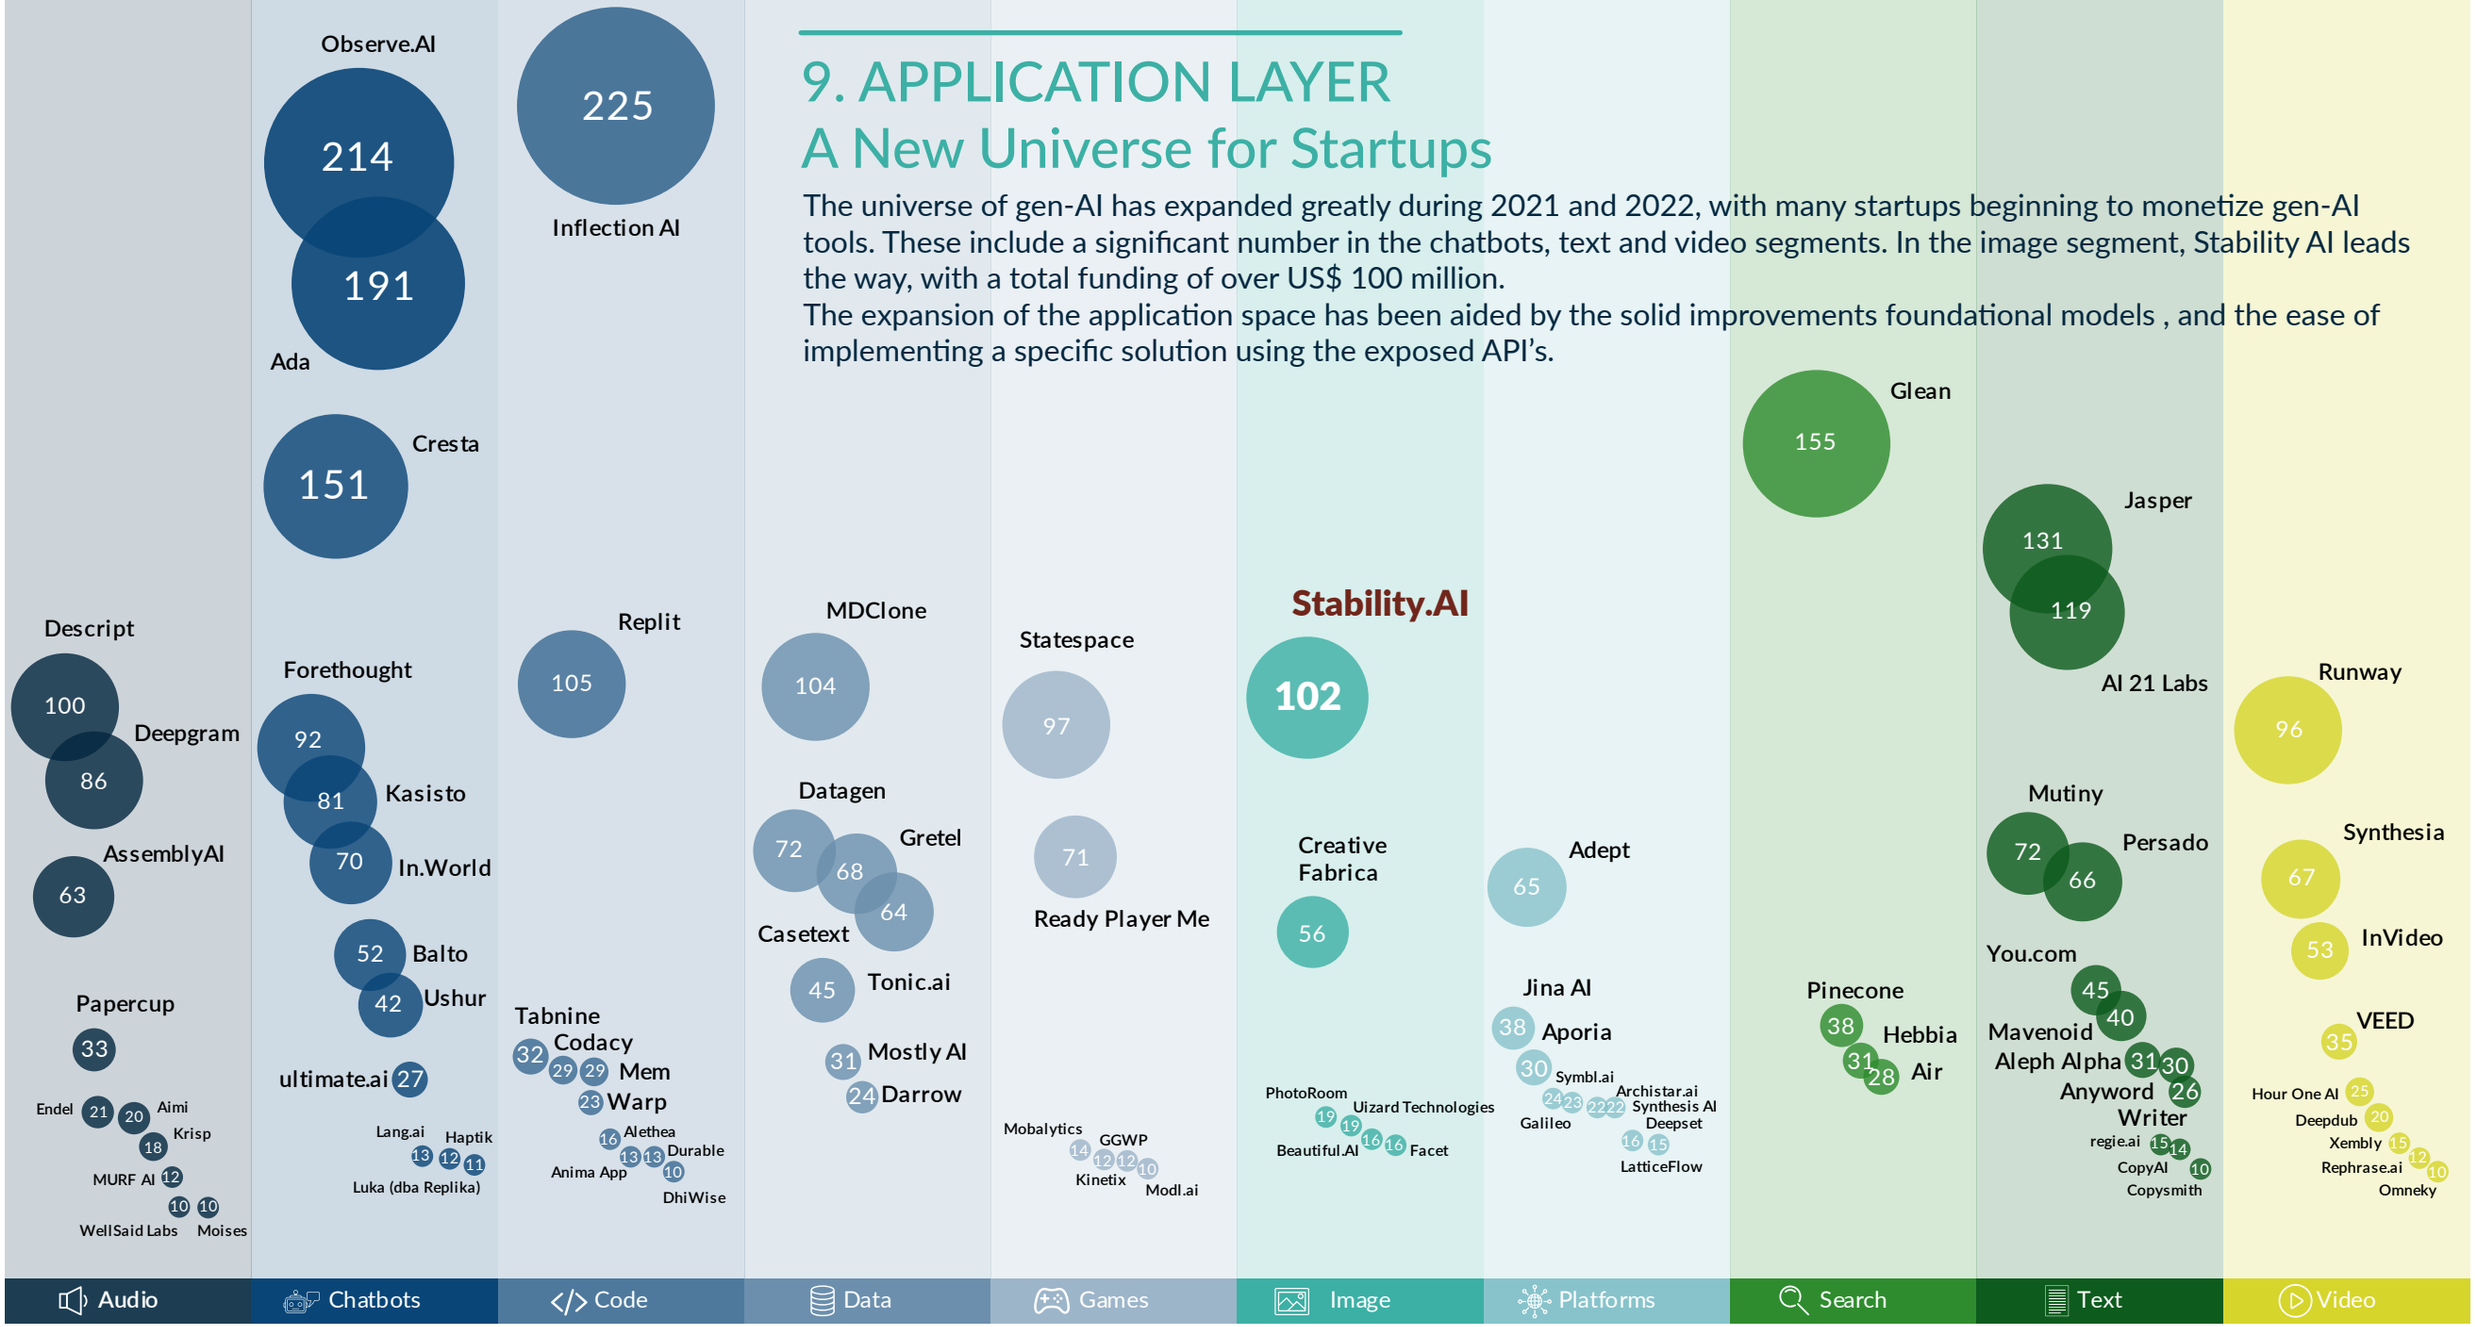

Scatter Plot per Business Type

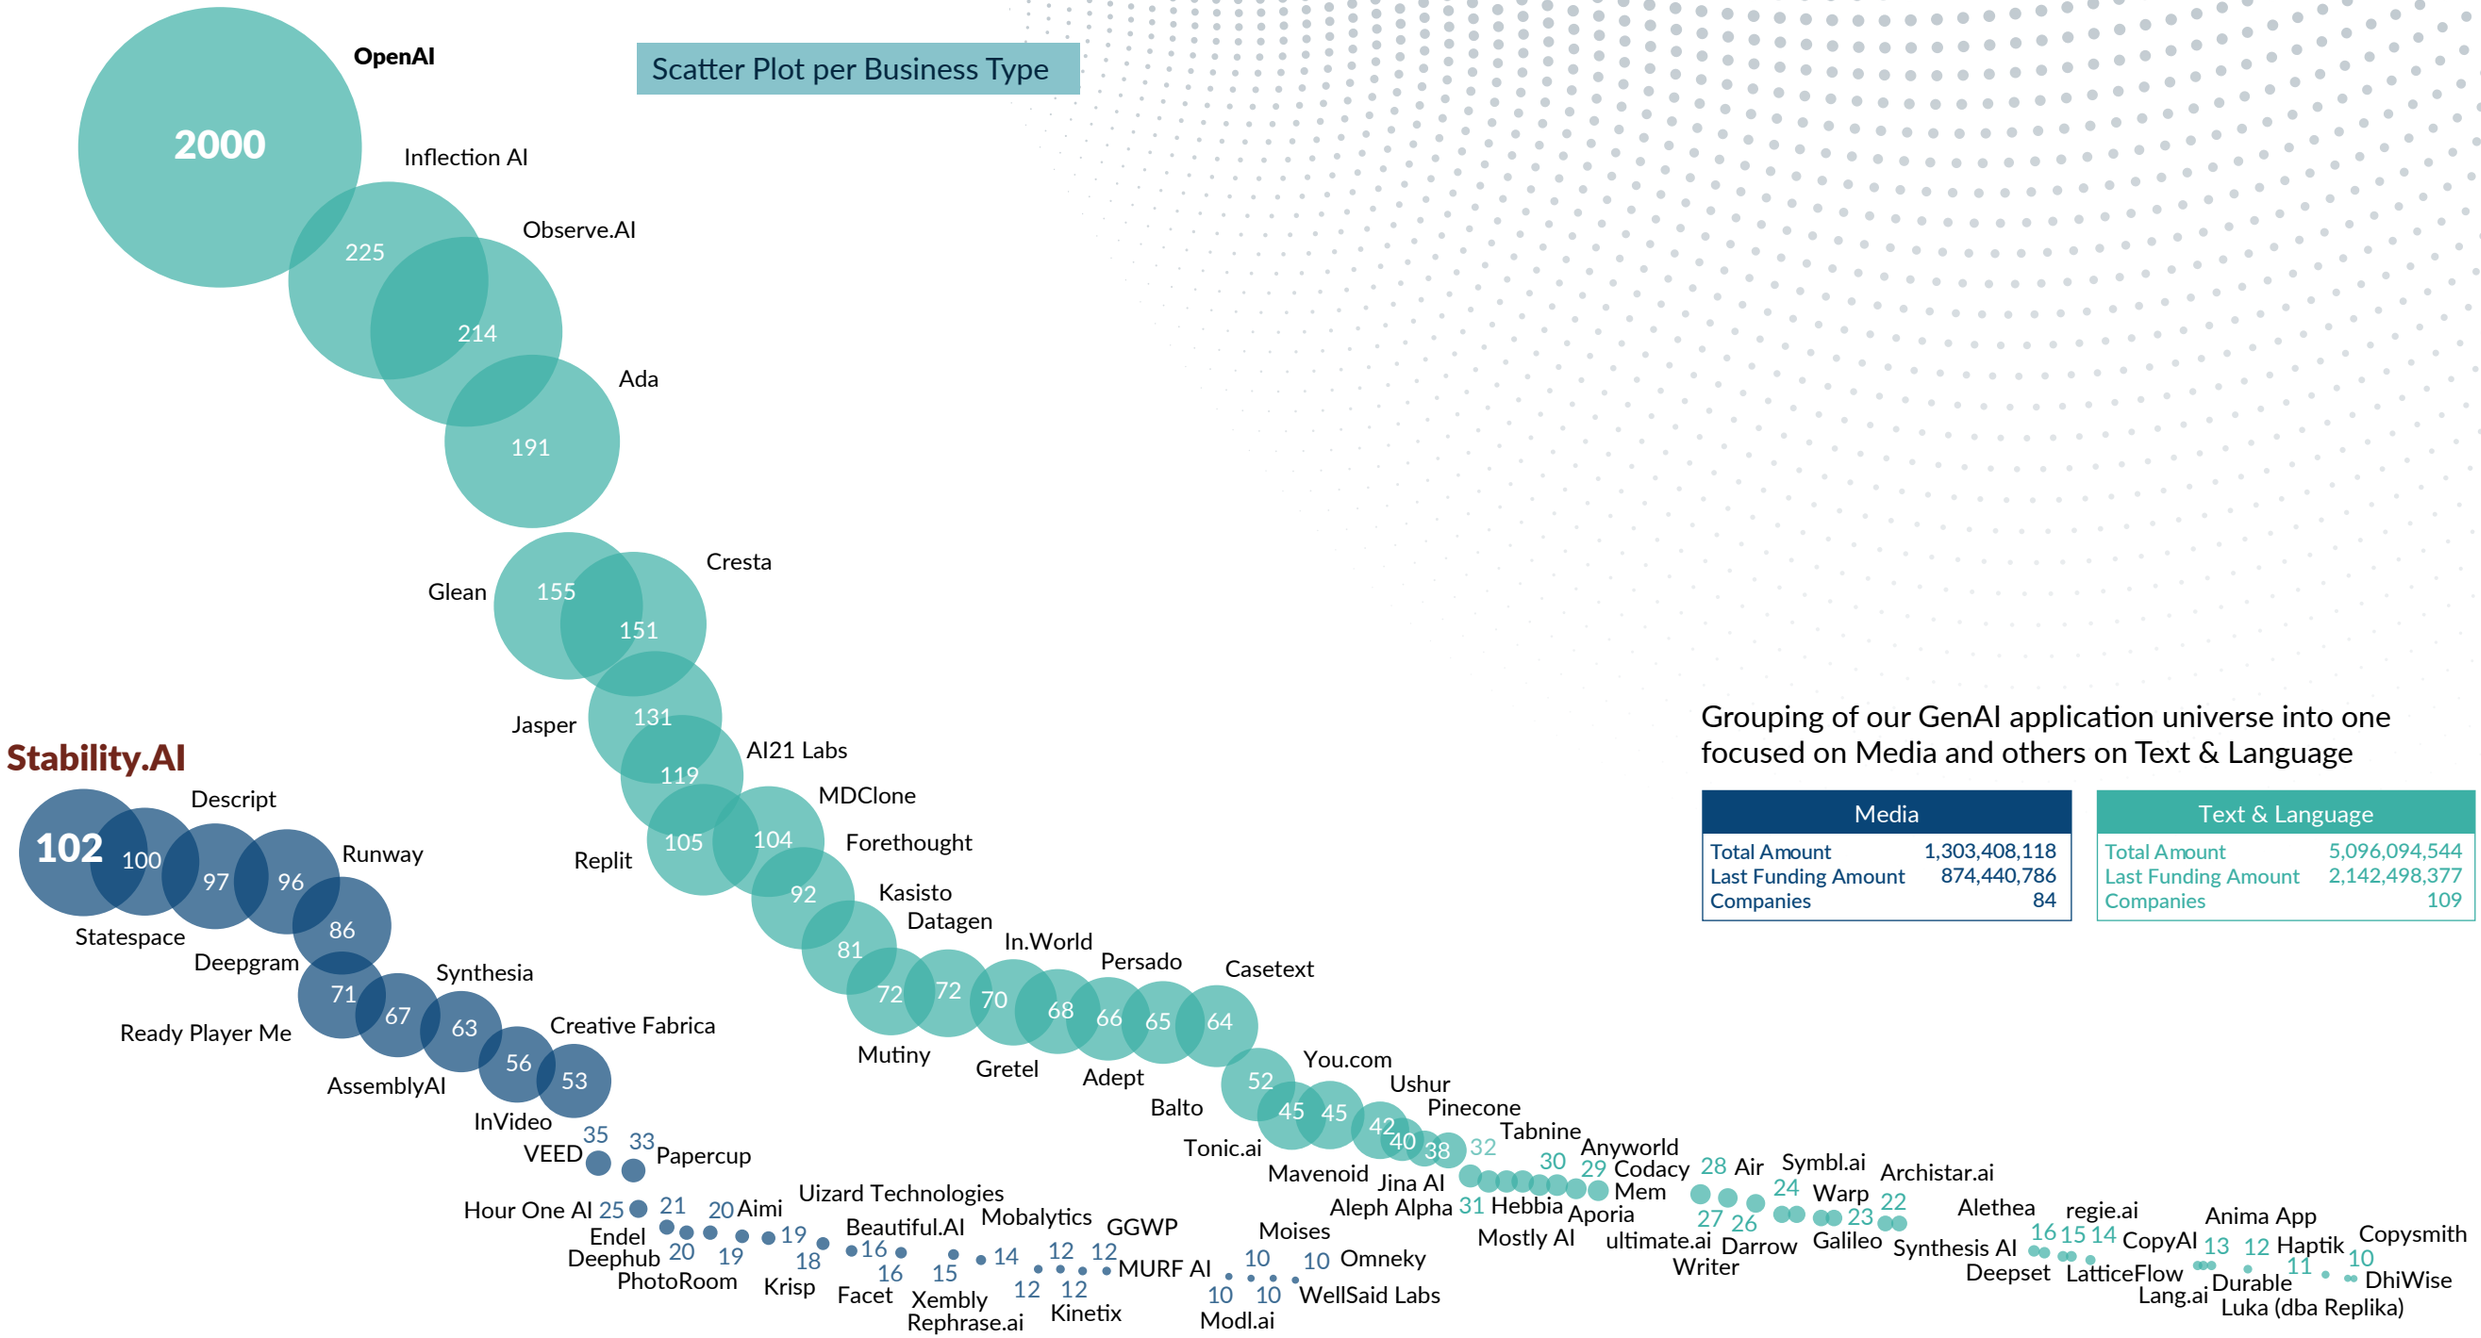

Grouping of our GenAI application universe into one focused on Media and others on Text & Language

Media	
Total Amount	1,303,408,118
Last Funding Amount	874,440,786
Companies	84

Text & Language	
Total Amount	5,096,094,544
Last Funding Amount	2,142,498,377
Companies	109

Applications per Funding Type

With investors focused on the Gen-AI space, we have tried to include most of companies surveyed in this report, grouping them by the different VC funding rounds. The funding amounts raised is shown (in Millions). This singles out as well the 8 unicorns appearing in this short span of time (marked with a *).

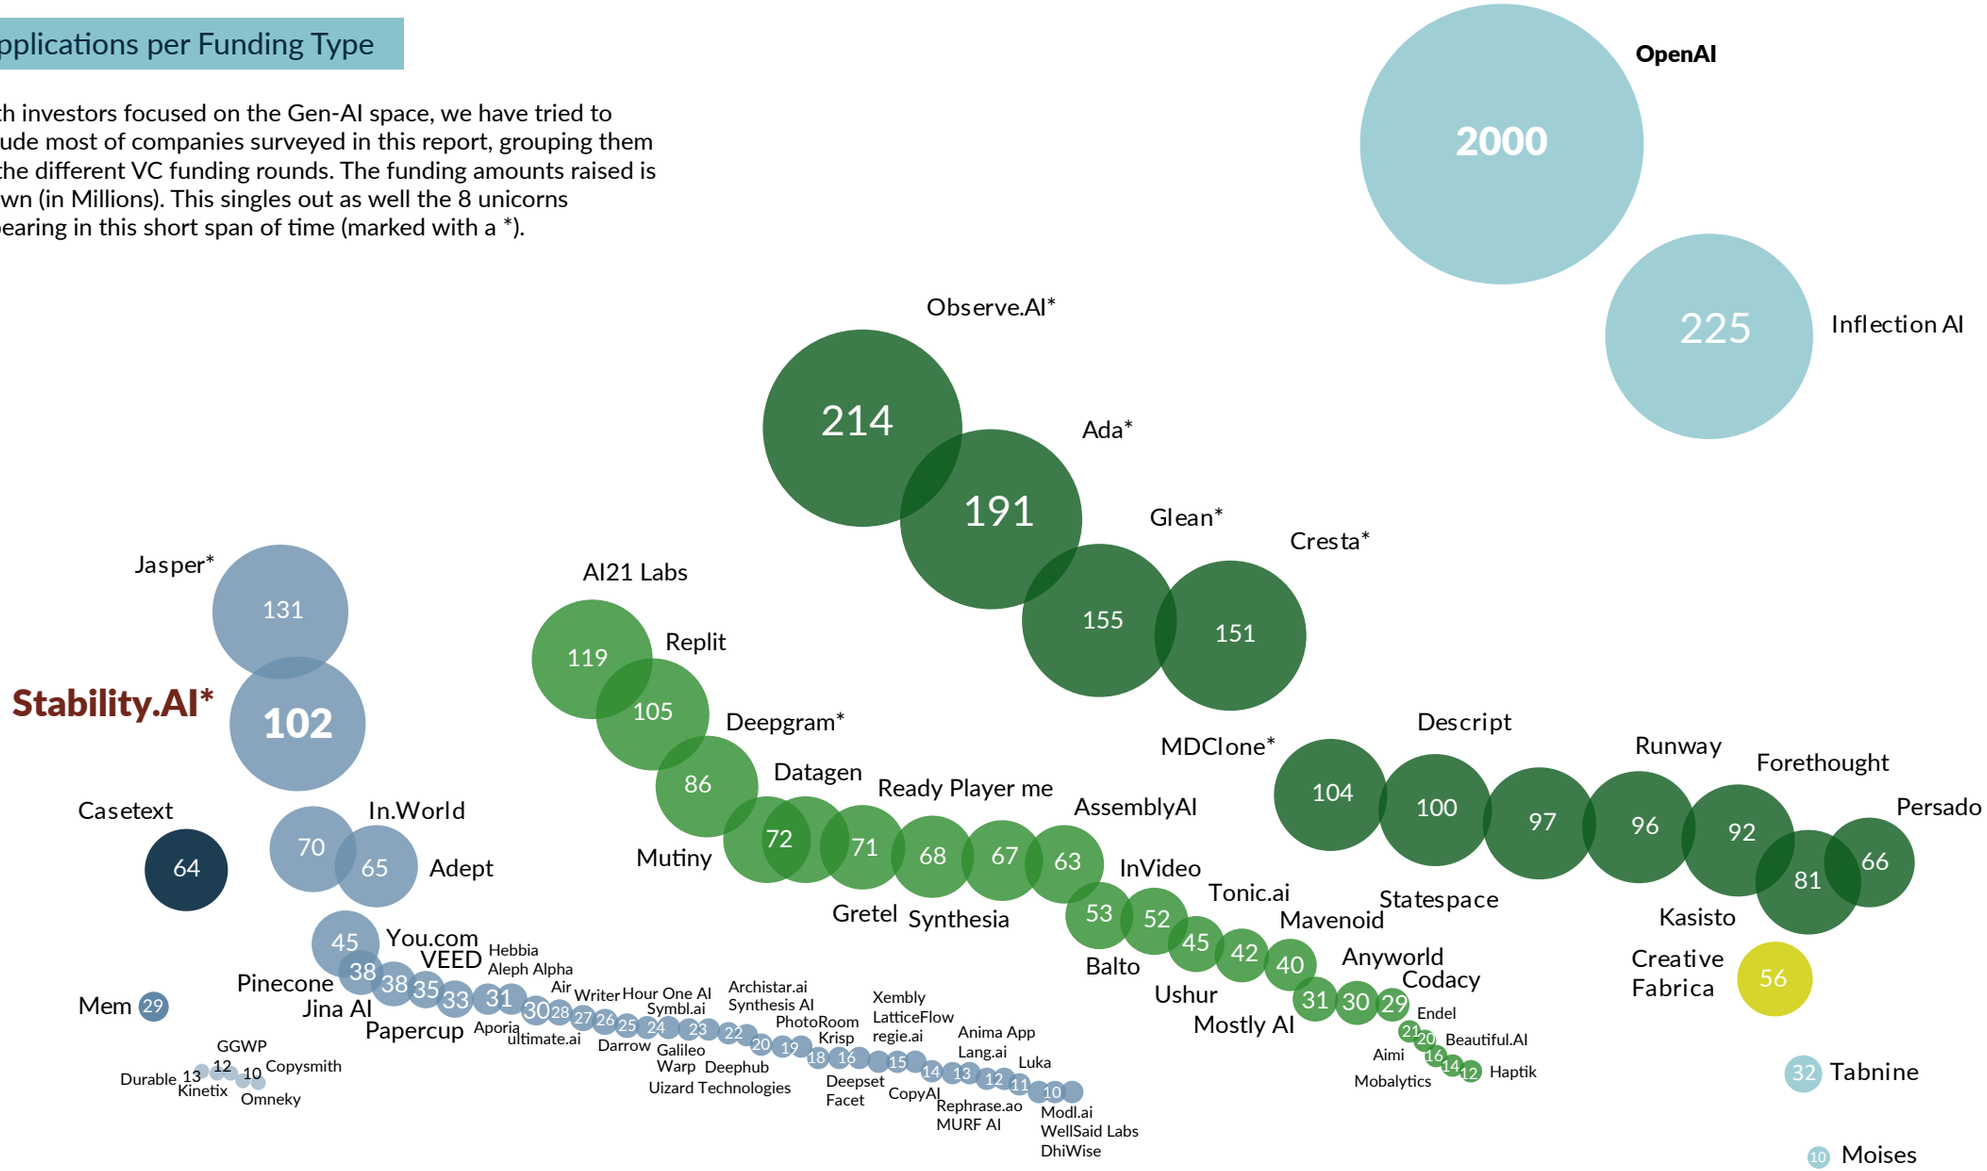

	Early VC	Late VC	Seed	Series A	Series B	Series C	VC Unknowwn	Debt
Total Amount	29,100,000	64,300,000	55,600,000	1,162,696,749	1,128,531,787	1,346,090,000	2,277,975,709	56,300,000
Last Funding Amount	23,500,000	25,000,000	46,250,000	952,266,642	952,266,642	758,500,000	249,325,709	12,200,000
Companies	1	1	5	44	23	12	8	1

Grant	
Total Amount	605,000
Last Funding Amount	604,000
Companies	2

Seed	
Total Amount	279,238,488
Last Funding Amount	207,935,265
Companies	68

Angel	
Total Amount	0
Last Funding Amount	0
Companies	1

In the next table, the massive influx of capital, nearly a 200% increase in inflows to Gen AI businesses across our market map between 2019-2021, is recorded. At US$ 6Mn in 2019 and US$ 65Mn in 2020, it jumped to almost 3 billion in 2021 and 2022. Unsurprisingly, this has been driven by a 500% growth in the specific segments of text and search, as well as 350% in media, 400% in image, 370% in data, and 350% in audio/video. Most of this growth occurred in 2022. Of the US$ 2 bn funding that poured into Open AI, the majority was recorded in 2021.

Year	Total Amount	Last Funding Amount	Companies
2022	2,957,215,800	2,056,971,280	100
2021	3,188,928,000	830,168,753	53
2020	65,016,862	47,777,130	13
2019	6,055,000	6,055,000	4
2018	20,367,000	15,367,000	4
2017	10,920,000	6,500,000	1
2016	78,200,000	41,200,000	2

Gen-AI revenue models: Turning algorithms into cashflows

As discussed in *Section 3*, Stability, well anchored in the foundation layer, still has many applications that are proprietary and could generate funds as per an internal revenue model. This is crucial to support future investments.

B2B, B2C

#	Category	Amount	Last Funding Amount	Companies
0	B2B	5,326,905,736.00	2,264,357,073.00	134.00
1	B2C	586,221,453.00	402,090,709.00	47.00
2	Both	486,375,473.00	350,491,381.00	14.00

OpenAI's US$ 2bn funding is considered part of the B2B category

B2G

Funding by Last Funding Round

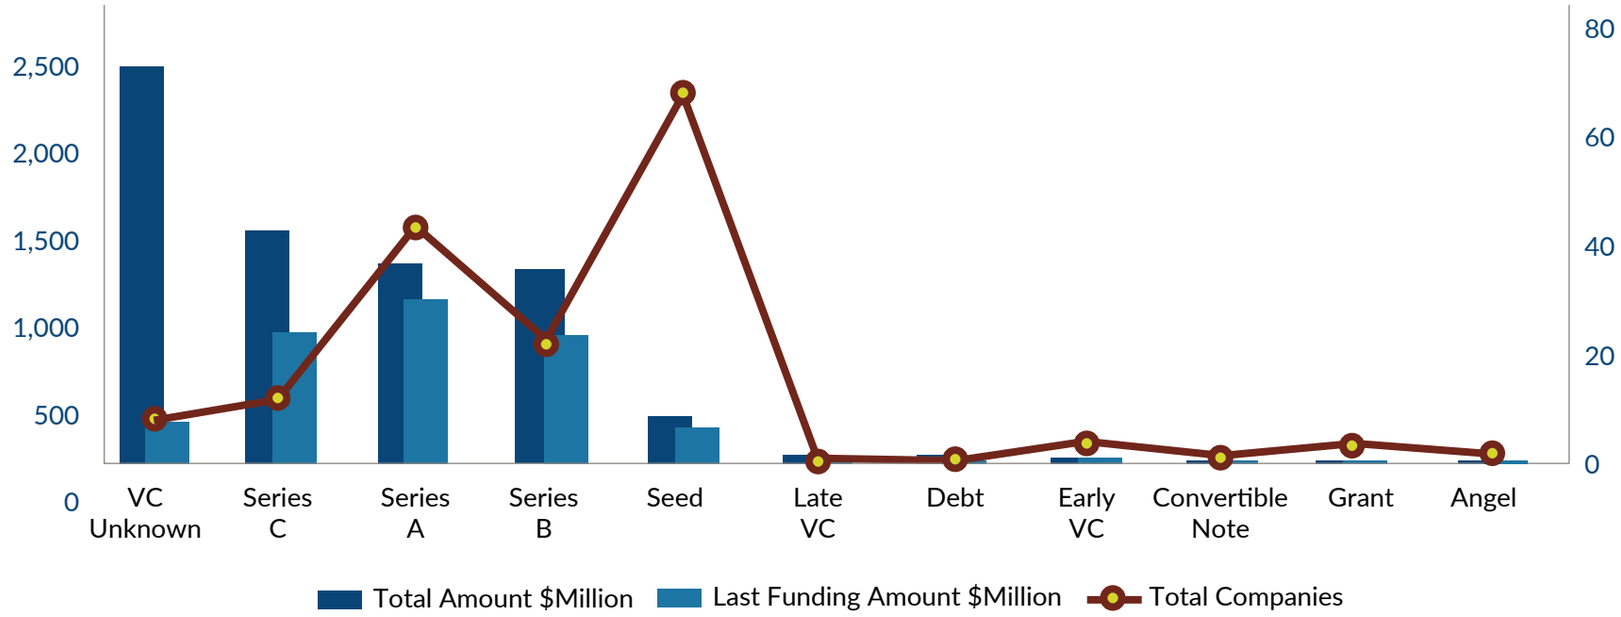

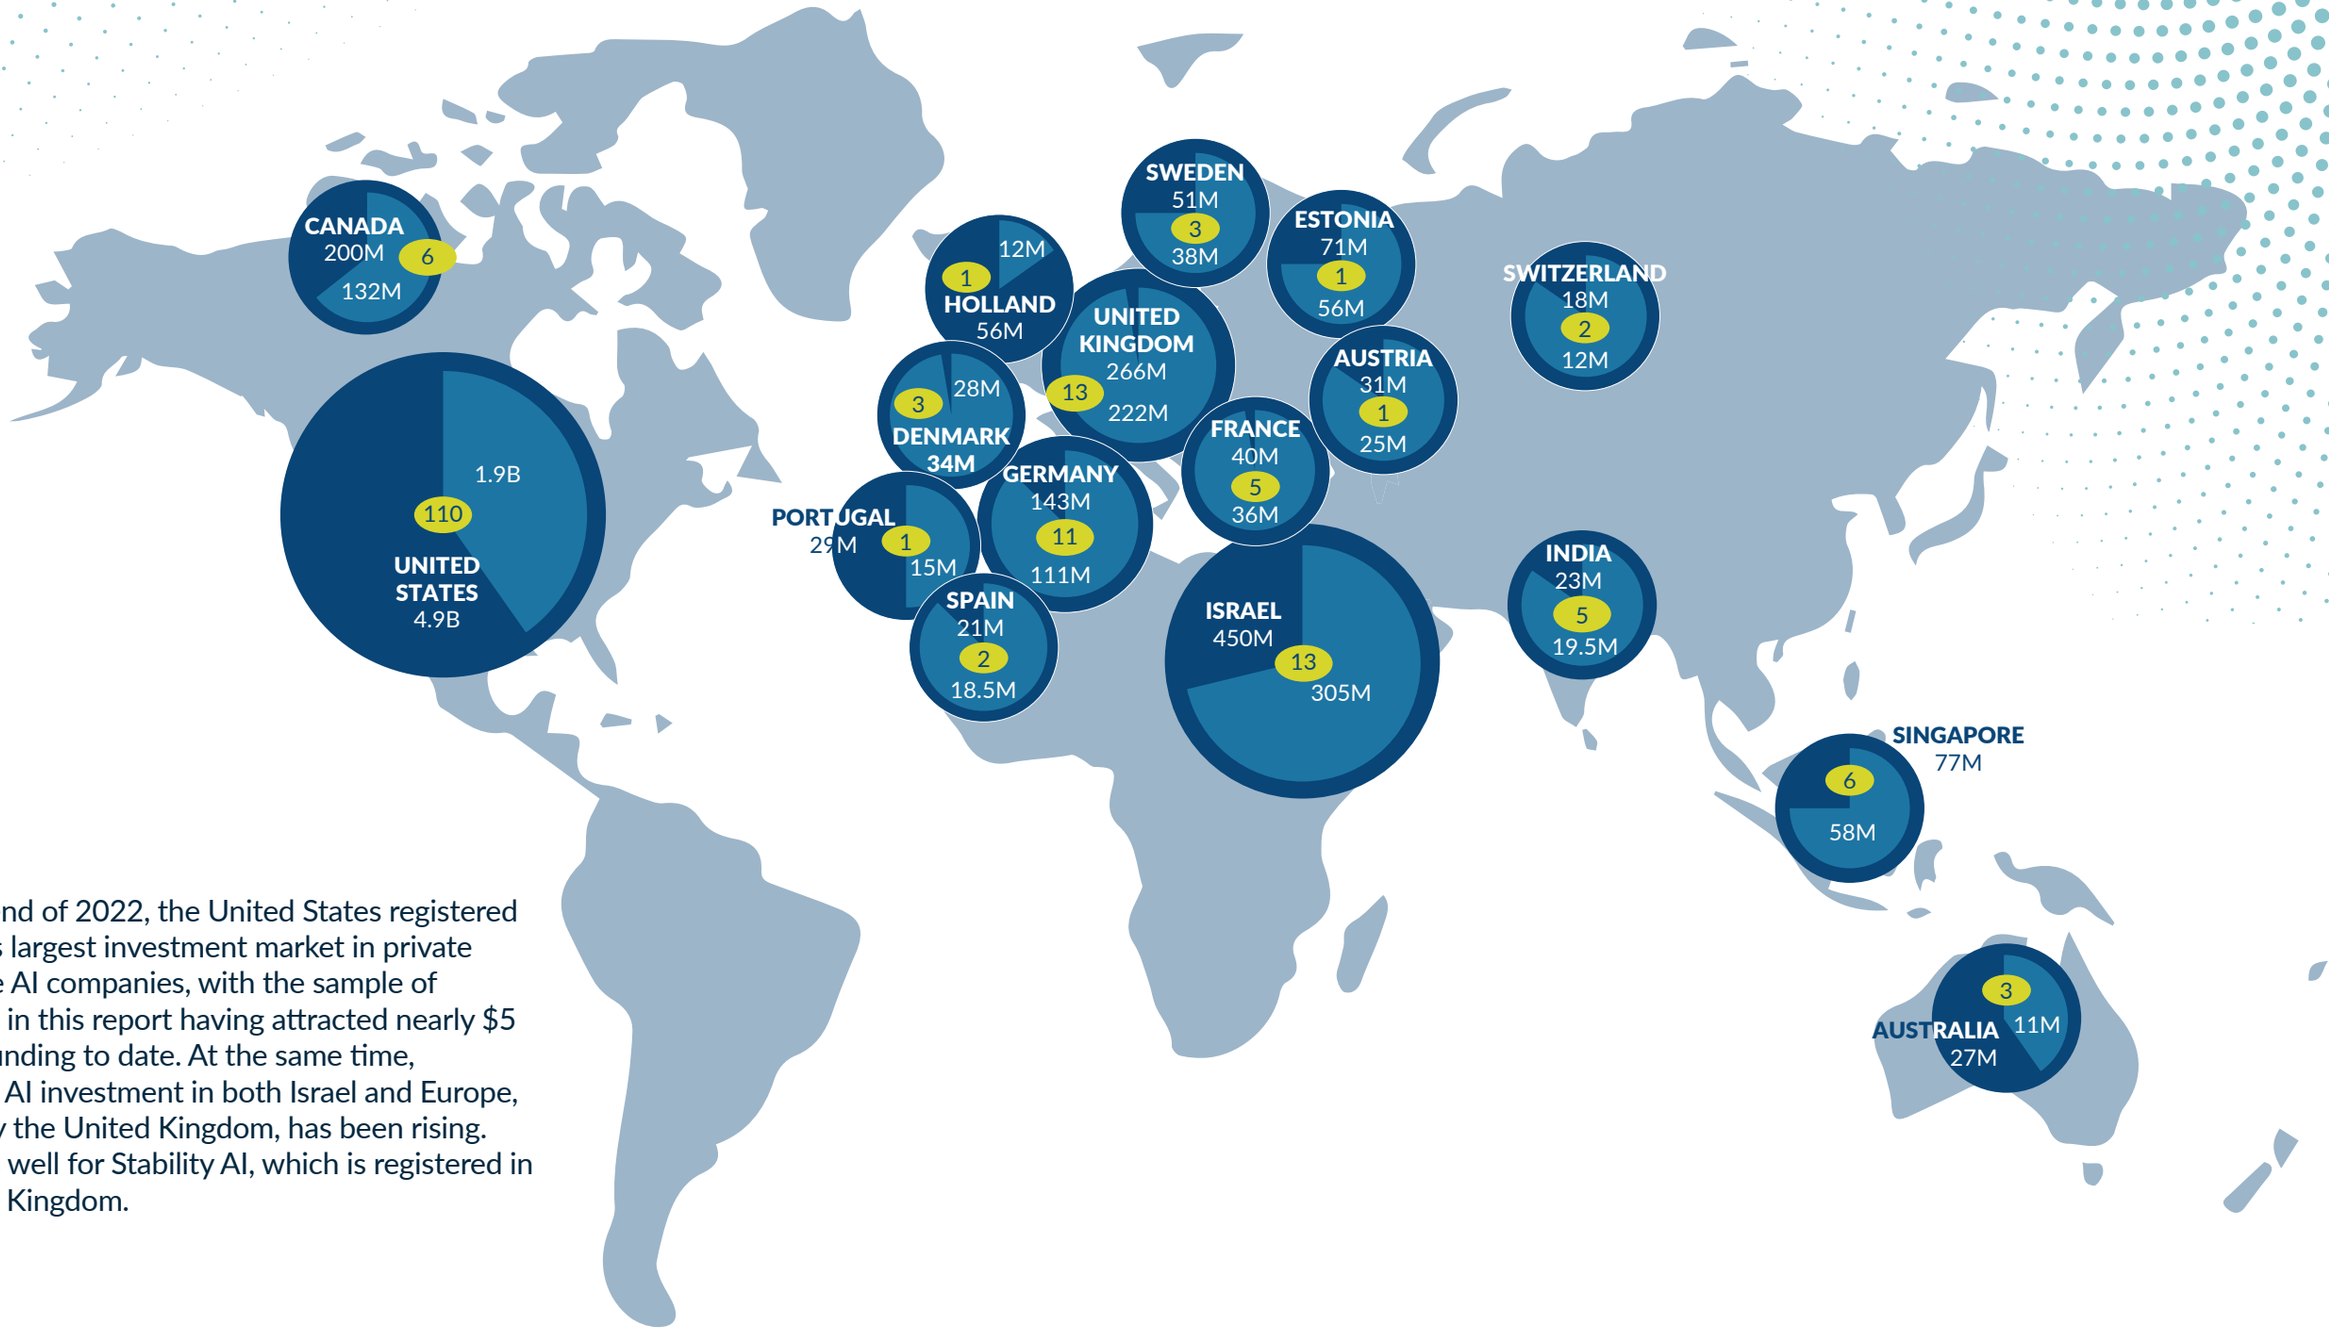

As of the end of 2022, the United States registered the world's largest investment market in private Generative AI companies, with the sample of companies in this report having attracted nearly $5 billion in funding to date. At the same time, generative AI investment in both Israel and Europe, particularly the United Kingdom, has been rising. This bodes well for Stability AI, which is registered in the United Kingdom.

10. FUTURE DIRECTION

The aggregate AI market currently stands at around US$ 300bn (Pitchbook). Gen AI platforms are expected to record significant returns to scale. Building high quality foundational models requires: 1) technical talent (e.g., researchers and data scientists), 2) large datasets (the larger the better), 3) data filtering and data quality, and 4) human and model feedback to improve new models. First-movers gain early advantages across all four of these.

Due to this, it is likely that two major players will take a major portion (65-70%) of the market. Stability AI, with its open-source model and a whole community of international researchers, has a distinct advantage and is likely to generate significant future revenues on the back of this.

We are still scratching the surface of what the future will still hold for this nascent branch of AI.
An eruption of new applications that sits in the application layer as shown in *Section 9*. by the time this report is delivered, probably the funding landscape would have changed, and new ventures having started from angel-level funding. The promise the field is showing means that the application layer, affordable and reachable through APIs will ensure that history repeats itself, but this time, on a considerably larger scale of depth and intelligence.

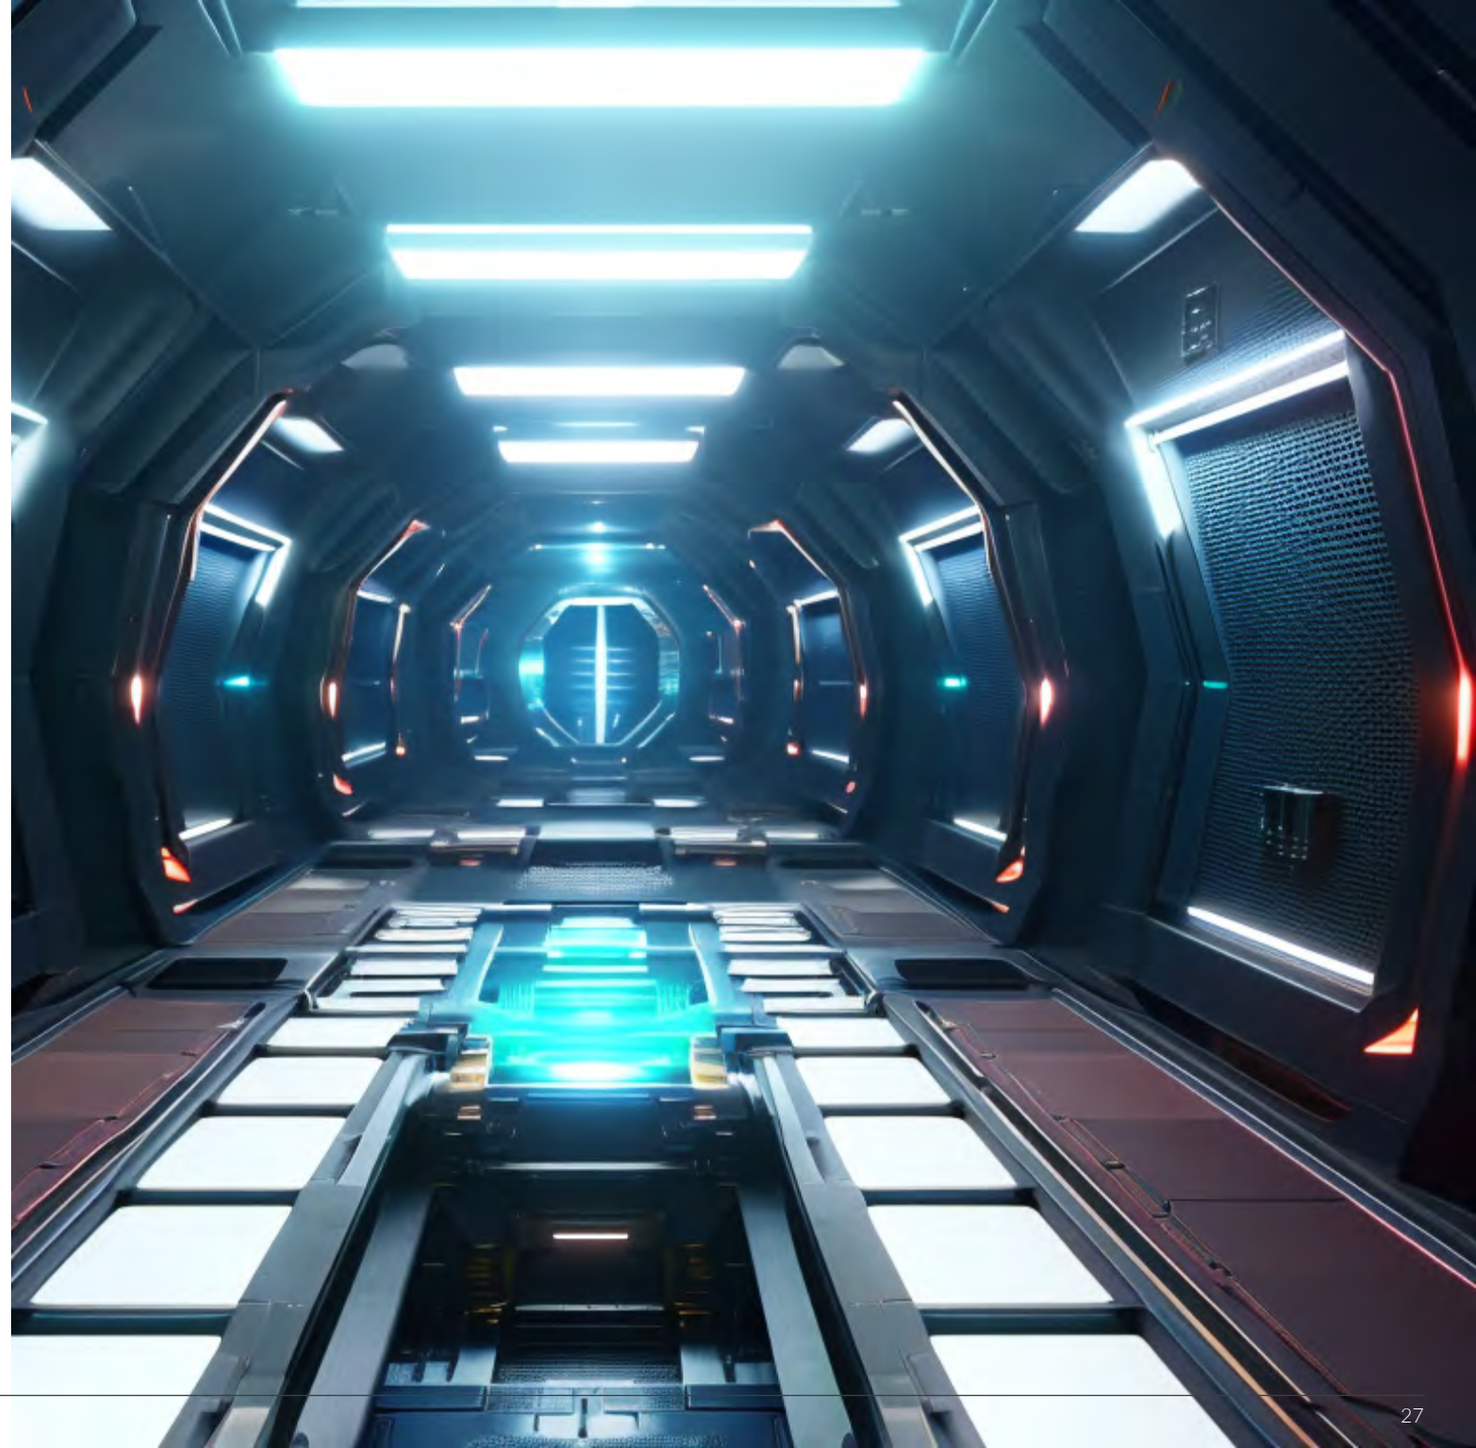

11. ECOSYSTEM

This section provides a brief overview of the major players in the field.

Inflection AI
Year Founded: 2022
Founders: Mustafa Suleyman
Total Funding: $ 225 million
Segment: Code

Inflection AI is an ML company seeking design products that allows people to simplify their ideas in communicating with machines, thus making human interaction with machines easier. The hope is that humans could ultimately be able to communicate with computers in ordinary language. The need for language adjustment, in the form of a command or search term will not be needed.

Observe AI
Year Founded: 2017
Founders: Akash Singh; Sharath Keshava Narayana; Swapnil Jain
Total Funding: $ 214 million
Segment: Chatbots

Observe.AI transforms contact centers using AI to analyze 100% of customer interactions enhancing agent performance and automating workflows.
It uses an intelligence engine to analyze customer interactions across every channel, granting visibility into what is happening inside a business.
The company has developed various NLP tools to track voice and text conversations. It works on augmenting rather than replacing, actual agents, a crucial in an AI approach.

Ada
Year Founded: 2016
Founders: David Hariri; Mike Murchison
Total Funding: $ 191 million
Segment Chatbots

Ada provides an AI-powered customer service automation platform. It works across channels and languages — with minimal human effort. With Ada, customer support service can be increased without increasing headcount through automating the most frequent interactions customers have with businesses. When human assistance is needed, Ada steps in with a hand-off strategy, giving the context needed for efficient resolution of a matter.

Glean
Year Founded: 2019
Founders: Arvind Jain; Piyush Prahladka; Tony Gentilcore; TR Vishwanath
Total Funding: $ 155 million
Segment: Search

Glean is an enterprise search built by a team of Ex-Google search engineers. It acts as a work assistant, searching across all of a company's collective knowledge, applications and content.
It helps personalize search results through multiple verticals (context, language, behavior and relationship with others).
The application also stores answers to frequently asked questions. It organizes related documents and links into a collection for easy sharing.
It also allows companies to restrict information access to certain departments and allows search rankings, and can be used for various teams, including sales, engineering, support etc.

Cresta
Year Founded: 2017
Founders: S. Zayd Enam; Sebastian Thrun; Tim Shi
Total Funding: $ 151 million
Field: Chatbots

Cresta is a cloud-based real time AI contact Centre for customer service. Cresta offers additional features then other such as contact management, live calls, customer support, and other administrative tools.
Cresta helps to turn real-time insights into real-time actions that help contact center performances and drive business results.

Jasper AI
Year Founded: 2021
Founders: Dave Rogenmoser; John Philip Morgan
Total Funding: $ 131 million
Field: Text

Jasper AI is a robotic writer powered by AI technology that can curate content 5x faster than an average human copywriter. It thus acts as a content generator for teams. This content stems from different sources across different media. It is also able to create video scripts for YouTube and TikTok and write blog articles as well as social media posts and more. It can generate educational articles that are keyword-rich and free of plagiarism.
Jasper AI has grossed considerable amount of funding and has got a rising star status in the GenAI community just lagging OpenAI and StabilityAI.

AI21 Labs
Year Founded: 2017
Founders: Ori Goshen; Yoav Shoham
Total Funding: $ 119 million
Field: Text

It is both a research lab, doing cutting-edge work on (NLP), and a commercial business, hoping to quickly push those state-of-the-art developments into products that real businesses can use—and pay for.
AI21 Labs specializes in Natural Language Processing (NLP) and is developing AI systems with a distinguished capacity to understand and generate natural language.
AI21 Studio was then released, along with a "pay-as-you-go" service that allows developers to apply for access to custom models fine-tuned on datasets unique to their requirements.
It released Jurassic-1 Jumbo, the world's largest and most sophisticated language model, to anyone interested in prototyping custom text-based AI applications in 2021.

Replit
Year Founded: 2016
Founders: Amjad Masad; Faris Masad; Haya Odeh
Total Funding: $ 105 million
Field: Code

Replit is a platform for creating and sharing software. It is a popular free online integrated development environment (IDE) that can be used to create projects with very little setup. This editor supports over 50 languages and many programming courses use it.
While traditional IDEs are software programs downloaded to a computer, Repl.it is completely web browser-based and can be accessed from any internet-enabled device such as a phone, tablet, or computer. It is the first-of-its-kind online development environment for students and developers to learn, collaborate and ship applications.

MDClone
Year Founded: 2015
Founders: Ziv Ofek
Total Funding: $ 104 million
Field: Data, Health

MDClone is a digital health company and a leader in synthetic data. It offers a free, secure, self-service platform for building queries and downloading computationally derived ("synthetic") data from the institute's research data core (RDC). It offers data analytics, enabling healthcare collaboration, research, and innovation, both cross-institutionally and globally.
It serves major health systems, customers in the United States, Canada, and Israel, including the VHA ecosystem (Department of Veterans Affairs), The National Institutes of Health, Intermountain Healthcare, Regenstrief Institute, The Ottawa Hospital, and Sheba Medical Center, among others.

Decsript
Year Founded: 2017
Founders: Andrew Mason
Total Funding: $ 117 million
Field: Audio

Descript, a podcast-maker application, a collaborative audio/video editor that works like a doc. It includes transcription, screen recording and publishing. It is an all-in-one tool for audio transcription, podcasting, screen recording, audio, and video editing. And also allows users to organize material, transcribe meeting notes, transcribe zoom video calls and record guests remotely for podcasts.
It also lets users edit their sound files using a word-processing-style document system. From Within Descript, sound files can be published to Twitter, Buzz sprout or Headliner, allowing users to select they type they would like to export.

Statespace
Year Founded: 2017
Founders: Wayne Mackey
Total Funding: $ 102 million
Field: Games

Statespace is a platform that trains gamers. launched in 2017 with a product called Aim Lab, which recreates the physics of popular first-person shooter video games to help players improve their aim.
It goes beyond the mechanics of aim itself to identify and assess various elements of a player's game, from visual acuity across the quadrants of the screen to reaction time.
The company various products like, Academy, which was designed in collaboration with a number of top streamers. Users can get advanced tutorials from these streamers, which include KingGeorge (Rainbox Six Siege), SypherPK (Fortnite), Valkia (Overwatch), Drift0r (CoD) and Launders (CS:GO).
Statespace has also partnered with the Pro-football Hall of Fame to develop the "Cognitive Combine." Just like the NFL, Combine measures general skills and abilities, such as speed, strength, agility, etc.

Runway
Year Founded: 2018
Founders: Alejandro Matamala; Anastasis Germanidis; Cristóbal Valenzuela Barrera
Total Funding: $ 95.5 million
Field: Video

Runway offers dozens of AI-powered creative tools to help users ideate, generate, and edit content. It simplifies the tedious, repetitive, and time-consuming facets of content creation and video editing, allowing users to edit and generate content in seconds using AI tools. It allows video backgrounds to be removed and object movements to be tracked, thus accelerating workflows.
Runway also serves as an asset hub, storing all videos, fonts, and content in a centralized location, thus eliminating the need for further downloads and saving time.

Datagen
Year Founded: 2018
Founders: Gil Elbaz; Ofir Chakon
Total Funding: $ 72 million
Field: Data

Datagen is a leader in synthetic images and video for computer vision use cases. The company offers a platform for generating synthetic data (data manufactured artificially), based on specific criteria and distributions, for virtual reality (VR), augmented reality (AR), computer vision (CV), and artificial intelligence (AI), namely, self-driving cars, robotics and IoT security. This allows users to train and test their machine learning models more efficiently.
Its solutions permit users to generate synthetic data for faces and humans-in-motion in the form of images and videos. This includes in-cabin automotive, security, fitness, digital communication or AR/VR/Metaverse.

Mutiny
Year Founded: 2018
Founders: Jaleh Rezaei; Nikhil Mathew
Total Funding: $ 71.6 million
Field: Text

Mutiny is a no-code AI platform that helps marketers transform their top of funnel demand into revenue, without engineers.
The idea behind Mutiny is to develop an AI system that can analyze a company's online data to offer guidance on underperforming customer segments. Specifically, Mutiny suggests segments for personalization and shows companies how others personalized for that segment. It might also suggest to an enterprise company, for instance, that small startups don't translate well on their website, and then show them how rivals personalized their homepages. Ultimately, Mutiny makes it simple to target and convert more B2B buyers on the business site with no code web personalization.

Ready Player Me
Year Founded: 2014
Founders: Timmu Toke; Rianer Selvet; Kaspar Tiri; Haver Jarveoja
Total Funding: $ 71.1 million
Field: Games

Ready Player Me is a cross-game avatar platform for the metaverse. It allows users to generate a 3D avatar with a selfie and utilize it in 3,000+ compatible apps and games. It also allows users to buy limited non-fungible tokens (NFTs) for their avatars and trade or sell them on a marketplace of your choice.
The Ready Player Me avatar system has two elements. The first is the Avatar Creator which allows users to generate avatars from inside the application, on the partner website or via the Ready Player Me account. The second is the Ready Player Me SDK, which is to be integrated into the application to find user-created avatars to be used in the user's game or application. This benefits users by eliminating the need to design and implement their own avatar system.

Forethought
Year Founded: 2017
Founders: Colm Doyle; Deon Nicholas; Sami Ghoche
Total Funding: $ 92 million
Field: Chatbots

Forethought serves as a GenAI platform for customer support automation. It improves efficiency by automating the entire support ticket lifecycle, lowering support costs, and providing top-tier service in customer interactions.
Forethought's generative AI uses Large Language Models (LLMs) to maximize efficiency and ROI for support teams.
The platform allows for the streamlining of the entire support ticket lifestyle. It resolves simple, repetitive ticks automatically, reducing customer wait times, improving self-service and reducing ticket volume and backlog for agents. It also accelerates first response time by using historical knowledge to predict ticket characteristics and detect sentiment and intent, routing it to the right team.

Deepgram
Year Founded: 2015
Founders: Noah Shutty; Scott Stephenson
Total Funding: $ 85.9 million
Field: Audio

Deepgram offers an AI speech platform that automatically transcribes real-time or pre-recorded audio and video into text with AI, in addition to formatting features for better readability. Deepgram builds tailored speech models to optimize their transcription accuracy.
Deepgram concentrates on building custom voice-recognition solutions for customers such as Spotify, Auth0 etc.. . The company's data scientists' source, build, identify and evaluate speech data to generate speech-recognition models that can understand brands and jargon, capture an assortment of languages and accents, and adjust to challenging audio environments. For example, for NASA, Deepgram constructed a model to transcribe communications between Mission Control and the International Space Station.

Kasisto
Year Founded: 2013
Founders: Dror Oren; Ruth Brown; Sasha Caskey; Zor Gorelov
Total Funding: $ 81.5 million
Field: Chatbots

Kasisto gives every financial institution the power to deploy best-in-class digital assistants that allow all channels to come alive. KAI lets clients adopt a set of AI-powered digital assistants trained to imitate the best bankers. They host human-like conversations, usually matching and often surpassing live agent performance. Deep conversational AI and financial expertise allow KAI to examine account activity, produce insights, and consider past interactions to weave in context from previous conversations. KAI then brings it all together to provide intuitive recommendations from a virtual best banker.

Cohere
Year Founded: 2019
Founders: Aidan Gomez; Ivan Zhang; Nick Frosst
Total Funding: $ 164.9 million
Field: Chatbots

Cohere offers access to affordable, easy-to-deploy advanced large language models and NLP tools through one easy-to-use API. The company's platform provides computers the ability to read and write, allowing customers to be better understood and compelling copies to be written for target audiences.
With Cohere's technology, businesses can help customers in seconds by taking instant control of their screen without any downloads or setup on the customer's end. Fundamentally, Cohere is designed to make it easier for teams, whether small founders or large enterprises, to help address their customers' issues. Instead of requiring a software download or complicated install process, customers can just click a button on a website to allow remote control of their screen. It is still though way behind, Stability or Open AI , in the foundation Layer.

Harmonai
Year Founded:
Founders:
Total Funding:
Field: Audio

Harmonai generates audio tools, that are open-source in nature, in order to help users in producing music. It allows users to produce infinite music libraries customized to their taste. It released Dance Diffusion, an AI tool that can create music clips by training on existing songs. It uses a process called diffusion to create new sounds and ultimately aims to help researchers and developers working in the field of music production.

EleutherAI
Year Founded: 2020
Founders: Connor Leahy
Total Funding:
Field: Chatbot

EleutherAI is a group of researchers, engineers and developers supported by Stability AI. They have developed GPT-neo and GPT-J, a 6 billion parameter model available publicly. It can be used to produce code.

Laion
Year Founded: 2021
Founders: Christoph Schuhmann; Jenia Jitsev; Richard Vencu; Robert Kaczmarczyk

Laion, (or Large-scale Artificial Intelligence Open Network,) is Not-for-Profit, organization, an open community, that maintains datasets, tools and models meant to serve the machine learning and AI community. Its primarily the custodian of LAION-5B, a 5.85 Billion multilingual image-text pairs, used to train Generative AI models.(Though 2Billion image-text pairs are in English). It's the most widely used dataset by open and closed source entities including Google and Meta. The Lain-5B precursor is the Laion -400 containing 400M image-text pairs. Laion is funded by donations and research grants.

Prepared for STABILITY.AI by
Avalanche Consultancy Services LTD

Building 280, Wstation 182A-194A, Taweelah, Abu Dhabi United Arab Emirates

Contact :
Camille Beyrouthy
C.Beyrouthy@avalanche-me.com
+971 (56) 330 4790

Sana Khan
s.khan@avalanche-me.com
+971 (56)5589090