AXA Venture Partners

Artificial Intelligence in 2023:

An Overview of the Current State of AI, Investment Trends and Recommendations for AVP

We invest in great entrepreneurs
We support outstanding companies

April 2023

Table of Contents

The (AI)ntroduction

A Welcome Message

What do these people have in common?









It’s AVP’s New York AI-Generated Team!





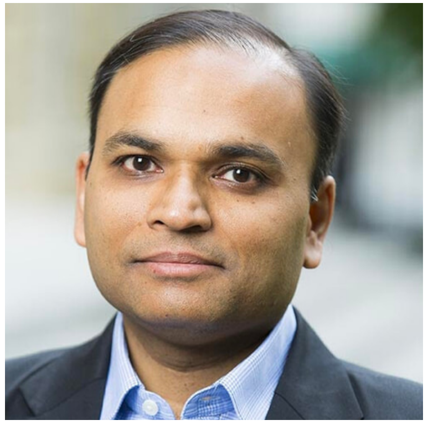



Executive Summary

Executive Summary

- Over the past five years, artificial intelligence (AI) has seen a step-function acceleration in technological progress due to three key factors
  - Advancements in ML architecture (e.g., Transformers)
  - Compute power has increased exponentially
  - Increased creation of and access to data (for model training)
- Wall Street Research analysts are bullish on the future of AI – reports suggest that AI is currently a >$30 billion market with 100% expected annual growth
  - Widespread AI adoption could eventually drive a 7%, or almost $7 trillion, increase in annual global GDP over a 10-year period
- Venture capital investment in AI has boomed since 2019
  - Between 2019 and 2022, there have been 14,000+ unique VC funding rounds globally for AI companies (~5,500 of these rounds occurred in either the United States, Canada, Europe or Israel)[a]
  - Over the same period, both the funding amount and valuation of rounds have steadily increased in conjunction with the broader tech bull market landscape
- For AVP, there is ample opportunity to invest in AI within the application layer (i.e., SaaS applications that use AI in their core product offering)
  - However, AVP should also opportunistically consider companies in the foundation model and AI operations layers as companies in these segments will provide the core infrastructure and tooling on which the application layer is built
- AI is still very much in the early innings and new advancements are happening at a faster pace – we view AI as the next generational technology shift and should seek to invest in the category in the near term

(a) Pitchbook.

Artificial Intelligence – Why Now?

Why is AI Interesting Now?

1 Advancements in machine learning architecture (e.g., Transformers)

- In 2017, Google introduced the Transformer model, a novel neural network architecture that requires less computation to train and is better suited for modern machine learning hardware
  - The key innovation in the Transformer model revolved around the "self-attention mechanism", which allows Transformer models to understand relationships between words at greater distances
  - Moreover, Transformers can be parallelized across both training and inference, significantly speeding up training times
- The Transformer model improved upon existing neural network architectures (CNNs and RNNs) by addressing limitations that each model faces
  - Convolutional Neural Networks (CNNs) focus on relative location and proximity (vs. an entire sequence of data)
  - Recurrent Neural Networks (RNNs), including Long Short-Term Memory (LSTMs), have memory limitations and tend to read in one direction
- The advent of transformers has led to the proliferation of large language models (LLMs), including OpenAI's GPT-4 and Google's LaMDA

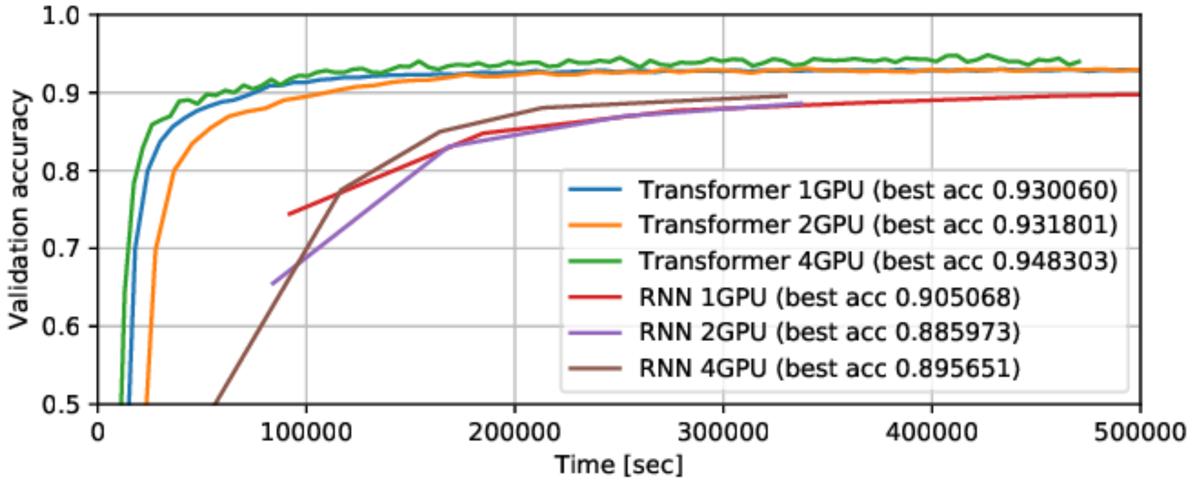

Source: Google Research and other publicly available information.

Why is AI Interesting Now? (cont'd)

2 Compute power has increased exponentially

- Since 2012, the amount of compute used in the largest AI training runs has been increasing exponentially with a 3.4-month doubling time (by comparison, Moore's Law had a 2-year doubling period)
  - This metric has grown by more than 300,000x (a 2-year doubling period would yield only a 7x increase)
  - Within many current domains, more compute seems to lead predictably to better performance, and is often complementary to algorithmic advances
- The increase in compute has been largely driven by two factors
  - Researchers repeatedly finding ways to use more chips in parallel and being willing to pay the economic cost of doing so
  - Custom hardware that allows more operations to be performed per second for a given price (GPUs and TPUs, or Tensor Processing Units)
- As compute power has increased, so too have the parameters of large language models
  - For example, GPT-3 was trained on 175 billion parameters (vs. 1.5 billion for GPT-2)

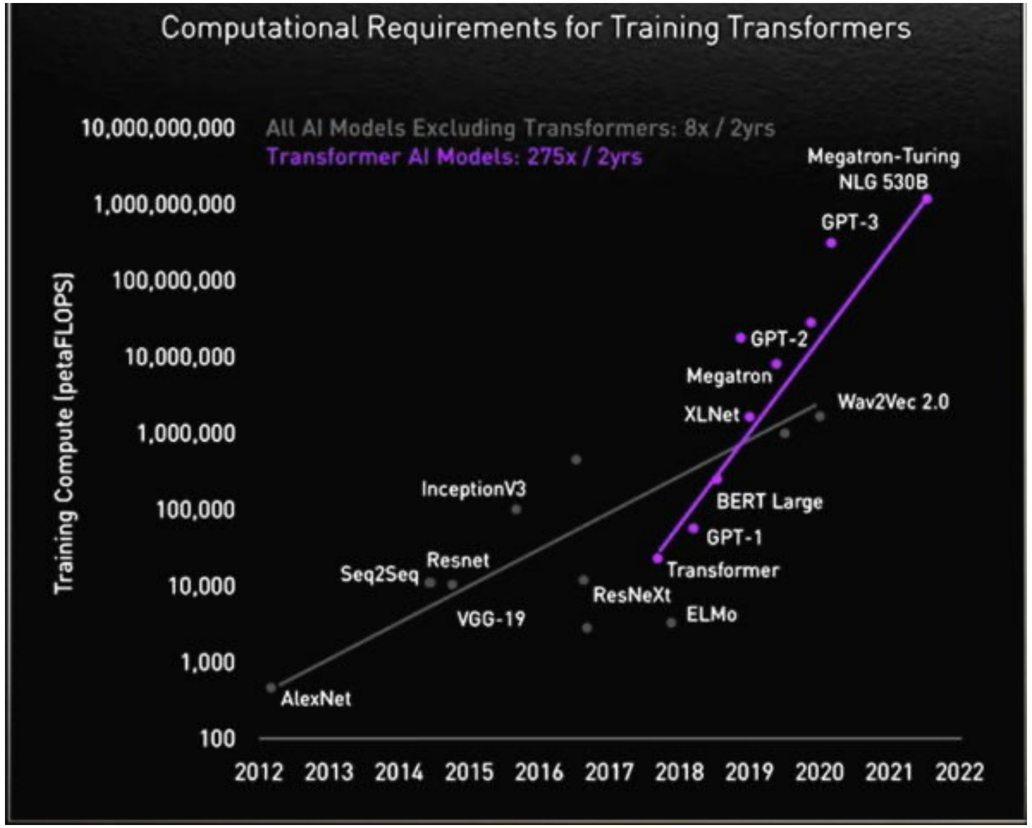

Source: OpenAI and other publicly available information.

Why is AI Interesting Now? (cont'd)

3 Increased creation of and access to data (for model training)

- Since the advent of the Transformer model in 2017, several companies including OpenAI, Google, Microsoft, Meta and Nvidia have invested resources towards building LLMs
  - The cost of training these models is incredibly expensive and requires extensive compute power – for the average company, access to both capital and compute is limited
    - For context, the cost of a single Nvidia GPU server can cost upwards of $200,000 with each server consuming up to 6.5 kilowatts
- Developers of LLMs have taken the route of both open source software and private APIs to democratize access to models and data for companies looking to build on top of their infrastructure
  - Open Source: Hugging Face offers over 120k models, 20k datasets and 50k demos which engineers can easily use to begin their software development
  - Private: Open AI offers its LLM capabilities via an API, which can cost anywhere from $0.0004 to $0.02 per API call

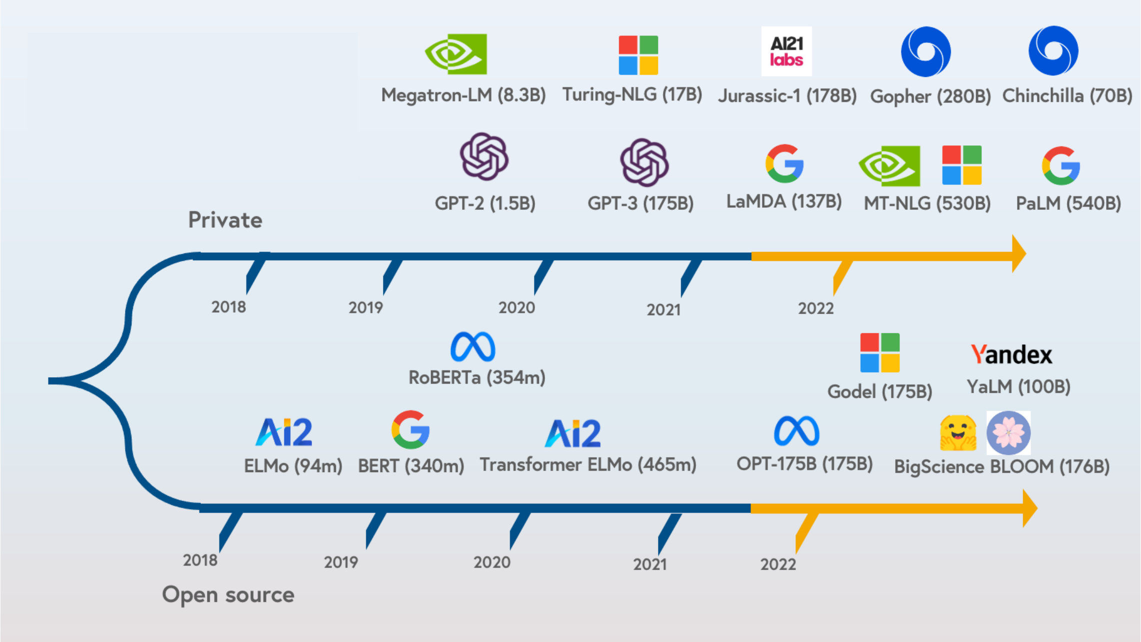

Source: Hugging Face and other publicly available information.

What are market leaders saying?

The development of AI is as ***fundamental as the creation of the microprocessor, the personal computer, the Internet, and the mobile phone***. It will change the way people work, learn, travel, get health care, and communicate with each other. Entire industries will reorient around it. Businesses will distinguish themselves by how well they use it. There will be an explosion of companies working on new uses of AI as well as ways to improve the technology itself. But one big open question is whether we'll need many of these specialized AIs for different uses or whether it will be possible to develop an artificial general intelligence that can learn any task.

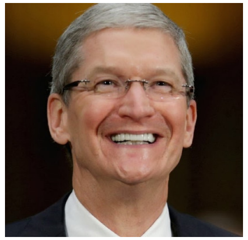

With regard to AI, it is a major focus of ours. It's incredible in terms of how it can enrich customers' lives. And you can look no further than some of the things that we announced in the fall with crash detection and fall detection or back a ways with ECG. I mean these things have literally save people's lives. And so we see an enormous potential in this space to affect virtually everything we do. It's obviously a horizontal technology, not a vertical. And ***so it will affect every product in every service that we have***.

The age of AI is upon us and Microsoft is powering it. We are witnessing nonlinear improvements in capability of foundation models, which we are making available as platforms. ***We fully expect us to sort of incorporate AI in every layer of the stack***, whether it's in productivity, whether it's in our consumer services. For the last 3.5, 4 years, we've been working very, very hard to build both the training supercomputers and now, of course, the inference infrastructure because once you use AI inside of your applications, it goes from just being training-heavy to inference.

First, the AI opportunity ahead. AI ***is the most profound technology we are working on today***. Our talented researchers, infrastructure and technology make us extremely well positioned as AI reaches an inflection point. Already, breakthroughs in everything from natural language understanding to generative AI are fueling our ability to deliver results that drive meaningful performance for advertisers and are useful to users. In fact, our Transformer research project and our field-defining paper in 2017 as well as our path-breaking work in diffusion models are now the basis of many of the generative AI applications you're starting to see today.

Source: Public filings and other publicly available information.

The Market for AI

Wall Street is Bullish on the Future of AI

<u>**Oppenheimer Equity Research (March 9, 2023)**</u>

- We estimate that AI will drive ***half the incremental GDP*** over the next decade, ***representing 20% of global GDP in 2032***
- The current ***AI market is only in the $30bn range, set to grow 100% per year*** for a while, but the impact on the economy will be ***10x*** this
- ChatGPT is an order of ***magnitude better than anything that has come before it*** and driving the reengineering of cloud infrastructure
- Over the next decade, most of the value in AI will accrue to platforms that have the ***capacity to support the training compute demand***

<u>**Goldman Sachs Equity Research (March 26, 2023)**</u>

- Widespread AI adoption could eventually drive a ***7%, or almost $7 trillion, increase in annual global GDP*** over a 10-year period
- Global private investment in AI represented $94 billion in 2021; if investment in AI continues to increase at the more modest pace that software investment grew during the 1990s, ***US investment in AI alone could approach 1% of US GDP by 2030***
- We estimate that generative AI could ***boost aggregate labor productivity by 1.5% in the US***; however, more powerful AI could drive productivity gains to 3%

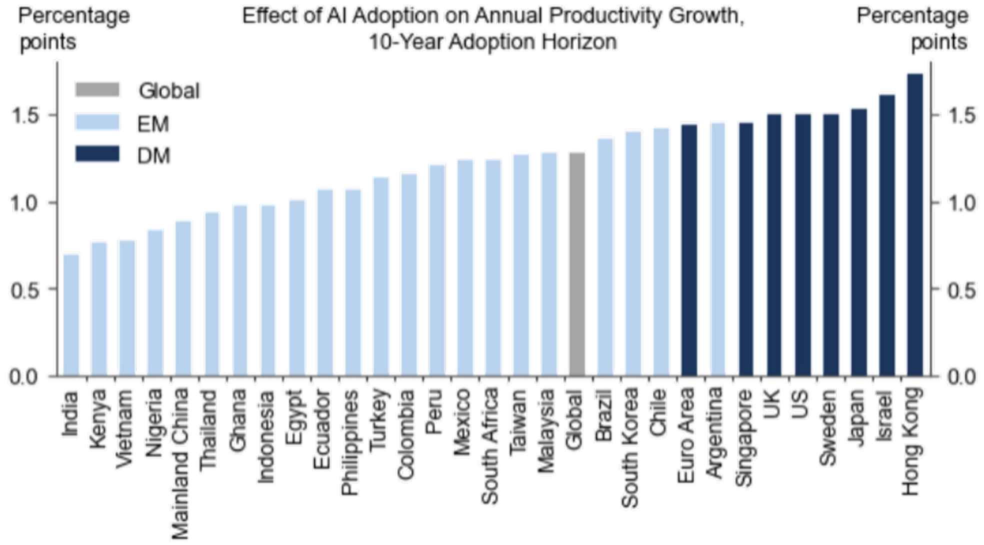

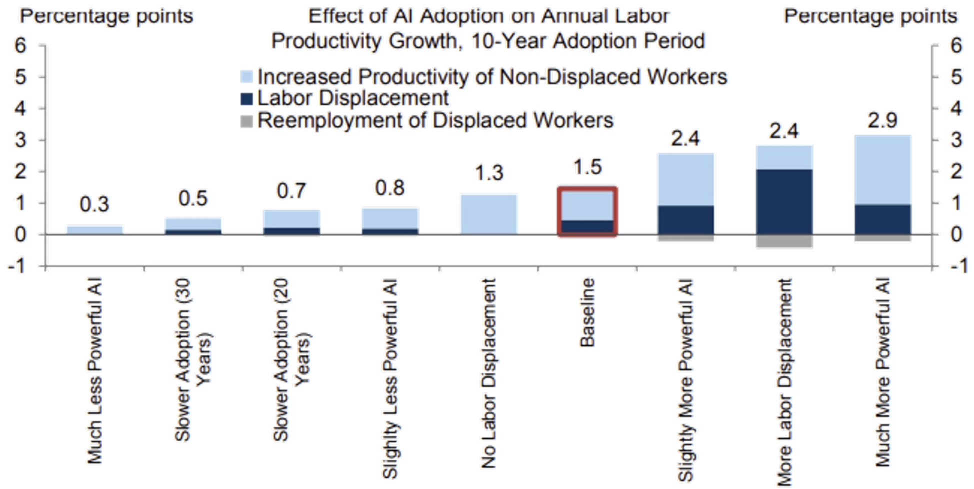

Source: Goldman Sachs Global Investment Research, Oppenheimer Equity Research.

AI as a Complement, Not a Substitute

- In the US and Europe, roughly 67% of current jobs are exposed to some degree of AI automation and generative AI could substitute up to 25% of current work
- Although the impact of AI on the labor market is likely to be significant, most jobs and industries are only partially exposed to automation and are thus more likely to be complemented rather than substituted by AI
  - — 7% of current US employment will be substituted by AI, 63% complemented and 30% unaffected
  - — Replacement in legal and administrative fields, little effect in manual and outdoor jobs and productivity-enhancement everywhere else
- The combination of labor cost savings, new job creation and higher productivity for non-displaced workers raises the possibility of a productivity boom that raises economic growth substantially

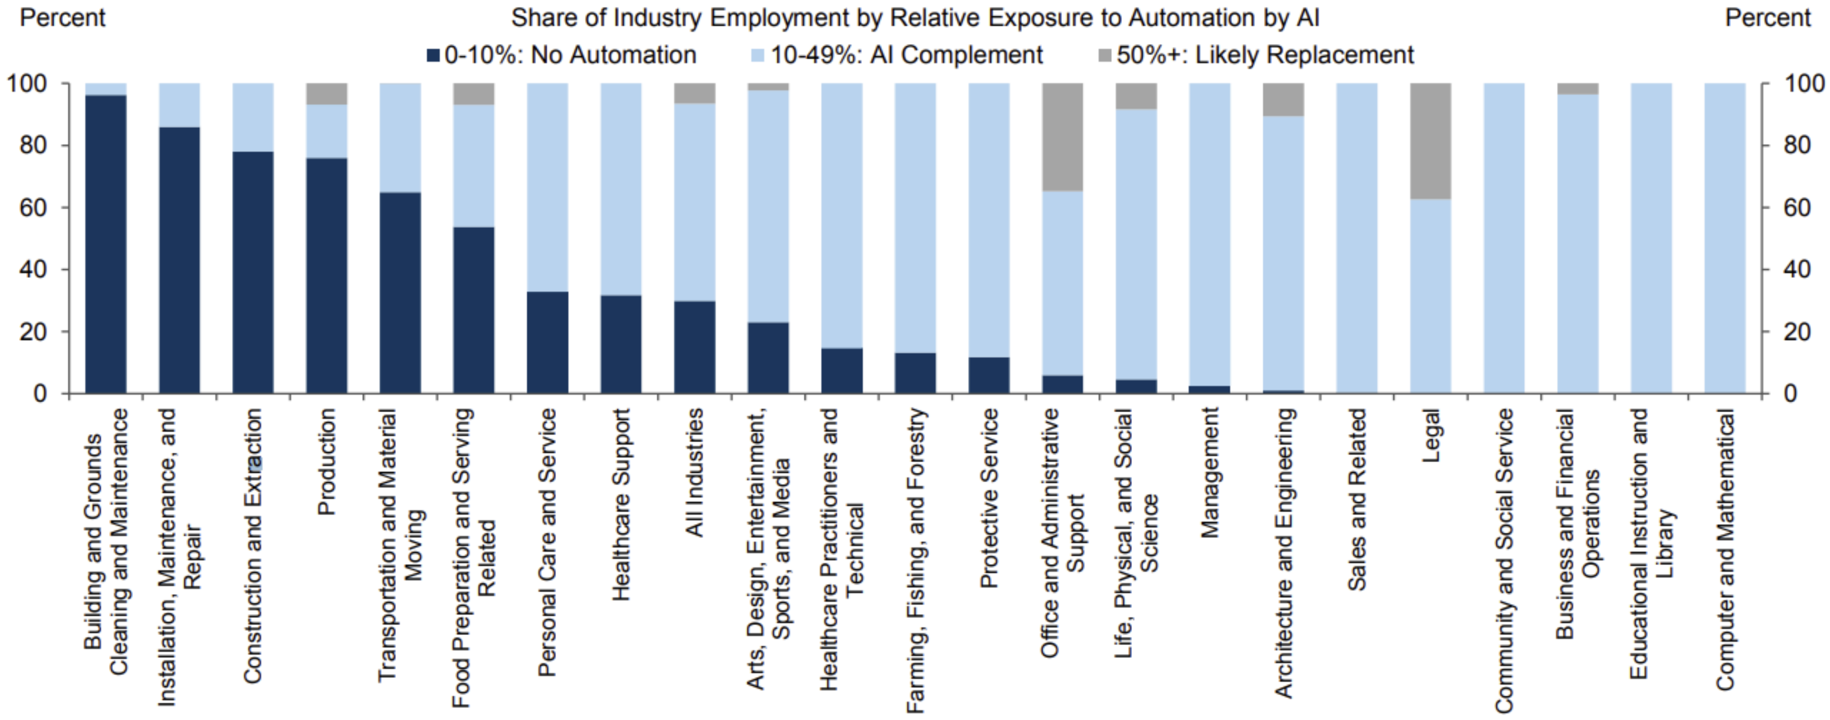

Source: Goldman Sachs Global Investment Research.

AI Investment Trends & Metrics

Trends and Stats in Artificial Intelligence Investing (Jan. 2019 – Feb. 2023)

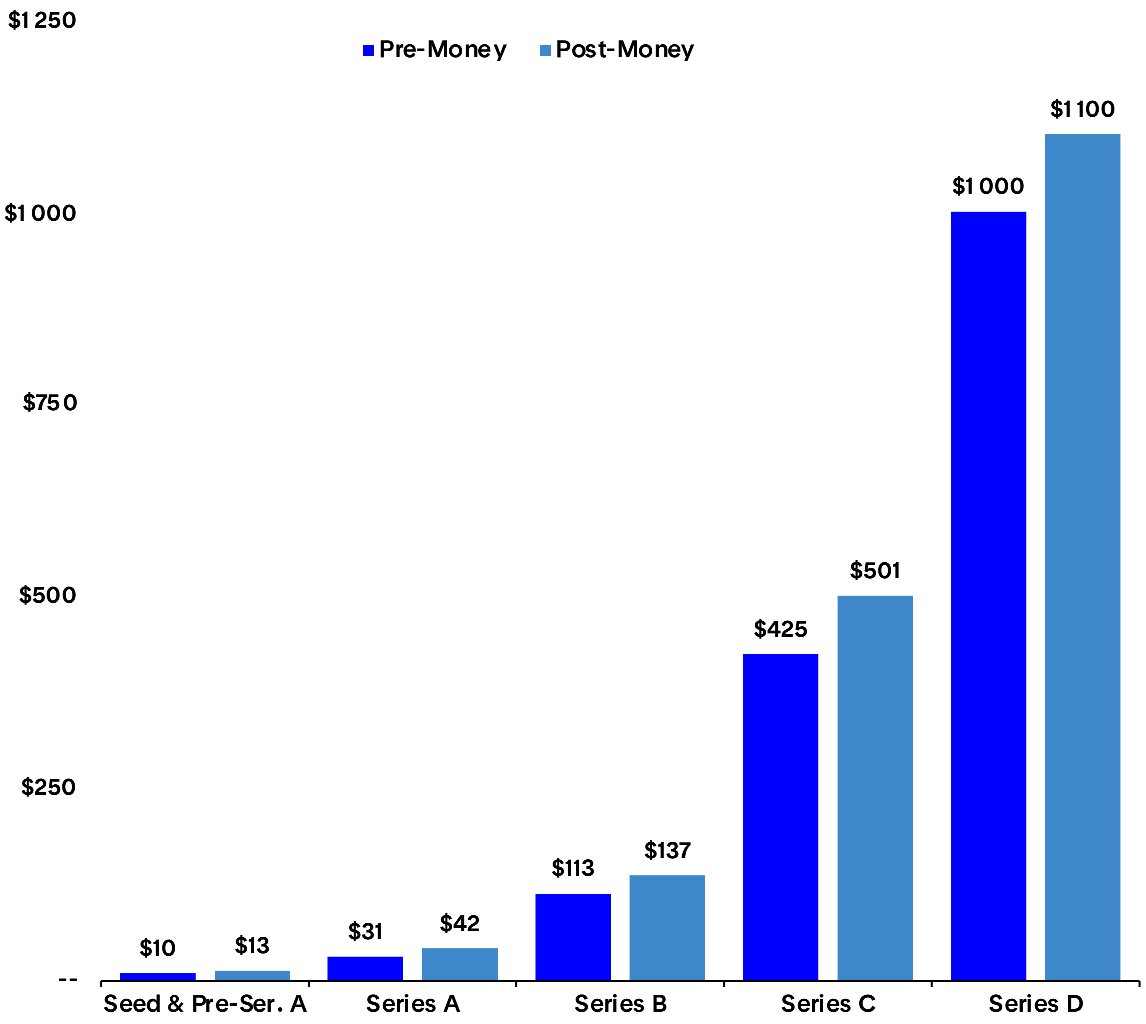

	Seed & Pre-Ser. A	Series A	Series B	Series C	Series D
# of Investments	703	787	315	116	36
Pre-Money Valuation	$10	$31	$113	$425	$1,000
Equity Raised	$3	$11	$24	$76	$100
Post-Money Valuation	$13	$42	$137	$501	$1,100
% Equity Acquired	24.5%	26.5%	17.3%	15.1%	9.1%
Valuation Step-Up		2.4 x	2.7 x	3.1 x	2.0 x
Time from Prior Round (yrs)		1.3	1.2	1.2	1.2
Company Age (yrs)	1.9	3.2	4.5	5.4	5.2

Source: Pitchbook. Screen represents all United States based companies who have received an equity investment of at least $1mm from January 2019 through February 2023. Incudes all investments received by a company during the period, not just the most recent investment. Does not include companies who did not disclose the amount raised or any details on valuation. Figures represent medians.

Trends and Stats in Artificial Intelligence Investing (Jan. 2019 – Feb. 2023) (cont'd)

Series A and B median entry valuations have doubled from 2019 levels; Series A and Series B median round sizes have increased by 33% and 40% since 2019, respectively

2019 (n = 362)

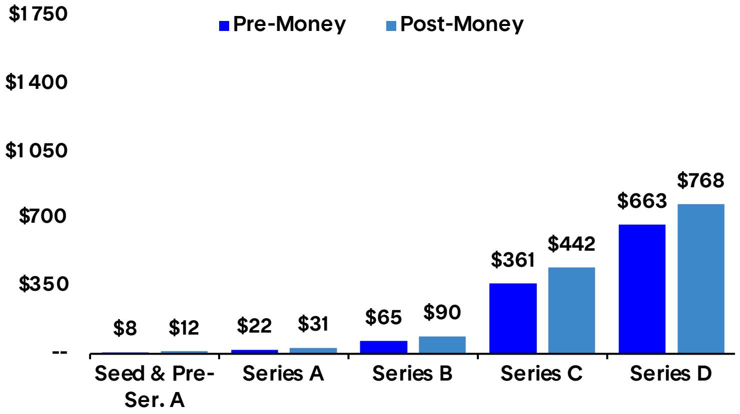

2020 (n = 410)

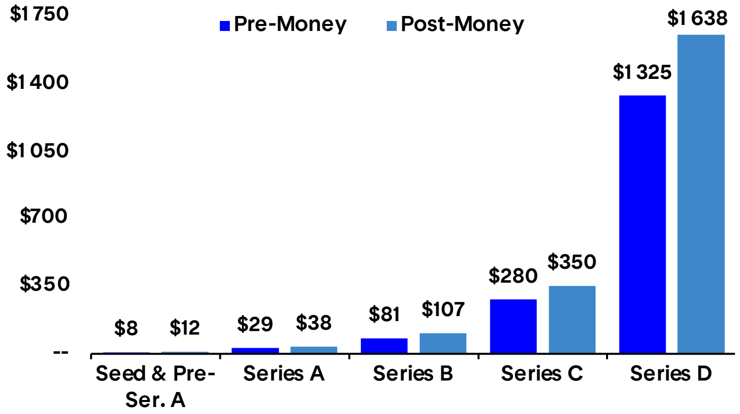

2021 (n = 650)

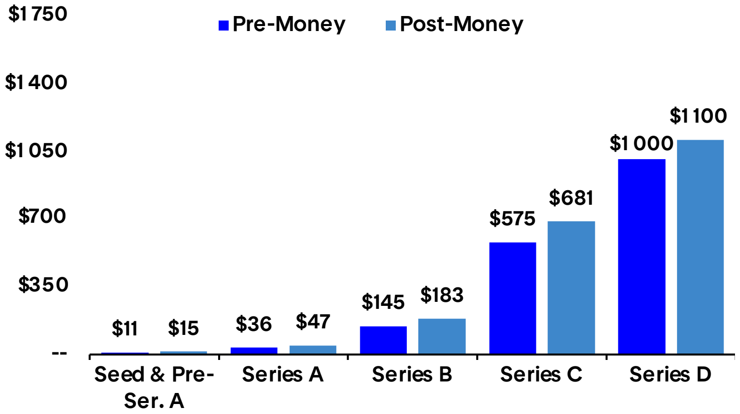

2022 (n = 518)

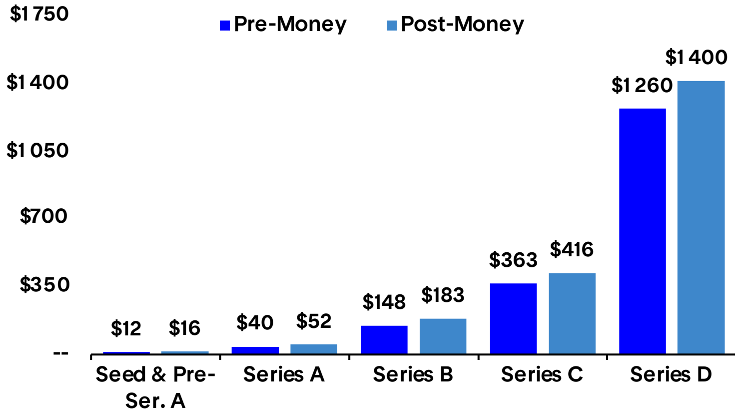

($ in mm)	2019	2020	2021	2022
Pre-Money Valuation				
Series A	$22	$29	$36	$40
Series B	65	81	145	148
Series C	361	280	575	363
Equity Raised				
Series A	$9	$9	$11	$12
Series B	25	26	38	35
Series C	82	70	106	53
% Equity Acquired				
Series A	30%	25%	23%	24%
Series B	28%	24%	21%	19%
Series C	18%	20%	16%	13%
Company Age (yrs)				
Series A	2.6	2.8	3.3	3.8
Series B	4.1	4.5	4.2	4.8
Series C	2.6	5.6	5.1	5.9

Source: Pitchbook. Screen represents all United States, artificial intelligence companies who have received an equity investment from January 2019 through February 2023. Incudes all investments received by a company during the period, not just the most recent investment. Does not include companies who did not disclose the amount raised. Figures represent medians.

Illustrative AI Investment Sensitivity – Entry Revenue Multiples (Jan. 2019 – Feb. 2023)

Series A

($ in mm)	ARR				
Pre-Money Valuation	$1.0	$1.5	$2.0	$2.5	$3.0
$25	25 x	17 x	12 x	10 x	8 x
$28	28 x	19 x	14 x	11 x	9 x
$31	31 x	21 x	16 x	12 x	10 x
$34	34 x	23 x	17 x	14 x	11 x
$37	37 x	25 x	19 x	15 x	12 x

Series B

($ in mm)	ARR				
Pre-Money Valuation	$2.0	$4.0	$6.0	$8.0	$10.0
$91	45 x	23 x	15 x	11 x	9 x
$102	51 x	26 x	17 x	13 x	10 x
$113	57 x	28 x	19 x	14 x	11 x
$125	62 x	31 x	21 x	16 x	12 x
$136	68 x	34 x	23 x	17 x	14 x

Series C

($ in mm)	ARR				
Pre-Money Valuation	$8.0	$11.0	$14.0	$17.0	$20.0
$340	42 x	31 x	24 x	20 x	17 x
$382	48 x	35 x	27 x	22 x	19 x
$425	53 x	39 x	30 x	25 x	21 x
$467	58 x	42 x	33 x	27 x	23 x
$510	64 x	46 x	36 x	30 x	25 x

Series D

($ in mm)	ARR				
Pre-Money Valuation	$20.0	$25.0	$30.0	$35.0	$40.0
$800	40 x	32 x	27 x	23 x	20 x
$900	45 x	36 x	30 x	26 x	23 x
$1,000	50 x	40 x	33 x	29 x	25 x
$1,100	55 x	44 x	37 x	31 x	28 x
$1,200	60 x	48 x	40 x	34 x	30 x

Note: Median pre-money valuation based on median values as shown on page 10. The additional pre-money valuations are based on a (20%) to 20% range, in 10% increments. ARR ranges are illustrative to show an implied multiple at each range.

Recent AI Investment Case Study: Jasper

- In October 2022, Jasper raised a $125 million Series A at a $1.5 billion post-money valuation
  - — The investment was led by Insight Partners with participation from Bessemer Venture Partners, IVP, Coatue and others
- Jasper leverages AI to generate content for blog articles, social media posts and website copy
  - — Using the platform, customers can describe in natural language what they want Jasper to write, whether a keyword-rich piece designed to rank well in search engines or existing content repurposed with additional context
  - — Jasper's language models are trained on 10% of the web and are fine-tuned for "customer specificity"
- According to the company, Jasper has 70,000+ customers and generated $45 million in revenue in 2021 and expects to reach $75 - $90 million in revenue in 2022
  - — Implied Pre-Money Multiple (2021): 31x
  - — Implied Pre-Money Multiple (2022): 15x – 18x

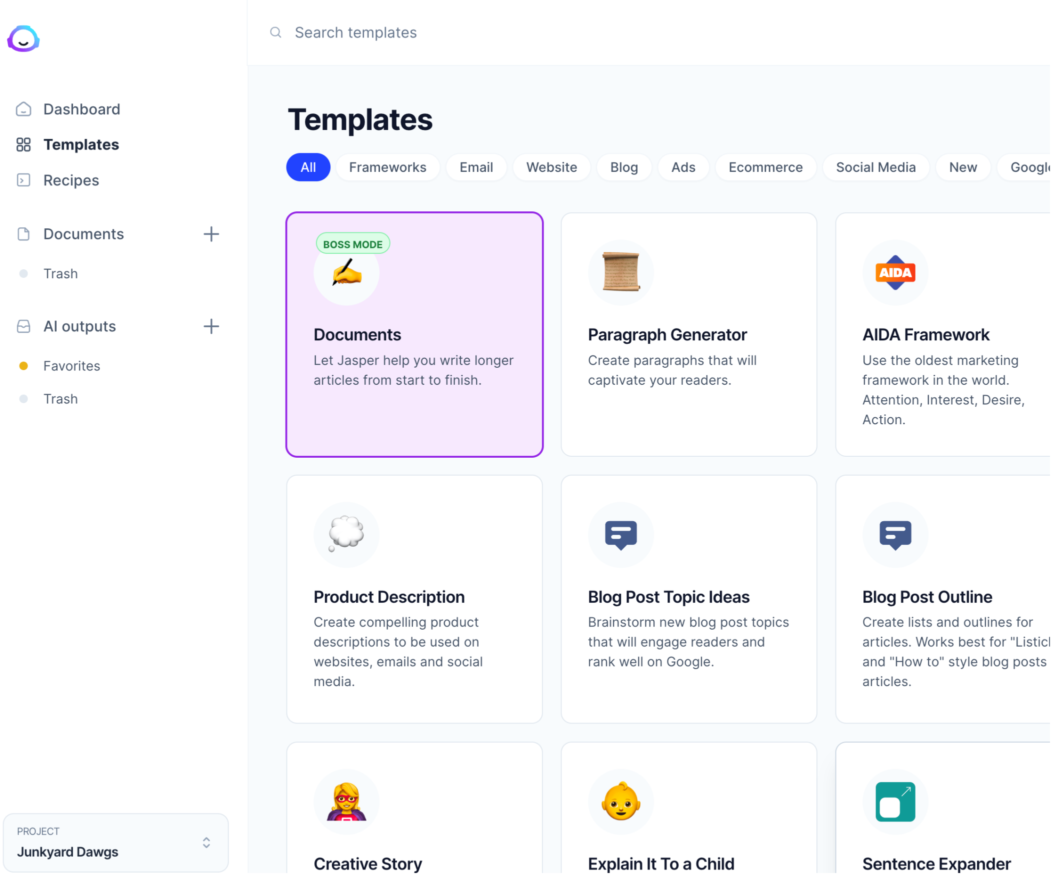

"With the advent of OpenAI's GPT-3, we saw an opportunity to launch an AI content platform that could help businesses and professional creators brainstorm and do their work more quickly and efficiently. The folks that will win at generative AI will be the ones that have the best feedback loops – we're committed to building the best feedback to AI loop." – Dave Rogenmoser, CEO

Source: TechCrunch.
Note: Implied Pre-Money Multiples are based on Revenue, not ARR. Based on ARR, the multiples would be lower as ARR would be higher than revenue.

AVP’s AI Market Map & Diligence Checklist

Basis of Presentation

- The following page presents AVP's AI Market Map[a], which is segmented into four categories
  - Cross-Industry Applications
  - Verticalized Solutions
  - Generative AI
  - Developer Tools, Model Deployment and Operations
- The market map was largely constructed based on detailed Pitchbook screening and supplemented by other publicly available AI market maps / industry reports
  - Pitchbook Screen Criteria:
    - Pitchbook Industry Tag / Keywords: "AI", "Artificial Intelligence", "Deep Learning"
    - Year Founded: 2012 – 2022
    - Geography (HQ): United States, Canada, Europe, Israel
    - Financing Round: Seed, Series A, Series B, Series C
    - Deal Size[b]: $2mm – $60mm
    - Total Capital Raised: $3mm – $75mm
    - Post-Money Valuation: $5mm – $400mm
    - Employee Count: 10 – 250
    - Investor-Backing: Venture Capital
  - The market map is specific to AVP in the sense that it generally focuses on companies that may be interesting for AVP's early stage and growth funds[c]

(a) Market map was made by AVP and is not a duplicate of an industry report / banker presentation. While third party reports were referenced as part of the creation of the market map, all research was conducted by AVP employees.
(b) Financing size for any round a company received, not just the latest funding round.
(c) For Generative AI specifically, some companies may be earlier stage than AVP typically looks at given the industry is very nascent.

AI Market Map – Sectors & Companies of Interest to AVP

Cross-Industry Applications

Search
miso, NewtonX, seek, Hebbia, kyndi, GOLDEN, Divinia, vectara, DASHWORKS, Pinecone, brave, neeva, Lucy, Twelve Labs, YOU

Unstructured Data Analysis
KLARITY, allganize, COACTIVE, acodis, automation hero_, super.ai, DOCUGAMI, INDICO DATA, veryfi

BI & Workflow Automation
intento, mem, fintastic, AVLINO, Tourmo, staircase, bardeen, Black Crow AI, socrates.ai, Harmonic, Faraday, DeepSee, RETINA, ALKYMI, avora, Continue AI, SWAY AI

Computer Vision
chooch, cogniac, VOXEL, KOGNIZ, Clearview.ai, Matroid, EVERGUARD, deepomatic, alcatraz ai, Superb AI, Ambient.ai, VOXEL51, alwaysAI, Mobius Labs, DRISHTI, constru

Speech & Voice
clinc, Gridspace, AVOMA, kea., dasha Artificial Intelligence, RASA, trint, Alan, talka, ava, Corti, ELSA

Customer Experience
pypestream, proton.ai, EVOLV AI, LEVITATE, netomi, remesh, TechSee, Chattermill, Neuron7, hello customer, Ultimate., Lily AI, truelark, Balto, Prodigal, dashbot, zingtree, worthix, Converse Now, Elise AI, posh, FLIP, UNSUPERVISED

Location & Mapping
BEANS.AI, SAFEGRAPH, unacast., cuebiq, FLOODBASE, pathr.ai, CARTO, MOBI, REPLICA, SLAMCORE, AtlasAI

Privacy & Security
Cape Privacy, privacy dynamics, skyflow, TRUU, Ketch, neosec, PRIVATEAI, THREAT WARRIOR, TripleBlind, Duality, ENVEIL, VENARI, Deepkeep, Opaque, DARKLIGHT, InsightFinder, Culture AI, Vouched, HORIZON3.ai, MixMode, SENSEON

Sales & Marketing
stratifyd, amplemarket, Intellimize, onescreen.ai, involve.ai, LAVENDER, UNIQUE, humanlinker, Databook, sturdy, Dragonfly AI, HeadsUp, MARQVISION, anyword, Quattr, BlueOcean

Generative AI

Copy & Ads
WRITER, HyperWrite, PODCASTLE, craftly.ai, textio, copysmith, autobound, copy.ai, memorable, Demandwell, MarketMuse, regie.ai

Photo & Video
tavus, Facet, synthesia, Colossyan, PAPERCUP, elai., Hour One, D-ID, PICTORY, Waymark, invideo, Rephrase.ai, windsor.ai, REMBRAND

Graphics & Design
Mutiny, Simplified, HYPAR, CSM, uizard, PRISMA LABS, Let's Enhance.io, BEAUTIFUL [AI], mirage, VIZCOM

Summarization & Insights
FATHOM, xembly, Fireflies, Grain, hume, agolo, Reduct Video, cogram, viable

TTS & STT
symbl.ai, MURF AI, LOVO, RESEMBLE.AI, Speechify, Speak, Otter.ai, WELLSAID, AssemblyAI

Text & Code Automation
Blaze, tabnine, mutable.ai, Fig, COMPOSE AI, warp, SimplyPut

Verticalized Solutions

Healthcare
leal, Rad AI, INCEPTO, theator, ONCOHOST, INFINITUS, HUMA.AI, healx, elephas, Curai Health, anumana, element5, Capitol, MENDEL, abridge, ZEPHYR AI, AQEMIA, Whiterabbit.ai, AIGNOSTICS, enlitic, regard, nucleai, Pearl, DeepScribe, ScienceIO, MACRO-EYES, UNLEARN

Commerce, CPG & Procurement
BOLD METRICS, supplypike, Particl, arena, arkestro, DEEPNORTH, onebeat, Oresma, syte, KLEVU, sensei, GATHER AI, Harmonya, tastewise, VALYANTAI, dataviva, Alloy.ai, Yogi, competera

Finance & Insurance
EvolutionIQ, convr, zesty.ai, chip, finally, FUNDGUARD, entera, JUSTT, GRADIENT AI, Counterpart, layr, LeaseLock, daloopa, PLANCK, SESAMm, THETARAY, abound, RIBBIT AI, ABAKA, surance.io, Anomaly

Gov't, Intelligence & Defense
ZERO EYES, AXION RAY, COGNITIVE SPACE, BLACKBIRD.AI, AIRSPACE INTELLIGENCE, spot.ai, actuate, MODERN INTELLIGENCE, Enabled Intelligence, Deep Sentinel, zencity, NOTRAFFIC, percipient.ai

Industrial & Logistics
Optimal Dynamics, flowlity, intenseye, Afresh, ORCA AI, SewerAI, noodle.ai, VERUSEN, HaydenAI, calculum, greyparrot, EXPEDOCK, DOXEL, siteaware, SWAPP, ALICE, Zendrive

Legal & Compliance
COMPLIANCE.AI, trellis, Litty, 4CRISK.ai, casetext, Luminance, MinervaAI, Hummingbird, Robin, InCountry, lexion, REGOLOGY, PACTUM, juro, laurel

Dev. Tools, Model Deployment and Ops

AI Model Operations
SELECTOR, Numbers Station, Arthur, PRYON, Lightning AI, Akkio, VIRTUALITICS, FedML, ROBUST INTELLIGENCE, arize, LatticeFlow, FIXIE, Galileo, WHYLABS, fiddler, Smartcat, mosaicML, Ozeni, Augtera

Data Quality & Observability
testRigor, zengines, bodo.ai, kubit, MANTIUM, GETVISIBILITY, aggua, iterative, lang.ai, diffblue AI for Code, Akridata

Labeling & Annotation
Heartex, ALEGION, Kili, Dataloop, SuperAnnotate, ENCORD, surge AI, datasaur.ai

Synthetic Data
TONIC The Fake Data Company, gretel, SYNTEGRA, Synthesis.ai, parallel domain, Diveplane, RENDERED.AI, datagen

Note: Market map includes companies in the U.S., Canada, Europe and Israel that may be interesting for AVP's early & growth stage funds. The market map purposefully does not include large market leaders which have raised significant capital. For presentation purposes, companies are mutually exclusive between categories; however, several companies could fall in multiple segments (e.g., TTS and STT, Speech & Voice; Sales & Marketing, Copy & Ads).

AI Market Map (cont'd)

Nearly 1,000 artificial intelligence companies generally fall within AVP's early stage and growth stage criteria

AI Companies by Location

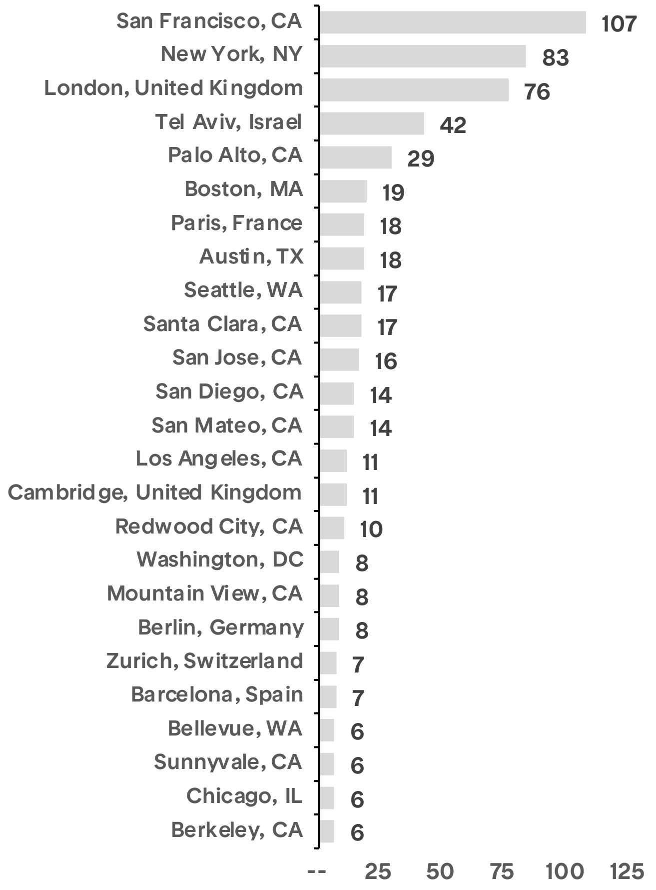

Select AVP Experience with AI Companies

D-ID

Sector: Generative AI – Photo & Video
Geography: Israel
Company Description: D-ID's creative AI technology takes images of faces and turns them into high-quality, photorealistic videos. D-ID's technology can combine images with audio or text to give them expression and speech.
AVP Experience: US Early Stage team led D-ID's Series A in 2020 (follow-on investment in Series B in Q1 2022)

INCEPTO

Sector: Verticalized Solutions - Healthcare
Geography: France
Company Description: Incepto offers artificial intelligence-based medical imaging applications designed to offer computed tomography scans, chest x-ray and detections for conditions such as fracture and bowel occlusion.
AVP Experience: EU Early Stage team led Incepto's Series A in 2019 (follow-on investment in Series B in Q3 2022)

phenom

Sector: Cross-Industry Applications – HR[(a)]
Geography: United States
Company Description: Phenom's AI-powered talent relationship management platform helps companies hire faster, develop better and retain employees longer, ensuring that organizations are well-positioned to grow.
AVP Experience: US Growth team led Phenom's Series B in 2018 (follow-on investments in Series C and Series D)

THETARAY

Sector: Verticalized Solutions – Finance & Insurance
Geography: Israel
Company Description: ThetaRay's AI-powered SaaS AML transaction monitoring and screening solution enables fintechs, banks and regulators to embed trust in cross-border and domestic payments while driving financial growth.
AVP Experience: EU Growth team evaluated ThetaRay during its Series C fundraise in Q1 2023.

INDICO DATA

Sector: Cross-Industry Applications – Unstructured Data Analysis
Geography: United States
Company Description: Indico uses AI and ML technology to automate the intake and understanding of unstructured documents, emails, images, videos and audio files.
AVP Experience: US Early Stage team evaluated Indico during its Series B fundraise in Q4 2022.

Source: Pitchbook.
Note: Pitchbook screen follows same methodology as described in the Basis of Presentation. Not all ~1,000 companies are included on the market map on the prior page.
(a) HR is not included as a designated segment on the market map, however, there are likely several companies within the space that could be of interest to AVP.

The AI Investment Diligence Checklist

For AI companies in the application layer, there are multiple factors to consider when evaluating differentiation

1. Speed of Iteration & Updates
What is considered state-of-the-art in AI is changing at a rapid pace; companies need to stay up to date with the latest technology, architecture and models in order to offer a differentiated solution

2. Access to Proprietary Data
LLMs are trained on billions of parameters, but are not specific to any given company or use case – without proprietary, private training data, companies cannot refine these models to suit their exact product / use case

3. Product & Customer Experience
With access to AI becoming increasingly accessible via OSS and private APIs, differentiation will increasingly be found in the product and customer experience itself

4. Private vs. Open Source Architecture
Private solutions (e.g., OpenAI's API) may be preferable over open source software (e.g., HuggingFace) as private solutions do not require continued maintenance, have higher levels of security and offer better support

5. Enterprise Grade Standards
Data privacy standards and regulatory compliance are non-negotiable when it comes to AI-enabled applications looking to sell to enterprise clients

6. Quantifiable ROI & Solution Relevance
It is important to understand how acutely a solution is solving a specific problem – market-leading solutions will explicitly reduce operating costs and increase worker efficiency / accuracy within a given niche

7. Commitment to Responsible AI (ESG)
When incorporating AI into SaaS applications, companies should ensure that steps are taken to reduce bias, increase transparency and uphold ethical standards

8. Multi-Platform
The ability to support or easily switch between multiple model platforms (e.g., Anthropic, OpenAI, Stable Diffusion) is important as different enterprise customers may require varying standards

9. Technical Founders
AI is complicated – while many companies in the application layer won't be building their own AI from scratch, the technical sophistication and experience of founders will still be important to consider

10. Data Feedback Loop & Network Effects
Model performance and accuracy will rapidly increase if there is a flywheel effect (e.g., real-time user engagement data is collected and used to further train / finetune models, more data is captured as users increase)

ESG Considerations

Energy Consumption & Climate Change

AI requires extensive energy to function; however, there are positive applications of AI which may help to lower our carbon footprint

Benefits

- Despite its massive energy requirements, AI does have practical applications that can help to lower our carbon footprint
  - It's a bit of a catch-22, but AI itself can analyze mass amounts of complex climate change data to help us better understand what is happening in the world around us and inform policy decisions that address specific problems
- 87% of private and public sector CEOs believe AI is an essential tool in the fight against climate change
  - Moreover, 43% of organizations confirm having a vision for using AI in their own climate change efforts
- Based on research by BCG, AI can help to achieve a 5 – 10% reduction in global carbon emissions (2.6 – 5.3 gigatons)
  - Companies can use AI-powered data engineering to automatically track emissions throughout their carbon footprint
  - Predictive AI can forecast future emissions across a company's carbon footprint, in relation to current reduction efforts, new carbon reduction methodologies and future demand
  - By providing detailed insight into every aspect of the value chain, prescriptive AI and optimization can improve efficiency in production, transportation, and elsewhere, thereby reducing carbon emissions and cutting costs

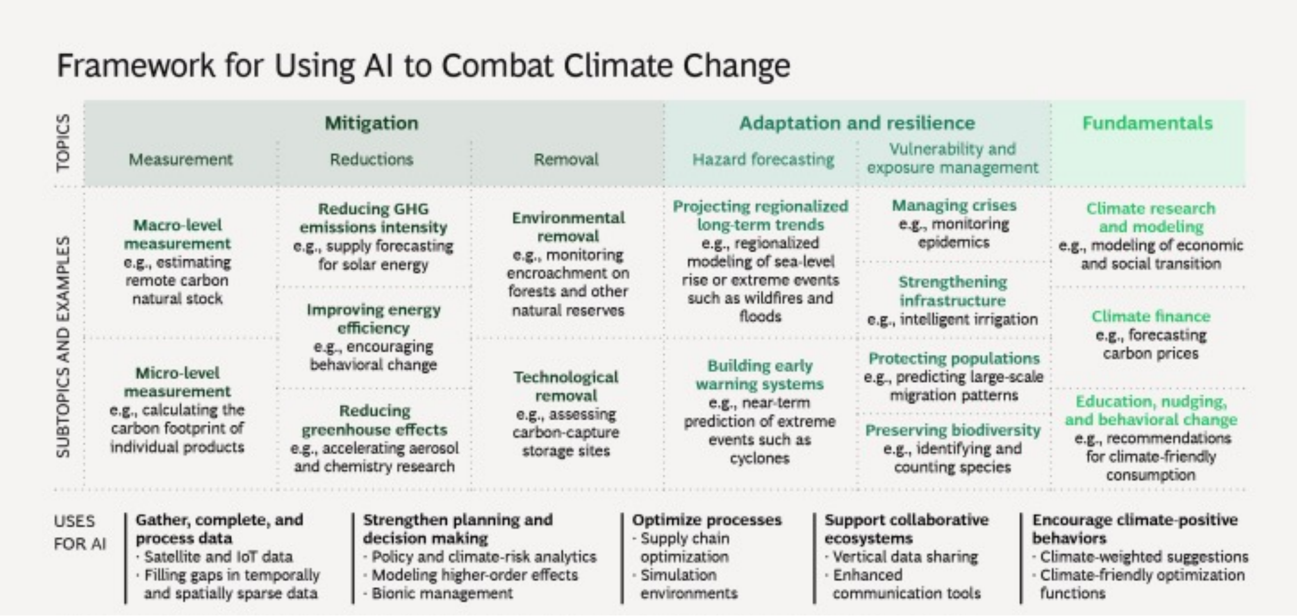

Source: Bloomberg, BCG and other publicly available information.

Energy Consumption & Climate Change (cont'd)

AI requires extensive energy to function; however, there are positive applications of AI which may help to lower our carbon footprint

Considerations

- As machine learning models have increased in scale, so too have the computational and energy requirements needed to create and sustain them
  - In 2020, computers consumed roughly 4-6% of the global electricity supply (up from 1-2% in 2018)
  - By 2030, this figure is projected to rise between 8-21%
- Taking OpenAI's GPT-3 model as an example, it took 1,287 MWh to train the model which generated a carbon footprint is 552 CO2e
  - For comparison, 1,287 MWh is comparable to powering 120 US homes in a year while 552 CO_2e is about as much as 110 US cars emit in a year
  - While training the model has significant upfront power costs, in some cases training is only 40% of the power burned by the model with the rest coming from billions of live requests to use the platform
- At the same time, the infrastructure required to power AI (i.e., data centers) has also accelerated
  - Between 2017 and 2020, energy consumption and carbon emissions associated with data centers doubled
  - Data centers are usually running near full utilization, meaning that a 20-megawatt facility would consume enough energy to power roughly 16,000 households
  - The location of a data center can greatly affect the amount of energy consumed; for example, a model trained in the Pacific Northwest would generate less carbon due to its use of clean hydropower

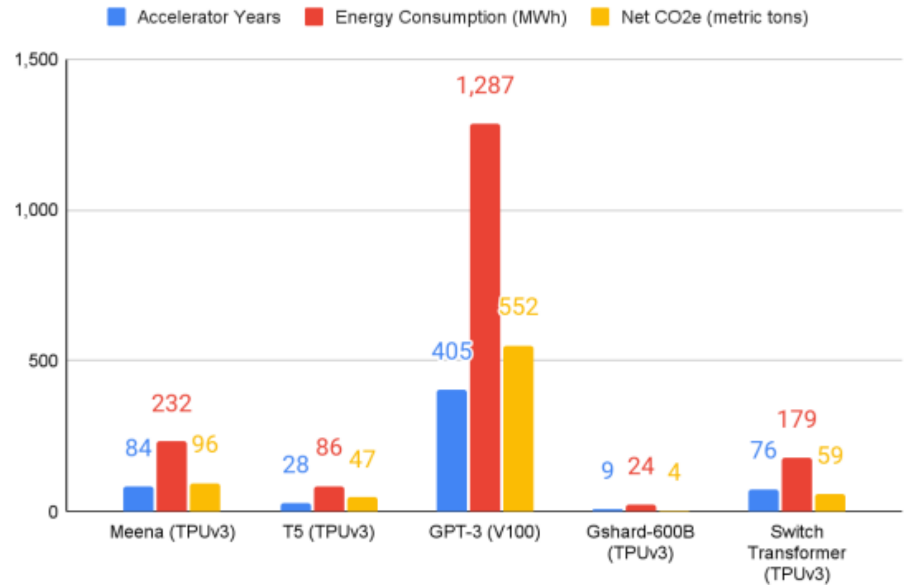

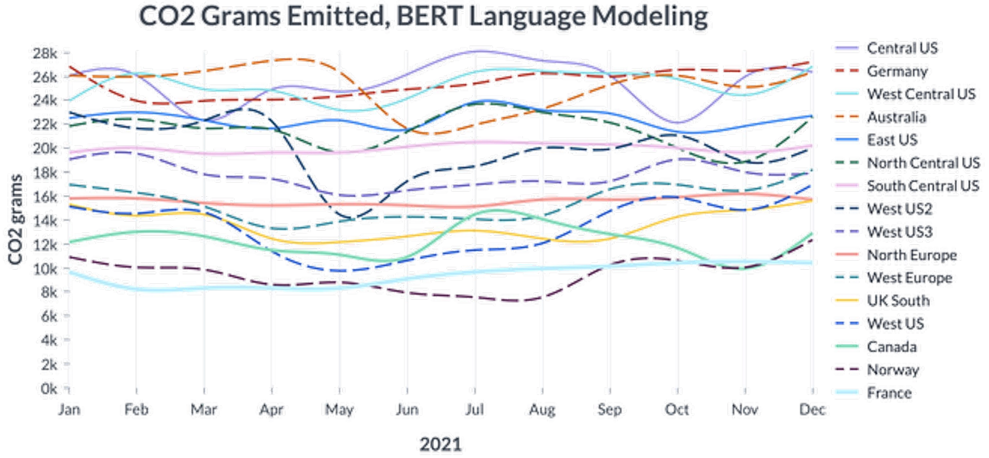

Source: "Carbon Emissions and Large Neural Network Training" (Patterson et al.), The University of Pennsylvania, Johns Hopkins, Forbes, Bloomberg, BCG and other publicly available information.

Responsible AI

Organizations scaling AI should be mindful of regulatory, privacy, security and ethical concerns

- Responsible AI is defined as the practice of designing, developing and deploying AI with the good intention to empower employees and businesses while fairly impacting customers and society
- AI is squarely on the minds of governments and regulators as an area of both promise and concern
  - In October 2022, the White House published its Blueprint for an AI Bill of Rights which included five main points
    - Automated systems should be safe and effective
    - Users should not experience algorithmic discrimination
    - Users should be protected from abusive data practices and have control over how their data is used
    - Users should know an automated system is being used and understand how it impacts them
    - Users should be able to opt out and talk to a human where possible
  - Similarly, the EU has proposed its own AI Act which would classify AI systems by risk and mandate various development and use requirements
- Organizations leading the push in AI, including Microsoft and Google, also have published (in 2017 / 2018) their own AI principles by which their systems will be measured

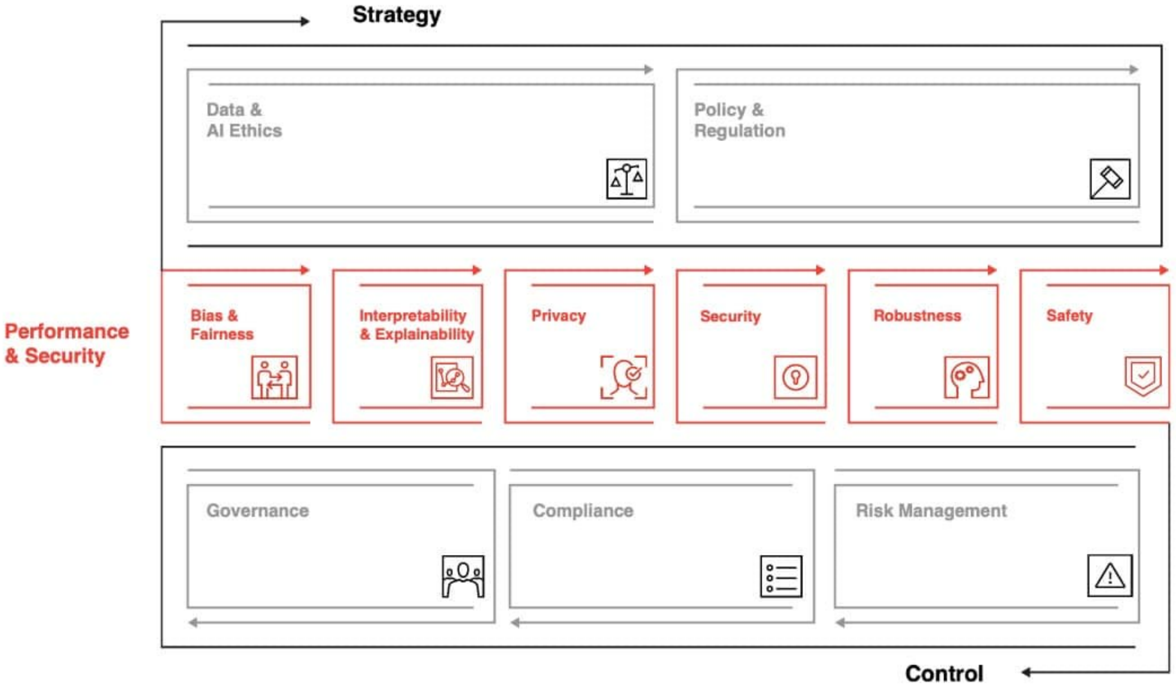

Source: Accenture, Google, Microsoft and PwC.

Final Thoughts & Next Steps

What's Next for AI?

Multi-Modal AI

To date, large language models have been just that – language. Companies like OpenAI are focused on multi-modal AI, or the idea that the input data can come in multiple forms, whether text, image or video. Through multi-modality, solutions including ChatGPT can become multi-purpose, combining skills in language and images to make the AI better and understanding both.

Algorithm Improvements & Next-Gen Hardware

With model improvements, experts contend that GPUs will pick up speed and remain an important part of the "computational power" variable in the formula that drives the next AI leaps. However, some AI hardware under development, such as neuromorphic chips or even quantum computing systems, could factor into the new equation for AI innovation.

Powerful Verticalized Applications

Within healthcare, for example, AI will allow for much more personalized medicine and bring a revolution in the use of large medical datasets. We'll see patient-specific treatments—for example, ones created using your genomic and expression data, which are much more likely to work.

Closing the Gap Between AI and Human Understanding

Deep learning continues to generate useful applications, but considerable work remains at the analytical level to understand how human cognition works in supporting problem-solving and critical thinking and creativity.

Efficient AI Models

Over the past five years, the following statement has largely held true – the larger the model, the more accurate the AI will be; however, a key breakthrough will occur when an extremely small dataset is proven to train AI systems as well as large datasets. Researchers are pushing to figure out ways to train systems on less data and are confident they'll find a viable solution. As a result, AI experts expect the "data" variable in the AI growth equation to flip on its head, with small datasets overtaking big data as drivers of new AI innovation.

Unsupervised Learning

Unsupervised learning is the idea that AI systems are able to learn without human guidance or labels on the data that is fed into the systems to train them. Experts contend that we're still a long way away from complete unsupervised learning, but that the next wave of AI innovation will likely be fueled by deep learning models trained using a method somewhere between supervised and unsupervised learning.

Appendix: AI Architecture & Techniques

Convolutional Neural Networks (CNNs) – 1980s

- CNNs are deep learning algorithms that can take in an input image, assign importance (learnable weights and biases) to various aspects / objects in the image and be able to differentiate one from the other
  - — From a use case perspective, CNNs are primarily used for image and video recognition tasks, including image classification, object detection, face recognition, medical imaging and natural language processing
- CNNs have three main types of layers, which are the Convolutional Layer, Pooling Layer and Fully-Connected (FC) Layers
  - — Convolutional Layer: Core building block of a CNN, and it is where the majority of computation occurs
    - — After each convolution operation, a CNN applies a Rectified Linear Unit (ReLU) transformation which allows for faster and more effective training by mapping negative values to zero and maintaining positive values
  - — Pooling Layer: Conducts dimensionality reduction, reducing the number of parameters in the input
  - — FC Layer: Performs the task of classification based on the features extracted through the previous layers and their different filters

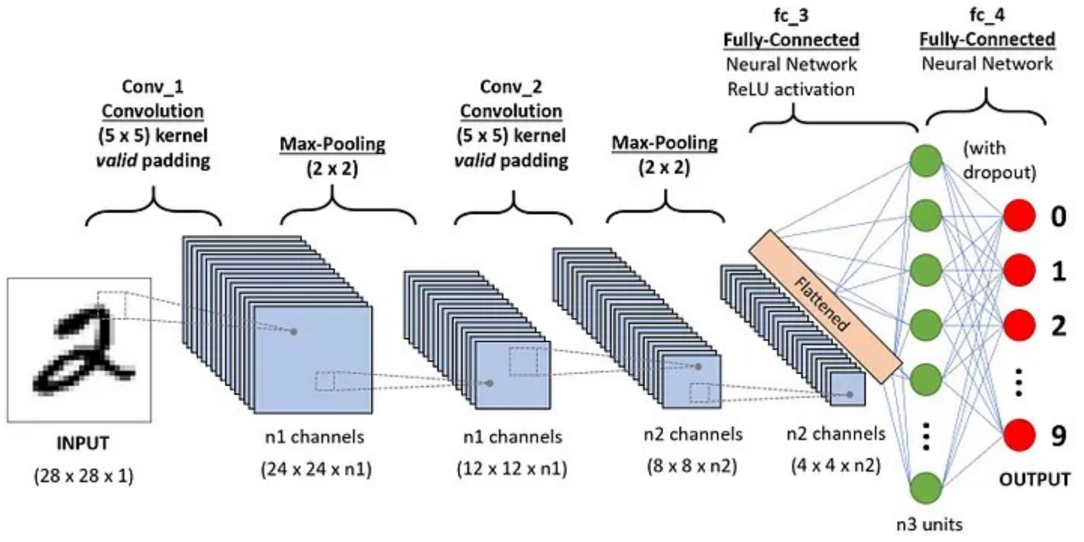

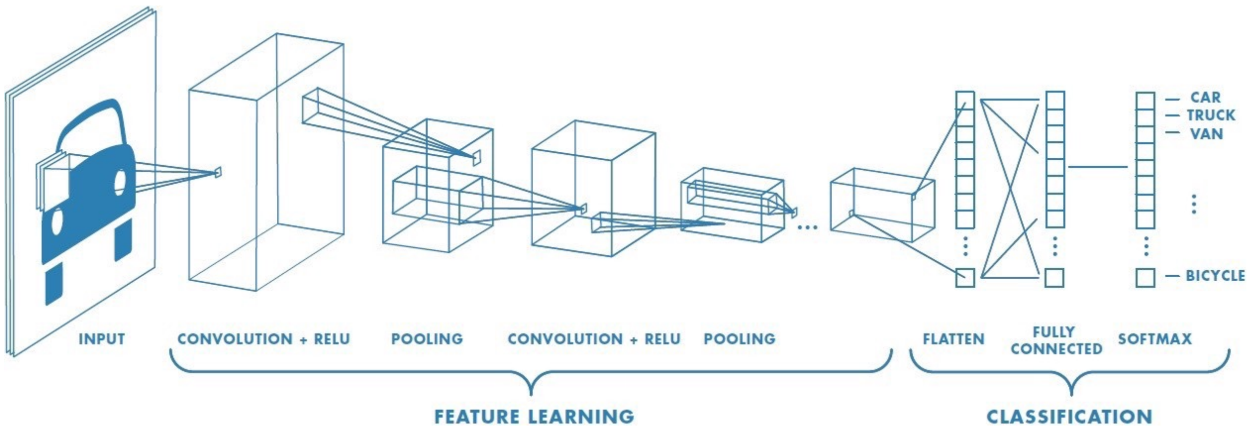

Source: IBM and other publicly available information.

Recurrent Neural Networks (RNNs) – 1986

- RNNs are a type of artificial neural network which use sequential data or time series data
  - These deep learning algorithms are commonly used for ordinal or temporal problems, such as language translation, natural language processing (NLP), speech recognition, and image captioning
- Like convolutional neural networks (CNNs), recurrent neural networks utilize training data to learn
  - They are distinguished by their "memory" as they take information from prior inputs to influence the current input and output
  - While traditional deep neural networks assume that inputs and outputs are independent of each other, the output of recurrent neural networks depend on the prior elements within the sequence
- Recurrent neural networks leverage backpropagation through time (BPTT) algorithm to determine the gradients, which is slightly different from traditional backpropagation as it is specific to sequence data
  - Through this process, RNNs tend to run into two problems, known as exploding gradients and vanishing gradients
  - One solution to these issues is to reduce the number of hidden layers within the neural network, eliminating some of the complexity in the RNN model

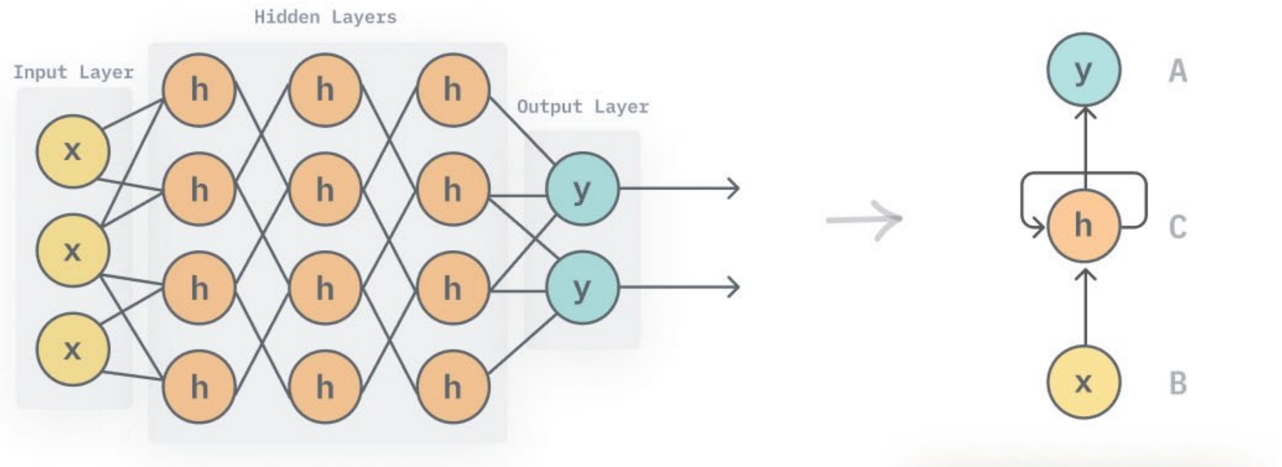

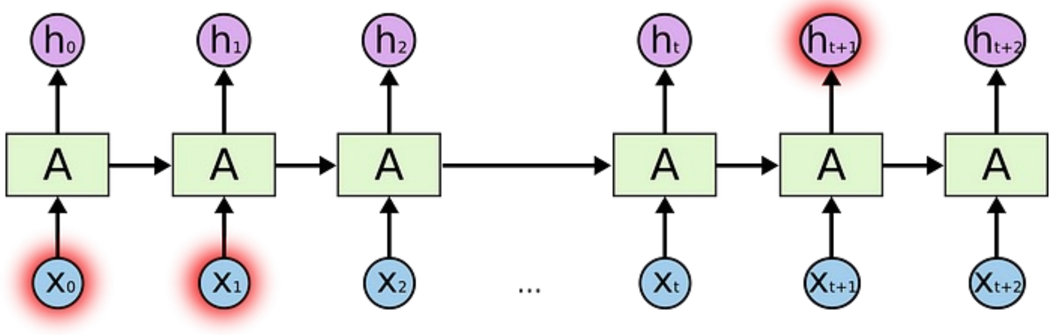

Source: IBM and other publicly available information.

LSTM (Long Short-Term Memory) – 1997

- LSTM is a popular type of RNN architecture which was introduced as a potential solution to the vanishing gradient problem
  - <u>Vanishing Gradient Problem</u>: When the gradient (i.e., a derivative of a function that has more than one input variable) is too small, it continues to become smaller, updating the weight parameters until they become insignificant (e.g., when this occurs, the algorithm is no longer learning)
  - The LSTM model works to solve the issue of long-term dependencies, or an RNN's inability to accurately predict the current state if the context to do so is well into the past
- For example, take the statement "Alice is allergic to nuts. She can't eat peanut butter."
  - The context of a nut allergy can help anticipate that peanut butter cannot be eaten because it contains nuts
  - However, if that context was a few sentences prior, then it would make it difficult, or even impossible, for the RNN to connect the information
- LSTMs have "cells" in the hidden layers of the neural network, which have three gates (input gate, output gate, forget gate)
  - These gates control the flow of information which is needed to predict the output in the network
  - For example, if gender pronouns, such as "she", was repeated multiple times in prior sentences, you may exclude that from the cell state

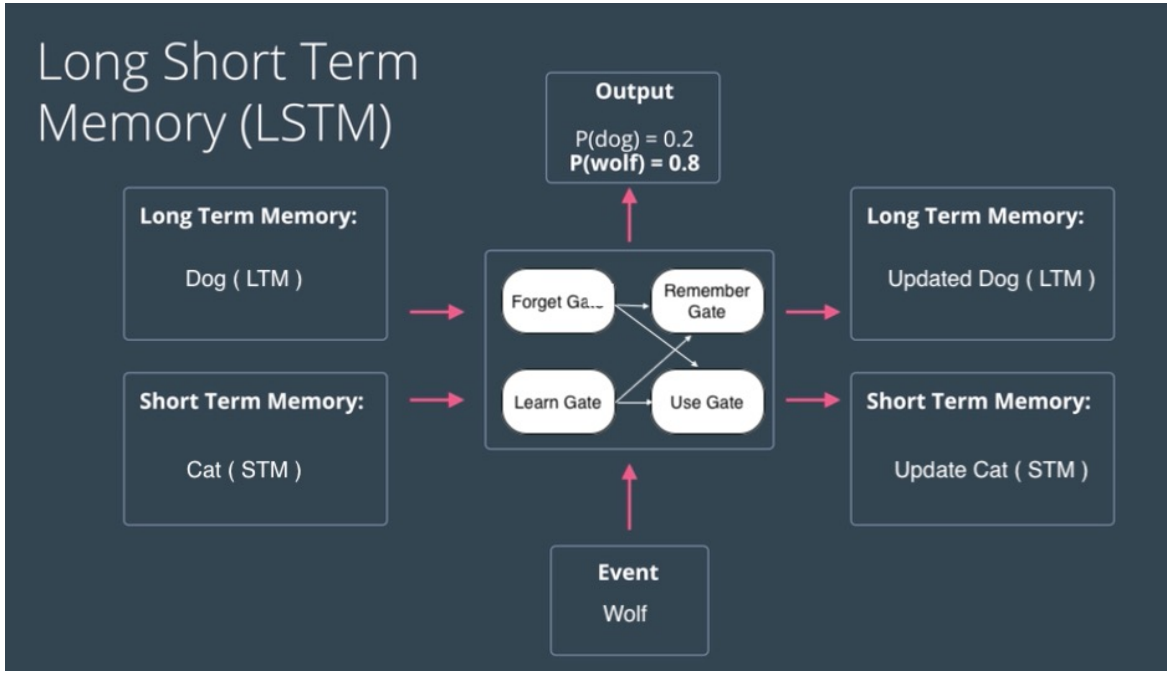

Source: IBM and other publicly available information.

Generative Adversarial Networks (GANs) – 2014

- GANs are algorithmic architectures that use two neural networks, pitting one against the other (thus the "adversarial") in order to generate new, synthetic instances of data that can pass for real data
  - — First described in a paper in 2014, GANs are used widely in image generation, video generation and voice generation
- One neural network (generator) generates new data instances, while the other (discriminator) evaluates them for authenticity (i.e., the discriminator decides whether each instance of data that it reviews belongs to the actual training dataset or not)
  - — Both networks are trying to optimize a different and opposing objective function, or loss function, in a zero-sum game – as the discriminator changes its behavior, so does the generator creating a double feedback loop
- However, there are several challenges when working with GANs
  - — Given that GANs involve two competing neural networks, there is double the amount of complexity vs. just training one neural network
  - — One network may overpower the other, such that neither can learn anymore
  - — GANs can also suffer from "mode collapse", or when the generator only learns a small subset of the possible realistic models

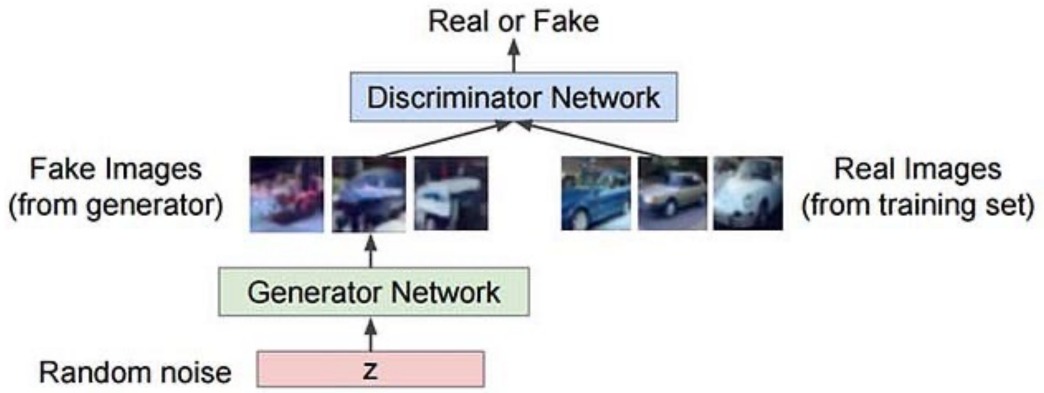

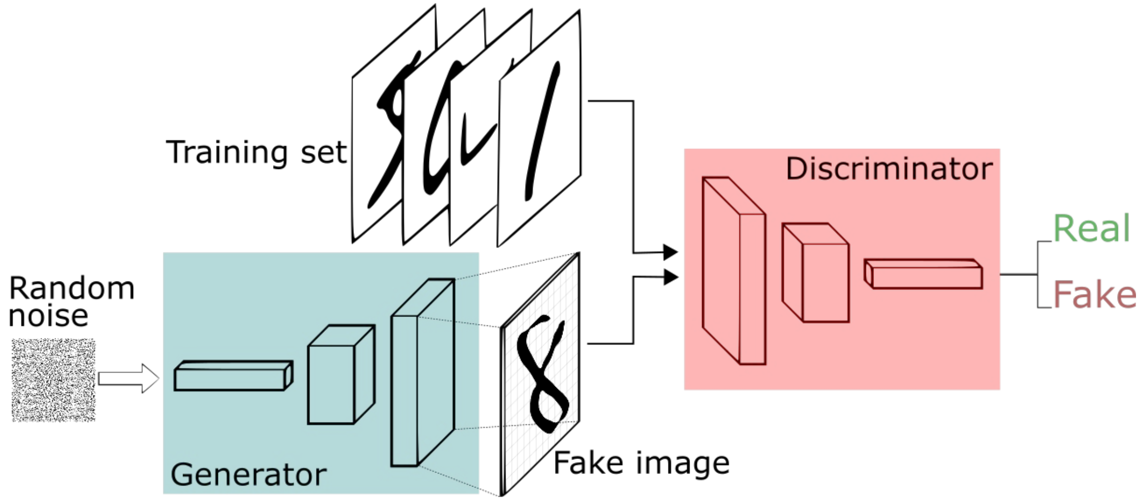

Source: Pathmind, Toptal and other publicly available information.

Transformers – 2017

- Transformer modes are neural networks that learn context and thus meaning by tracking relationships in sequential data (like the words in this sentence)
  - First described in a 2017 paper from Google, transformers are among the newest and one of the most powerful classes of models invented to date
- Transformer models apply an evolving set of mathematical techniques, called attention or self-attention, to detect subtle ways even distant data elements in a series influence and depend on each other
  - Transformers use positional encoders to tag data elements coming in and out of the network
  - Attention units follow these tags, calculating a kind of algebraic map of how each element relates to the others
  - Attention queries are typically executed in parallel by calculating a matrix of equations in what's called multi-headed attention
- Transformers are in many cases replacing CNNs and RNNs for NLP
  - Before transformers arrived, users had to train neural networks with large, labeled datasets that were costly and time-consuming to produce
  - By finding patterns between elements mathematically, transformers eliminate that need, making available the trillions of images and petabytes of text data on the web and in corporate databases

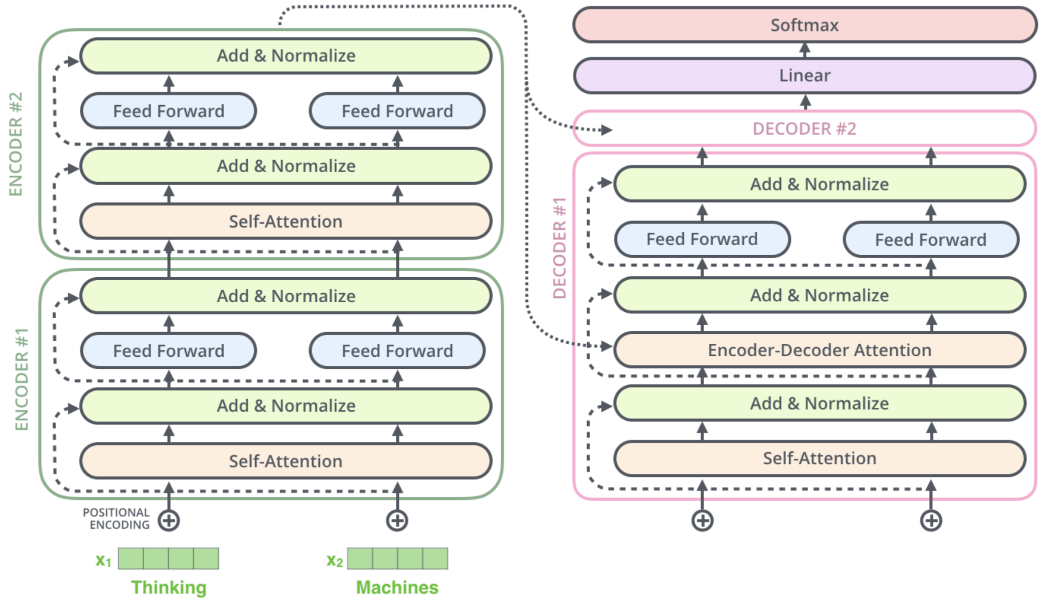

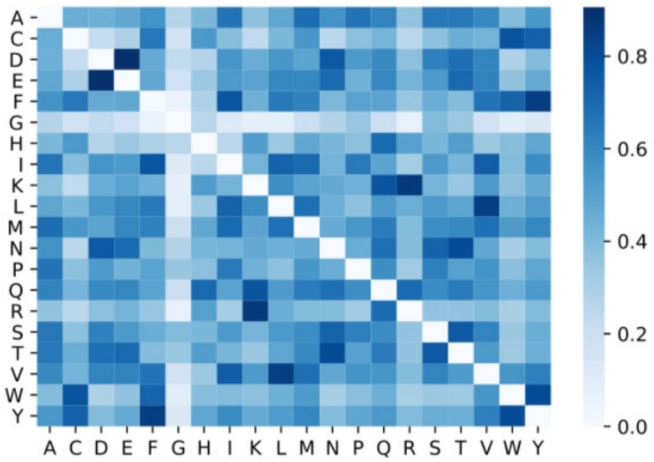

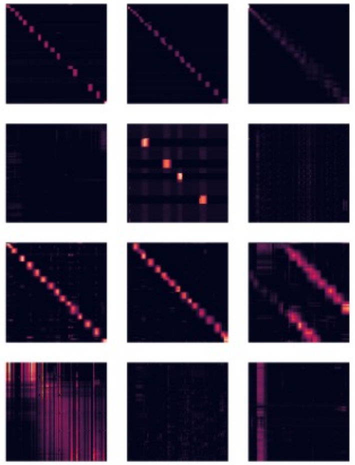

Source: Nvidia and other publicly available information.

Transformers – 2017 (cont'd)

1

Input Embeddings
- Embeddings are numerical representations (i.e., vectors) of pieces of information
- The representation captures the semantic meaning of what is being embedded and helps ML models understand relationships and similarities between words (e.g., hotel and motel have similar meaning)

2

Positional Encoding
- Describes the location of a piece of data in a sequence, storing information about word order in the data itself (vs. the network)
- By encoding data, the neural network learns the importance of word order from the data

3

Multi-Head Attention
- The attention mechanism allows a neural network to look at every single word in an input sentence at the same time before generating an output sentence (vs. reading each word sequentially in one direction)
- Multi-head attention in the encoder specifically uses a type of attention called self-attention which relates different positions of a single sequence in order to compute a representation of the sequence (i.e., self-attention allows a neural network to understand a word in the context of the words around it)

Self-Attention Example
- Take the following to sentences: 1) "What book are you reading?" and 2) "Did you book the appointment?"
- When looking at each sentence, it is clear to us that the word "book" has a differing meaning given the broader context surrounding it
- Using the self-attention mechanism, Transformer models can recognize that "book" is being used in two separate parts of speech by attending to "reading" and "appointment"

4

Decoder
- The decoder takes the vectors (text converted to vectors by the encoder) and predicts the next word in the output sequence based on the context learned from the encoder and the previously generated words in the output sequence (this process occurs until it has generated the entire output sequence)
- The decoder is trained using a process called backpropagation, which adjusts the model's parameters to minimize the difference between the predicted output sequence and the actual output sequence

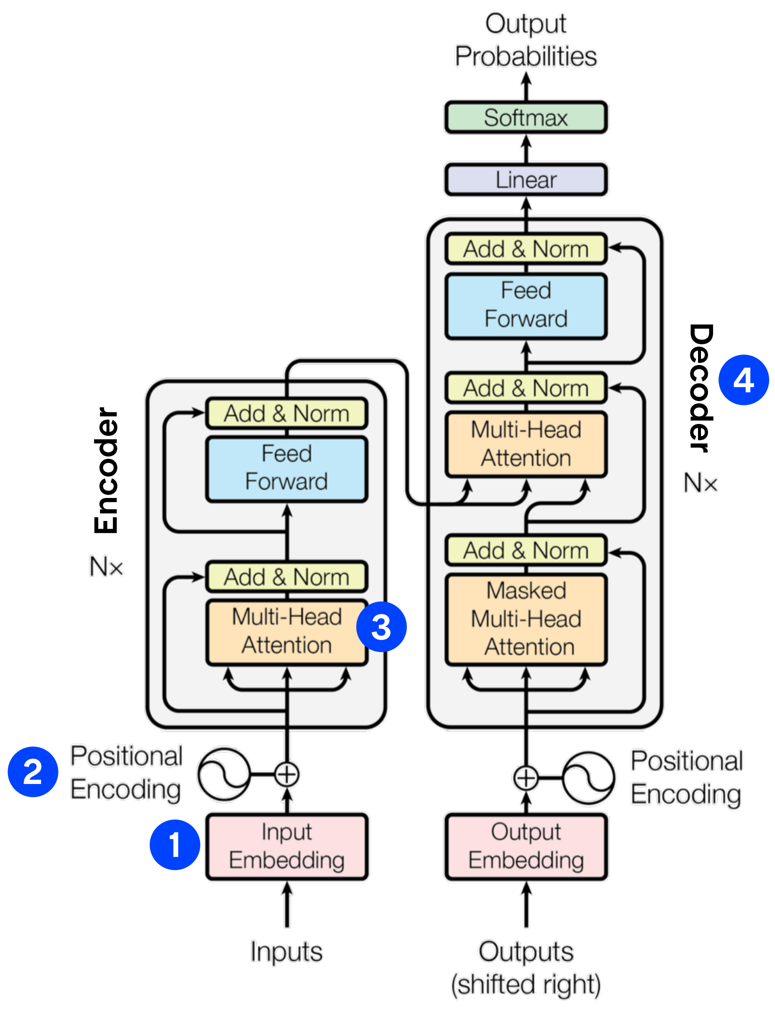

Figure 1: The Transformer - model architecture.

Source: Google, Hugging Face and other publicly available information.

Foundation Models – 2018 to present

- A foundation model is an AI neural network — trained on mountains of raw data, generally with unsupervised learning — that can be adapted to accomplish a broad range of tasks
  - — Foundation models generally learn from unlabeled datasets, saving the time and expense of manually describing each item in massive collections.
- With the previous generation of AI techniques, if you wanted to build an AI model that could summarize bodies of text for you, you'd need tens of thousands of labeled examples just for the summarization use case
  - — With a pre-trained foundation model, we can reduce labeled data requirements dramatically
    - — First, we could fine-tune it to domain-specific unlabeled data to create a domain-specific foundation model
    - — Then, using a much smaller amount of labeled data, potentially just a thousand labeled examples, we can train a model for summarization
    - — The domain-specific foundation model can be used for many tasks as opposed to the previous technologies that required building models from scratch in each use case
- Since the advent of the Transformer, there has been a rush to produce larger and larger foundation models (LLMs), including OpenAI's GPT and Google's BERT

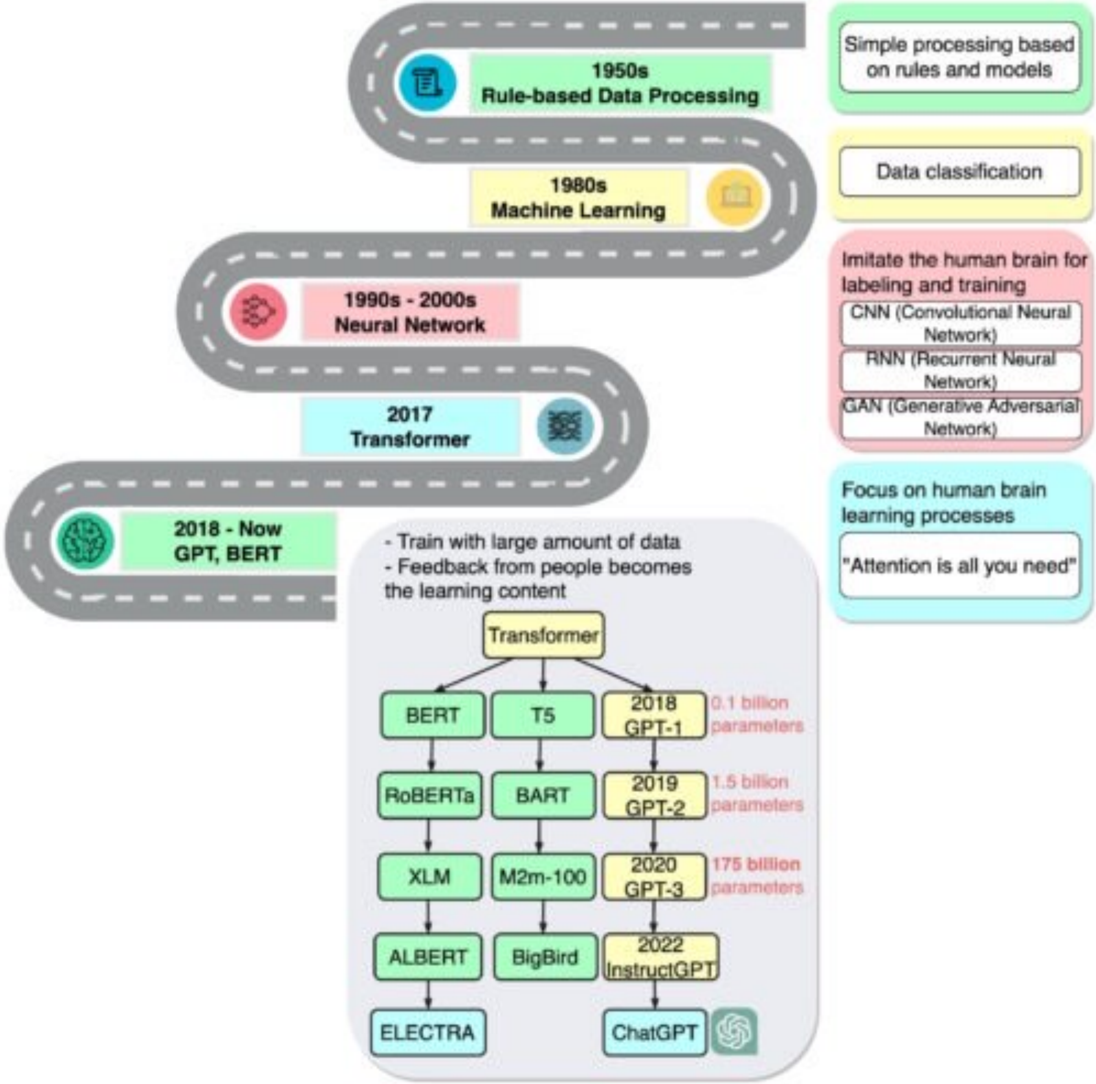

Source: Nvidia, IBM and other publicly available information.

Appendix: Overview of Generative AI

Overview of Generative AI

- Generative AI refers to a category of artificial intelligence algorithms that generate new outputs based on the data they have been trained on
  - Unlike traditional AI systems that are designed to recognize patterns and make predictions, generative AI creates new content in the form of images, text and audio
- Generative AI uses a type of deep learning called generative adversarial networks (GANs) to create new content
  - A GAN consists of two neural networks: a generator that creates new data and a discriminator that evaluates the data
  - The generator and discriminator work together, with the generator improving its outputs based on the feedback it receives from the discriminator until it generates content that is indistinguishable from real data
- Generative AI has a wide range of applications, including:
  - Images: Generative AI can create new images based on existing ones, such as creating a new portrait based on a person's face or a new landscape based on existing scenery
  - Text: Generative AI can be used to write news articles, poetry, and even scripts. It can also be used to translate text from one language to another
  - Audio: Generative AI can generate new music tracks, sound effects, and even voice acting

The Generative AI Tech Stack

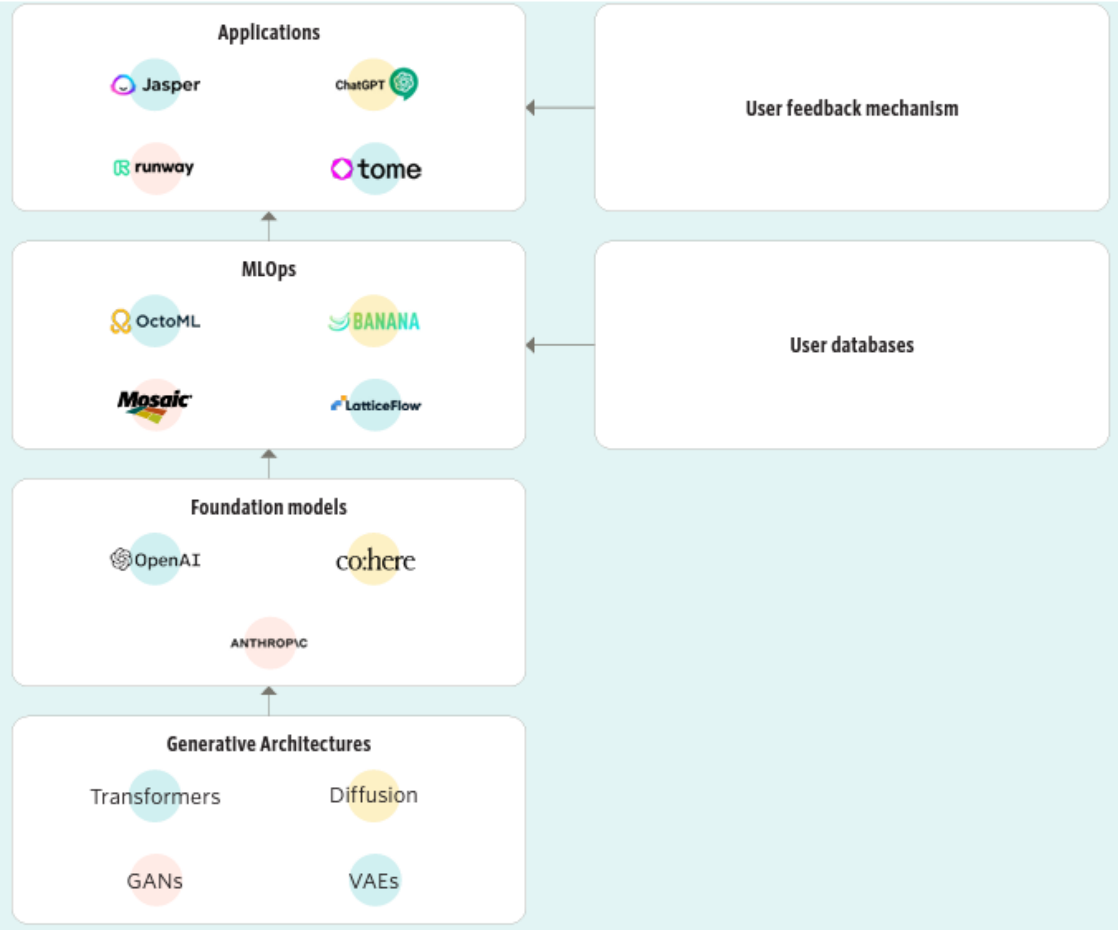

Source: World Economic Forum and Pitchbook.

Industry Drivers

1. **The cost of foundation model training has been decreasing**

   — While GPT-3 was costly to train (estimates placed a single training run at $10 million), newer solutions including Stable Diffusion have brought the cost for state-of-the-art models down to ~$600k

   — Continued advancements in hardware (e.g., Nvidia's H100 Tensor Core GPU) will continue to decrease the cost of compute

2. **The cost to use foundation models by startups remains accessible**

   — The cost of a typical response from an API (such as those provided by OpenAI, Cohere and Midjourney) are priced competitively to enable startups to launch applications on top of their software

3. **CIOs are pushing for AI adoption across their business departments**

   — 60% of CIOs plan for AI to gain widespread adoption across departments by 2025

4. **AI-enabled productivity increases can help to solve labor shortages and cut costs**

   — AI-powered tools, such as chatbots, assistants and copilots, can automate time-consuming tasks, freeing employees to focus on higher-level work and increasing productivity per worker

5. **AI advancements have increasingly shifted from research to industry**

   — AI researchers have steadily moved from academia to industry as opportunities have continued to grow in the private sector, leading to more companies focused on practical AI innovation

6. **Increasing virality of consumer-facing products**

   — As evidenced by ChatGPT, which reached 100 million MAUs faster than any technology within its first two months, consumer-facing AI products have the potential to go viral quickly upon launch

Source: Pitchbook.

VC Activity in Generative AI

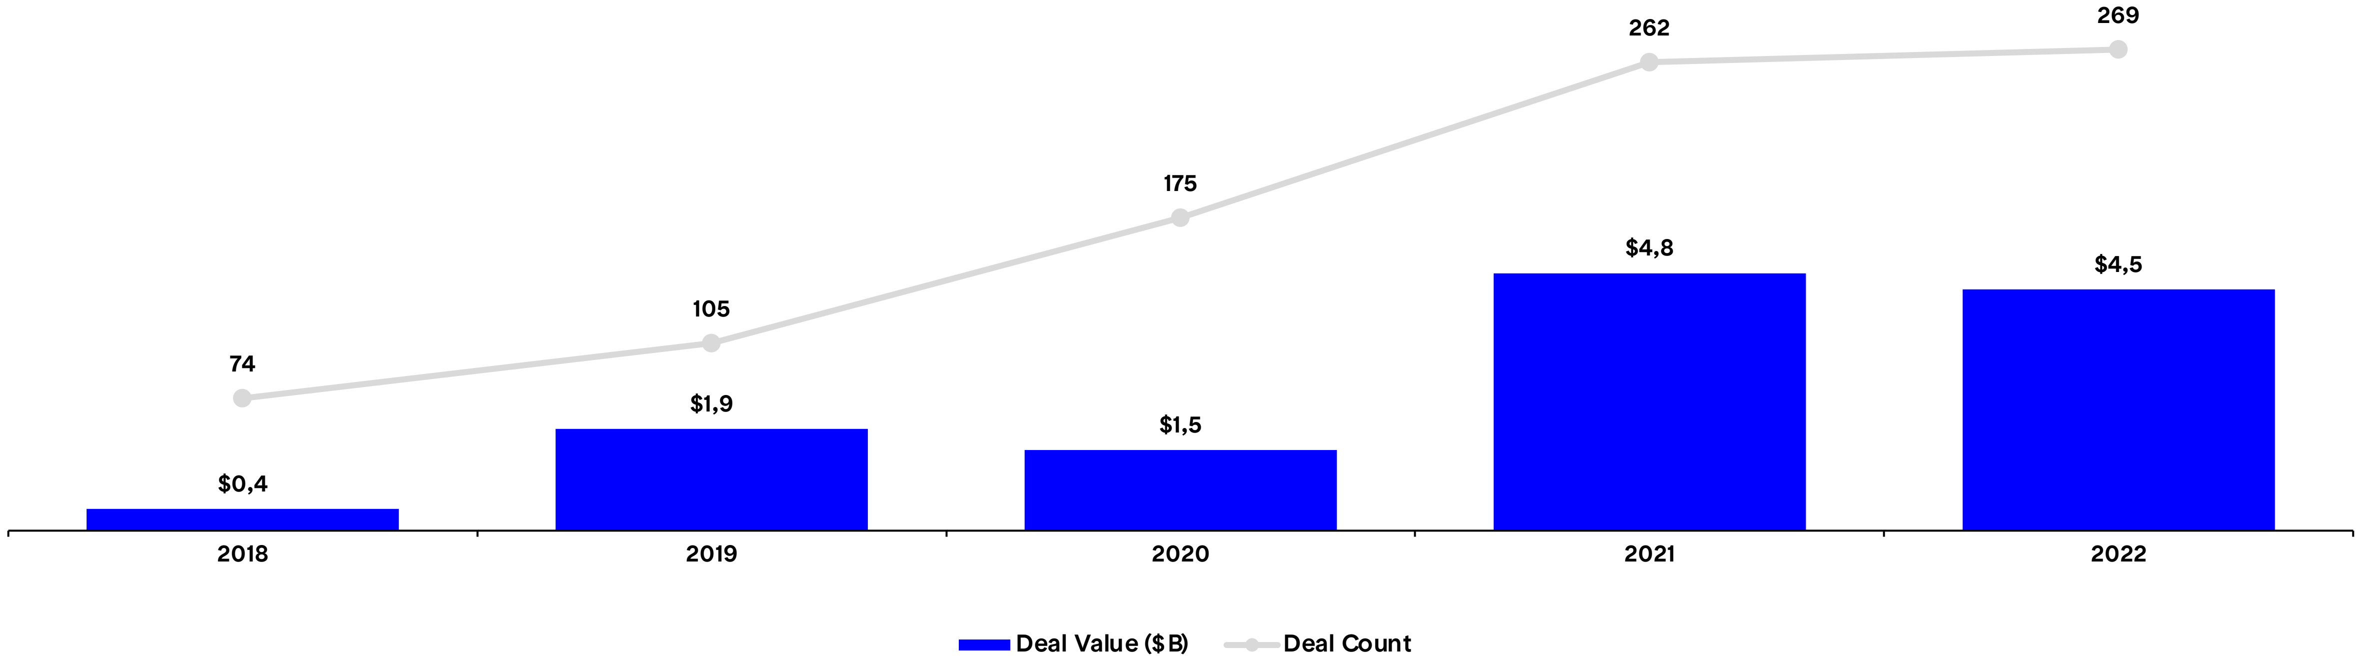

Top Generative AI VC Investors (since 2018)					
Investor	Deal Count	Seed	Early Stage VC	Late Stage VC	Venture Growth
a16z	21	5	9	5	2
Tiger Global	20	1	9	8	2
Sequoia	20	5	12	3	--
Amplify Partners	17	5	9	3	--
Khosla Ventures	17	5	6	6	--
Bloomberg Beta	15	4	8	3	--
Index Ventures	15	4	9	1	1
Soma Capital	15	11	3	1	--
South Park Commons	14	5	6	3	--
Alumni Ventures	14	7	5	2	--

Source: Pitchbook and other publicly available information as of 2022.

Generative AI Tech Stack

Apps
Models
Infrastructure

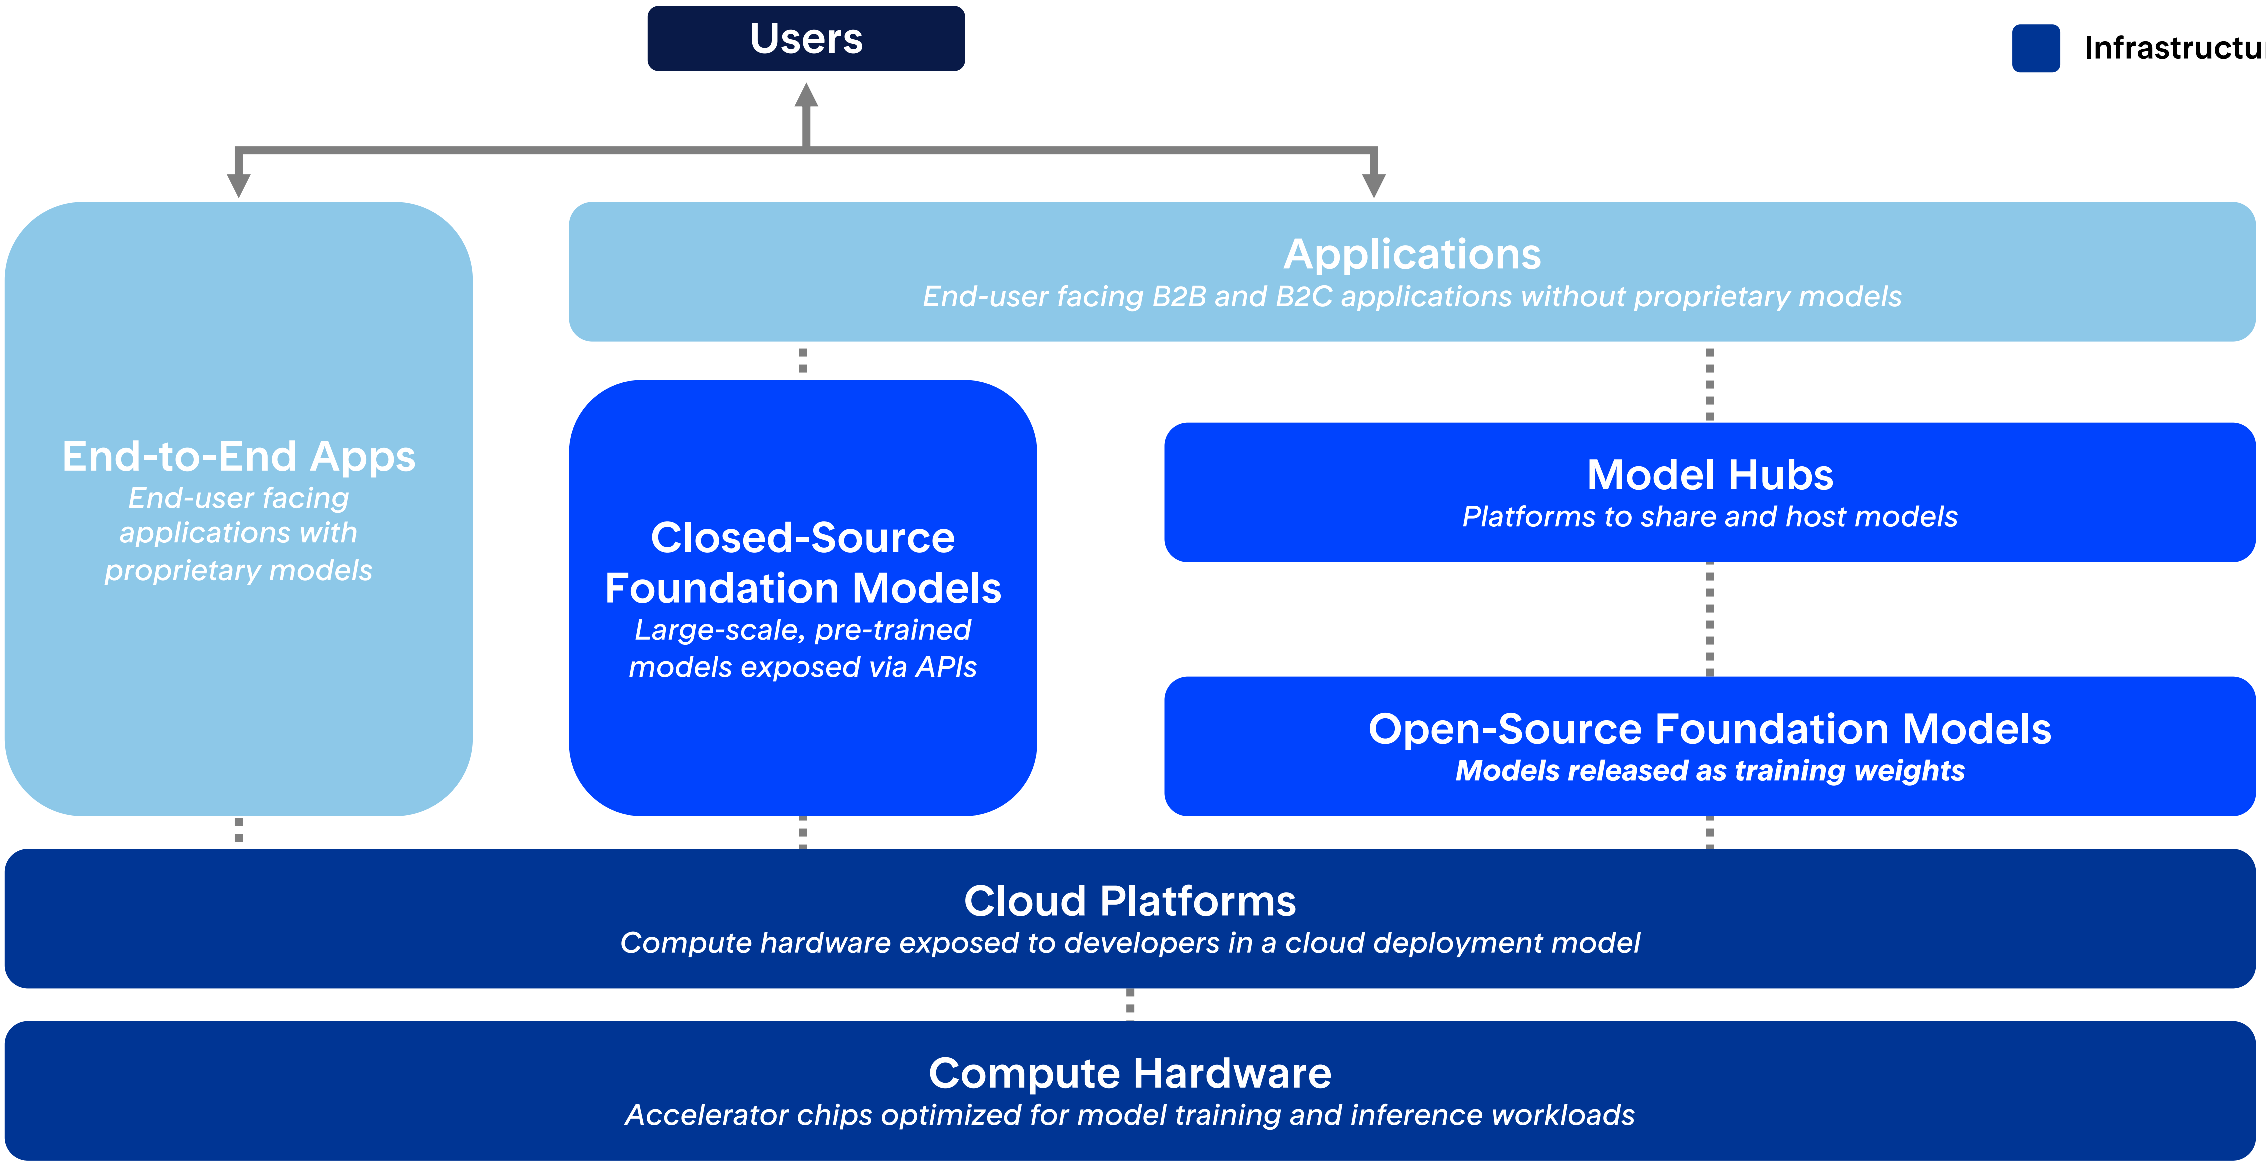

Select Generative AI Applications & Use Cases

2D Media	Code	Vertical Applications
Avatars	Generation	Health
Content Suite	Documentation	Legal
Images	Web App Builders	Finance
Product Design	**Audio**	HR
Video	TTS / STT	**Horizontal Applications**
3D Media	Music	Cybersecurity
Object Synthesis	Editing	Sales & Marketing
Product Synthesis	Voice Dubbing	Customer Experience
Space Synthesis	Summarization	Search
Avatars & NPCs	Translation	Workflow Productivity

Key Generative AI Partnerships with Cloud Hyperscalers

Google & Anthropic

In February 2023, Google invested $400 million in Anthropic, becoming its preferred cloud provider through its Google Cloud Platform

Partnership Overview:

- Anthropic will leverage Google's custom-developed machine learning systems designed to run computationally intensive workloads
- Google Cloud intends to build large-scale, next-generation TPU and GPU clusters that Anthropic plans to use to train and deploy its AI systems

Microsoft & OpenAI

In January 2023, Microsoft invested an additional $10 billion in OpenAI (OpenAI and Microsoft began their partnership in July 2019)

Partnership Overview:

- Microsoft will deploy OpenAI's models across its consumer and enterprise products, introducing new categories built on OpenAI's technology
- Azure will power all OpenAI workloads across research, products and API services
- Microsoft will increase its investments in development and deployment of supercomputing systems to accelerate OpenAI's research

AWS & Stability AI

In November 2022, Stability AI announced that it had selected AWS as its preferred cloud provider to build and scale its AI models

Partnership Overview:

- Stability AI uses Amazon's end-to-end machine learning service, SageMaker, as well as AWS' compute infrastructure and storage to accelerate its generative AI models
- Stability AI will also make its open-source models available on Amazon SageMaker JumpStart, the model hub of Amazon SageMaker, for all AWS customers

Exclusive access to best-in-class compute is critical for businesses building the infrastructure for LLMs

Source: Pitchbook and other publicly available information.

Appendix: Multiples Sensitivity by Year

Illustrative AI Investment Sensitivity – Entry Revenue Multiples (2019 investments only)

Series A

($ in mm)		ARR				
		$1.0	$1.5	$2.0	$2.5	$3.0
Pre-Money Valuation	$17	17 x	12 x	9 x	7 x	6 x
	$19	19 x	13 x	10 x	8 x	6 x
	$22	22 x	14 x	11 x	9 x	7 x
	$24	24 x	16 x	12 x	10 x	8 x
	$26	26 x	17 x	13 x	10 x	9 x

Series B

($ in mm)		ARR				
		$2.0	$4.0	$6.0	$8.0	$10.0
Pre-Money Valuation	$52	26 x	13 x	9 x	6 x	5 x
	$58	29 x	15 x	10 x	7 x	6 x
	$65	32 x	16 x	11 x	8 x	6 x
	$71	36 x	18 x	12 x	9 x	7 x
	$78	39 x	19 x	13 x	10 x	8 x

Series C

($ in mm)		ARR				
		$8.0	$11.0	$14.0	$17.0	$20.0
Pre-Money Valuation	$289	36 x	26 x	21 x	17 x	14 x
	$325	41 x	30 x	23 x	19 x	16 x
	$361	45 x	33 x	26 x	21 x	18 x
	$397	50 x	36 x	28 x	23 x	20 x
	$433	54 x	39 x	31 x	25 x	22 x

Series D

($ in mm)		ARR				
		$20.0	$25.0	$30.0	$35.0	$40.0
Pre-Money Valuation	$530	27 x	21 x	18 x	15 x	13 x
	$596	30 x	24 x	20 x	17 x	15 x
	$663	33 x	27 x	22 x	19 x	17 x
	$729	36 x	29 x	24 x	21 x	18 x
	$795	40 x	32 x	27 x	23 x	20 x

Note: Median pre-money valuation based on median values from the analysis as shown on page 10 (page 10 shows the median values over the entire evaluation period). The additional pre-money valuations are based on a (20%) to 20% range, in 10% increments. ARR ranges are illustrative to show an implied multiple at each range.

Illustrative AI Investment Sensitivity – Entry Revenue Multiples (2020 investments only)

Series A

($ in mm)	ARR				
Pre-Money Valuation	$1.0	$1.5	$2.0	$2.5	$3.0
$23	23 x	15 x	11 x	9 x	8 x
$26	26 x	17 x	13 x	10 x	9 x
$29	29 x	19 x	14 x	11 x	10 x
$32	32 x	21 x	16 x	13 x	11 x
$34	34 x	23 x	17 x	14 x	11 x

Series B

($ in mm)	ARR				
Pre-Money Valuation	$2.0	$4.0	$6.0	$8.0	$10.0
$65	32 x	16 x	11 x	8 x	6 x
$73	36 x	18 x	12 x	9 x	7 x
$81	40 x	20 x	13 x	10 x	8 x
$89	44 x	22 x	15 x	11 x	9 x
$97	49 x	24 x	16 x	12 x	10 x

Series C

($ in mm)	ARR				
Pre-Money Valuation	$8.0	$11.0	$14.0	$17.0	$20.0
$224	28 x	20 x	16 x	13 x	11 x
$252	32 x	23 x	18 x	15 x	13 x
$280	35 x	25 x	20 x	16 x	14 x
$308	39 x	28 x	22 x	18 x	15 x
$336	42 x	31 x	24 x	20 x	17 x

Series D

($ in mm)	ARR				
Pre-Money Valuation	$20.0	$25.0	$30.0	$35.0	$40.0
$1,060	53 x	42 x	35 x	30 x	27 x
$1,193	60 x	48 x	40 x	34 x	30 x
$1,325	66 x	53 x	44 x	38 x	33 x
$1,458	73 x	58 x	49 x	42 x	36 x
$1,590	80 x	64 x	53 x	45 x	40 x

Note: Median pre-money valuation based on median values from the analysis as shown on page 10 (page 10 shows the median values over the entire evaluation period). The additional pre-money valuations are based on a (20%) to 20% range, in 10% increments. ARR ranges are illustrative to show an implied multiple at each range.

Illustrative AI Investment Sensitivity – Entry Revenue Multiples (2021 investments only)

Series A

($ in mm)		ARR				
		$1.0	$1.5	$2.0	$2.5	$3.0
Pre-Money Valuation	$29	29 x	19 x	14 x	12 x	10 x
	$32	32 x	22 x	16 x	13 x	11 x
	$36	36 x	24 x	18 x	14 x	12 x
	$40	40 x	26 x	20 x	16 x	13 x
	$43	43 x	29 x	22 x	17 x	14 x

Series B

($ in mm)		ARR				
		$2.0	$4.0	$6.0	$8.0	$10.0
Pre-Money Valuation	$116	58 x	29 x	19 x	15 x	12 x
	$131	65 x	33 x	22 x	16 x	13 x
	$145	73 x	36 x	24 x	18 x	15 x
	$160	80 x	40 x	27 x	20 x	16 x
	$174	87 x	44 x	29 x	22 x	17 x

Series C

($ in mm)		ARR				
		$8.0	$11.0	$14.0	$17.0	$20.0
Pre-Money Valuation	$460	58 x	42 x	33 x	27 x	23 x
	$518	65 x	47 x	37 x	30 x	26 x
	$575	72 x	52 x	41 x	34 x	29 x
	$633	79 x	58 x	45 x	37 x	32 x
	$690	86 x	63 x	49 x	41 x	35 x

Series D

($ in mm)		ARR				
		$20.0	$25.0	$30.0	$35.0	$40.0
Pre-Money Valuation	$800	40 x	32 x	27 x	23 x	20 x
	$900	45 x	36 x	30 x	26 x	23 x
	$1,000	50 x	40 x	33 x	29 x	25 x
	$1,100	55 x	44 x	37 x	31 x	28 x
	$1,200	60 x	48 x	40 x	34 x	30 x

Note: Median pre-money valuation based on median values from the analysis as shown on page 10 (page 10 shows the median values over the entire evaluation period). The additional pre-money valuations are based on a (20%) to 20% range, in 10% increments. ARR ranges are illustrative to show an implied multiple at each range.

Illustrative AI Investment Sensitivity – Entry Revenue Multiples (2022 investments only)

Series A

($ in mm)		ARR				
		$1.0	$1.5	$2.0	$2.5	$3.0
Pre-Money Valuation	$32	32 x	21 x	16 x	13 x	11 x
	$36	36 x	24 x	18 x	14 x	12 x
	$40	40 x	27 x	20 x	16 x	13 x
	$44	44 x	29 x	22 x	18 x	15 x
	$48	48 x	32 x	24 x	19 x	16 x

Series B

($ in mm)		ARR				
		$2.0	$4.0	$6.0	$8.0	$10.0
Pre-Money Valuation	$118	59 x	30 x	20 x	15 x	12 x
	$133	66 x	33 x	22 x	17 x	13 x
	$148	74 x	37 x	25 x	18 x	15 x
	$162	81 x	41 x	27 x	20 x	16 x
	$177	89 x	44 x	30 x	22 x	18 x

Series C

($ in mm)		ARR				
		$8.0	$11.0	$14.0	$17.0	$20.0
Pre-Money Valuation	$290	36 x	26 x	21 x	17 x	15 x
	$327	41 x	30 x	23 x	19 x	16 x
	$363	45 x	33 x	26 x	21 x	18 x
	$399	50 x	36 x	29 x	23 x	20 x
	$436	54 x	40 x	31 x	26 x	22 x

Series D

($ in mm)		ARR				
		$20.0	$25.0	$30.0	$35.0	$40.0
Pre-Money Valuation	$1,008	50 x	40 x	34 x	29 x	25 x
	$1,134	57 x	45 x	38 x	32 x	28 x
	$1,260	63 x	50 x	42 x	36 x	32 x
	$1,386	69 x	55 x	46 x	40 x	35 x
	$1,512	76 x	60 x	50 x	43 x	38 x

Note: Median pre-money valuation based on median values from the analysis as shown on page 10 (page 10 shows the median values over the entire evaluation period). The additional pre-money valuations are based on a (20%) to 20% range, in 10% increments. ARR ranges are illustrative to show an implied multiple at each range.

AXA Venture Partners

www.axavp.com