



Accelerated Computing is the Path Forward

Al is Software that Writes Software





Data Center is the New Unit of Computing





Al-on-5G Kickstarts the 4th Industrial Revolution

Autonomous Systems in Real and Virtual Worlds





Omniverse Isaac



NEW NVIDIA TECHNOLOGIES

Megatron Drug Discovery Quantum Computing

Jarvis Merlin Maxine Morpheus NVIDIA AI

DGX Grace BlueField DOCA

EGX 5G





NVIDIA IS A COMPUTING PLATFORM





GIVING SCIENTISTS A TIME MACHINE





SCIENTISTS AT GTC



Yoshua Bengio University of Montreal Quebec Al Institute



Yann LeCun Facebook New York University



Geoffrey Hinton University of Toronto Google Vector Institute





Jürgen Schmidhuber Dalle Molle Institute for Al Research



Raquel Urtasun University of Toronto



Alvy Ray Smith Pixar Altamira





Kim Libreri Epic Games



Rommie Amaro University of California, San Diego



Soumith Chintala Facebook



FOR THE DA VINCIS OF OUR TIME





NVIDIA OMNIVERSE



AI Path-Tracing USD Materials Physics









NVIDIA OMNIVERSE



















NVIDIA OMNIVERSE



















COLLABORATING AND SIMULATING IN OMNIVERSE





10M DESIGNERS | 20M CREATORS | 1M SCIENTISTS | 2M DEVELOPERS | 40M ENGINEERS

Image Courtesy of Industrial Light & Magic. © and TM Lucasfilm Ltd. All Rights Reserved.











NVIDIA ISAAC DIGITAL TWIN IN OMNIVERSE











CONNECTING AND CREATING WITH OMNIVERSE































BLUEFIELD SECURES AND ACCELERATES GEFORCE NOW CLOUD GAMING Isolate and Secure Infrastructure | High Quality-of-Service | More Concurrent Users



<100ms







BLUEFIELD SECURES AND ACCELERATES GEFORCE NOW CLOUD GAMING Isolate and Secure Infrastructure | High Quality-of-Service | More Concurrent Users







BLUEFIELD SECURES AND ACCELERATES GEFORCE NOW CLOUD GAMING Isolate and Secure Infrastructure | High Quality-of-Service | More Concurrent Users





ANNOUNCING DOCA 1.0 AND BLUEFIELD PARTNER ECOSYSTEM

Data Center on a Chip Architecture

Software Development Framework for BlueField DPUs

Offload, Accelerate, and Isolate Infrastructure Processing

Support for Hyperscale, Enterprise, Supercomputing and Hyperconverged Infrastructure

Transparent Offload and Acceleration of VMware ESX

Software Compatibility for Generations of BlueField DPUs



ORCHESTRATION

MANAGEMENT

TELEMETRY



ACCELERATION LIBRARIES





ANNOUNCING NVIDIA BLUEFIELD-3 400 Gbps Data Center Infra Processor

Offloads and Accelerates Data Center Infrastructure

Isolates Application from Control and Management Plane

Powerful CPU – 16x Arm A78 Cores

Process Networking, Storage, and Security at 400 Gbps

22 Billion Transistors

DATA PATH ACCELERATOR



DDR5 MEMORY INTERFACE



EXPONENTIAL GROWTH IN DATA CENTER INFRASTRUCTURE PROCESSING Cloud-Native | Disaggregation | Micro-Services | AI | Zero-Trust Security



2020



BlueField-3 22 Billion Transistors 42 SPECint 1.5 TOPS 400 Gbps

BlueField-2 7 Billion Transistors 9 SPECint 0.7 TOPS 200 Gbps

DOCA — ONE ARCHITECTURE

2022



BlueField-4 64 Billion Transistors 160 SPECint 1000 TOPS 800 Gbps



NVIDIA DGX The Computer of AI Researchers



Top U.S. Banks



ANNOUNCING NVIDIA DGX STATION 320G Workgroup Al Supercomputer-in-a-Box

Plug-into-the-Wall Instant AI Infrastructure

2.5 petaFLOPS

320 GB at 8 TB/sec

7.68 TB NVMe

28 MIGs

1500W and < 37db

\$149,000 or \$9,000/Month Subscription







ANNOUNCING THE NEW DGX SUPERPOD World's First Cloud-Native Supercomputer | Secured by NVIDIA BlueField | Multi-Tenant Bare-Metal Performance





NVIDIA MEGATRON TRAINS TRANSFORMERS





Input sequence length=128 tokens (average of 102 words), Output sequence length=8 tokens (average of 6 words) GPU: Megatron GPT-3 on DGX-A100-80GB, Batch size=16, FP16, FasterTransformer 4.0, Triton 2.6 CPU: OpenAI GPT-3 on Xeon Platinum 8280 2S, 755GB System memory, Batch size=1, FP32, TensorFlow 2.3

NVIDIA MEGATRON TRITON



ANNOUNCING NAVER ADOPTS DGX SUPERPOD TO CREATE LANGUAGE UNDERSTANDING **AI SERVICES**

#1 Internet Tech Platform in Korea

Construct AI R&D Belt Across Asia and Europe

Develop Giant AI Language Models

DGX SuperPOD to Build Global-Scale AI Supercomputing Infrastructure









MegaMolBART



ATAC-Seq



AlphaFold1



GatorTron™



NVIDIA CLARA DISCOVERY





ANNOUNCING OXFORD NANOPORE TECHNOLOGIES ADOPTS NVIDIA DGX


ANNOUNCING SCHRODINGER AND NVIDIA ACCELERATE DRUG DISCOVERY \$1.25 Trillion Dollar Industry | 10+ Years Development | 3,000 Pharma Companies

Accelerating Schrodinger's Physics and Machine Learning-Based Computational Drug Discovery Tools with NVIDIA Clara Discovery Acceleration Libraries on NVIDIA DGX A100

Achieve Experimentally-Accurate Simulations



SCHRÖDINGER®

ANNOUNCING RECURSION BUILDS PHARMA AI SUPERCOMPUTER WITH NVIDIA DGX SUPERPOD BioHive-1 Aims to Decode Biology and Industrialize Drug Discovery

Recursion OS Built on NVIDIA DGX SuperPOD Generates, Analyzes, and Derives Insight from Biological and Chemical Datasets

Generate up to 9 Million Images, or Approximately 80 Terabytes of Data, Across up to 1.5 Million Experiments per Week







QUANTUM COMPUTING **RESEARCH ECOSYSTEM** Major Academic and Non-Profit/Governmental Quantum Computing Research Institutions Jniversities Harvard Quantum Initiative Caltech OXFORD 1958 JAR Physics Berkeley universität innsbruck LAB jqi Joint Suantum Institute Ludwig Maximilian Universi University of Science and of-Munich Technology - China University of Maryland Non-Profit/Governmental Research Institute 🜌 Fraunhofer leti UK MCOST Munich Center for Quantum Science and Technology Materials Resear NQIT BROOKHAVEN NATIONAL LABORATORY 北京量子信息科学研究院 ijing Academy of Quantum Information Science Gemini Center on Quantum Computing DECECOR The Institute of Photonic The **EQUS** Fermilak Los Alamos Mational LABORATORY Berkeley Centre for Quantum Technologies Australia ingapore School of Physical Science National Institute of Standards and Technology RRI Raman Research Institute Korea Institute of Science and Technology THE KIST, Making New History Source: IDC 2021 Quantum Computing Initiatives at Universities and Research Institutes Note: This is not an exhaustive listing Quantum Computing Inst Oak Ridge National Laborate

GLOBAL QUANTUM COMPUTING RACE



ANNOUNCING NVIDIA CUQUANTUM Research the Computer of Tomorrow on the Most Powerful Computer Today



Footnotes: State Vector- 1,000 circuits, 36 qubits depth =10, complex 64 | CPU: Qiskit (IBM) on Dual AMD EPYC 7742 | GPU: Qgate (NVAITC) on DGX-A100 Tensor Network - 53 qubits, depth 20 | CPU: Estimated Quimb (Caltech) on Dual AMD EPYC 7742 | GPU: Quimb (Caltech) on DGX-A100

E VECTOR SIMULA Scales to 10's of Qubits	ATION 5	TENSOR NETWORK Scales to 1000's c
Irs -	10 Days	4 Days ←
	Dual-CPU DO	SX A100

Using the Cotengra/Quimb packages, NVIDIA's new cuQuantum SDK, and the Selene supercomputer, we've generated a sample of the Sycamore quantum circuit at depth=20 in record time (less than 10 minutes). This sets the benchmark for quantum circuit simulation performance and will help advance the field of quantum computing by improving our ability to verify the behavior of quantum circuits."

> —Johnnie Gray, Research Scientist, Caltech Garnet Chan, Bren Professor of Chemistry, Caltech





DIVERSE DATA CENTER ARCHITECTURES



ENTERPRISE COMPUTING



HYPERSCALE



STORAGE SERVERS





ACCELERATED HYPERSCALE







SCIENTIFIC COMPUTING

ACCELERATED HPC

SYSTEM BANDWIDTH

DATA-COMPUTE DEMAND **GROWING FASTER THAN** GPU Starved by CPU Memory and PCIE Bandwidth -

GPU	8,000	GB/sec
CPU	200	GB/sec
PCIE Gen 4	16	GB/sec
Mem-to-GPU	64	GB/sec

HBM2e



DDR4

SYSTEM BANDWIDTH

DATA-COMPUTE DEMAND **GROWING FASTER THAN** GPU Starved by CPU Memory and PCIE Bandwidth

GPU	8,000	GB/sec
CPU	200	GB/sec
PCIE Gen 4	16	GB/sec
Mem-to-GPU	64	GB/sec

HBM2e



DDR4

A NEW COMPUTING ARCHITECTURE FOR AI AND DATA SCIENCE

30X Increase System Memory to GPU

GPU	8,000	GB/sec
CPU	500	GB/sec
NVLINK	500	GB/sec
Mem-to-GPU	2,000	GB/sec

30X

LPDDR5x



HBM2e







ANNOUNCING NVIDIA GRACE CPU Designed for Giant-Scale AI and HPC Accelerated Computing

ANNOUNCING THE WORLD'S FASTEST SUPERCOMPUTER FOR AI

20 Exaflops of AI

Powered by NVIDIA Grace CPU and Next Generation NVIDIA GPU

HPC and AI for Scientific and Commercial Apps

Advance Weather, Climate, and Material Science



CSCS Centro Svizzero di Calcolo Scientifico Swiss National Supercomputing Centre











SUPERCOMPUTING COMMUNITY EMBRACES ARM

Alps will use NVIDIA's novel Grace CPU to converge AI technologies and classic supercomputing in one single powerful data center infrastructure.



CSCS Swiss National Supercomputing Centre

-Thomas Schulthess Director CSCS



Thanks to NVIDIA's new Grace CPU, we'll be able to deliver advanced scientific research using high-fidelity 3D simulations and analytics with data sets that are larger than previously possible.

-Thom Mason **Director LANL**







Ampere



DPU



2020

BlueField-2

3 CHIPS. YEARLY LEAPS. ONE ARCHITECTURE.



2021





Ampere Next Next



Grace Next



BlueField-4

2024

2025



EXPANDING ARM IN THE CLOUD

Bringing Together AWS Graviton2 CPU and NVIDIA GPU in 2H21 Android Gaming and AI Inference in the Cloud AWS Graviton2 Delivers Significantly Better Price Performance NVIDIA GPUs Deliver Performance to Scale Streaming Easiest Way to Move Android Gaming to the Cloud

AWS Source of the second secon

EXPANDING ARM ECOSYSTEM BEYOND MOBILE







AI Computing





Cloud



5G Industrial Edge



Robotics







NVIDIA A

ANNOUNCING NVIDIA EGX ENTERPRISE PLATFORM

ANNOUNCING NVIDIA EGX ENTERPRISE PLATFORM

ANNOUNCING NVIDIA AERIAL A100 Al-on-5G

ANNOUNCING GOOGLE CLOUD AND NVIDIA PARTNER TO DELIVER AI-ON-5G

Anthos-Enabled 5G Edge to Deliver Low-Latency, Secure and Mission-Critical Edge AI Applications

Enables the Rapid Delivery of New Services and Applications at the 5G Edge

Provides a Consistent Platform for Application Deployments from Cloud to the Edge

Google Cloud

EVERY USER AND WORKLOAD OF A DATA CENTER IS A SECURITY THREAT

ANNOUNCING NVIDIA MORPHEUS

Security Policies

"who": "32205660 pVycjWHH", "thing": "within there ImXxoVGN", "condition": "Nnpdtcvs 89234632", "key": "nearly appear n1kLSr8A 41142934 actually QZnPjQyd"

INFRASTRUCTURE

CANONICAL

Google Cloud

Red Hat

Mware[®]

INDUSTRIAL EDGE

DEMATIC

everseen

kinetic vision

OSARO

NVIDIA EGX ENTERPRISE ECOSYSTEM

ANNOUNCING NVIDIA EGX ENTERPRISE PLATFORM Enterprise-Ready Suite of AI and Data Science Software

Convolutional Network (CNN)

Autoencoders

Transformer

Long Short-Term Memory (LSTM)

Reinforcement Learning

Generative Adversarial Network (GAN)

Face Vid2Vid Audio-Driven Character Animation

NVIDIA'S AI

NGC PRE-TRAINED MODELS Production-Quality AI Models

Trained by Experts for Enterprise Deployment

Credentials to Find Models You Trust

Continuously Updated to be State-of-the-Art

Adapt with NVIDIA TAO and Orchestrate with NVIDIA Fleet Command

Reference AI Code Samples to Ease Application Development

INTIDIA. NGC | CATALOG

GAZE ESTIMATION

Application

Gaze detection for a person - point of regard (X, Y, Z) and gaze vector (theta and phi).

Popularity

Domain Computer Vision

Usage Unrestricted

License **TLT Licence**

Training Dataset Proprietary dataset with more than 220k images.

Performance T4 (1698 FPS)

Xavier (704 FPS)

NX (510 FPS)

CV Computer Vision DL Deep Learning Gaze TLT

Expand Credentials

Input

The model described in this card detects a person's eye gaze point of regard (X, Y, Z) and gaze vector (theta and phi). The eye gaze vector can also be derived from eye position and eye gaze points of regard.

Model Architecture

GazeNet is a multi-input and multi-branch network. The four input for GazeNet consists: Face crop, left eye crop, right eye crop, and facegrid. Face, left eye, and right eye branch are based on AlexNet as feature extractors. The facegrid branch is based on fully connected layers. Please see the paper in the citations for an example of the model architecture.

Training Algorithm

The training algorithm optimizes the network to minimize the root mean square error between predicted and ground truth point of regards.

How to use this model

Primary use case for this model is to detect eye point of regard and gaze vector. The model can be used to detect eye gaze point of regard by using appropriate video or image decoding and pre-processing. In the TLT Computer Vision Inference Pipeline, gaze estimation network results are used to determine whether the subjects are looking at the camera. See the following image for an illustration of eye gaze estimation usage.

GazeNet is a multi-input network, which takes in face crop image, left eye crop image, right crop image, and facegrid.

Face Image which is gray scale. 224 x 224 x 1

Output

3D point of regards (X, Y, Z) and gaze vector (theta and phi)

ANNOUNCING NVIDIA TAO FRAMEWORK Train | Adapt | Optimize

Customize Pre-Trained Models for **Domain-Specific Applications**

Federated Learning Enables Model Training Collaboration while Protecting Data Privacy

Produce State-of-the-Art Models in Hours

NVIDIA NGC

NVIDIA TAO

ANNOUNCING NVIDIA FLEET COMMAND Securely Orchestrate AI Fleet at the Edge of the Network

Control and Manage Millions of AI-Powered Devices from Any Cloud

Secure from Boot, Attestation, Uplink and Downlink, to Confidential AI Enclave

Centrally Monitor Health and Remotely Fix Edge Systems

NVIDIA NGC

NVIDIA TAO

NVIDIA FLEET COMMAND

NGC PRE-TRAINED MODELS Production-Quality AI Models

Trained by Experts for Enterprise Deployment

Credentials to Find Models You Trust

Continuously Updated to be State-of-the-Art

Adapt with NVIDIA TAO and Orchestrate with NVIDIA Fleet Command

Reference AI Code Samples to Ease Application Development

NGC | CATALOG

ক্ট EMR CXR AI MODEL

Application

Detect likelyhood that a COVID-19 patient in the emergncy room will need supplemental oxygen.

Popularity

Domain Healthcare

Usage For research purposes only

License Clara License

Training Dataset 20 different FL client sites with 16,000 cases.

Performance AUC - 24 hours (94%)

AUC - 72 hours (91%)

Expand Credentials

Input Output

The ultimate goal of this model is to predict the likelihood that a person showing up in the emergency room will need supplemental oxygen, which can aid physicians in determining the appropriate level of care for patients, including ICU placement.

Model Architecture

The model uses a pre-trained ResNet34 [1] for image feature extraction together with a deep & cross network [2] to combine it with EMR features.

Training Algorithm

The models were developed within the EXAM consortium using federated learning. The chest x-ray feature extraction branch of the network was Pretrained on >200,000 images from CheXpert dataset (on pneumonia vs. rest task) & fine-tuned on ~500 images from Mass Gen Brigham to predict **RALE** [3] score to evaluate lung edema on CXR.

NVIDIA JARVIS State-of-the-Art Conversational Al GPU-Accelerated ASR, NLU, TTS Interactive Performance Customize with NVIDIA TAO Orchestrate with NVIDIA Fleet Command Scale-Out with NVIDIA Triton Run in Every Cloud and at the Edge

INTIDIA. NGC | CATALOG

NV
Jarvi
optin
com
The
visio
Fusi
Tush
(Cheven
Get
This
for te
DO
Initia
co
ba

Jarvis Overview Jarvis Documentation Jarvis Developer Forum

vis is a fully accelerated application framework for building multimodal conversational AI services that use an end-to-end deep learning pipeline. velopers at enterprises can easily fine-tune state-of-art-models on their data to achieve a deeper understanding of their specific context and imize for inference to offer end-to-end real-time services that run in less than 300 milliseconds (ms) and delivers 7x higher throughput on GPUs pared with CPUs.

a Jarvis framework includes pre-trained conversational AI models, tools in the NVIDIA AI Toolkit, and optimized end-to-end services for speech, on, and natural language understanding (NLU) tasks.

ing vision, audio, and other sensor inputs simultaneously provides capabilities such as multi-user, multi-context conversations in applications

ng Started

lease of Jarvis includes Quick Start scripts to help you get started with Jarvis AI Services. These scripts are meant for deploying the services locally ing and running the example applications.

ngc registry resource download-version nvidia/jarvis/jarvis_quickstart:1.0.0-b.3	Ð
Initialize and start Jarvis. The initialization step downloads and prepares Docker images and models. The start script launches the server.	
cd jarvis_quickstart_v* bash jarvis_init.sh bash jarvis_start.sh	Ð
Start a container with sample clients for each service.	
bash jarvis_start_client.sh	
Next Steps	Ð

Jarvis transcription

We're also making tremendous progress in translation and now offer support for five languages.

You can see how fluid the translation is into Japanese.

And it's running in real time with under one hundred milliseconds latency for each sentence. す。 す。

また、翻訳も飛躍的に進歩しており、現在5カ国語に対応していま

日本語の翻訳の流動性がわかりま

そして、**各文**で 100 ミリ秒未満の レイテンシでリアルタイムで**実行** されています。

ANNOUNCING NVIDIA MERLIN END-TO-END ACCELERATED **RECOMMENDER SYSTEM**

GPU-Accelerated ETL, Data Loading, Training, Inference

Scales Transparently to Target Data Sets and More Complex Models

Run in Every Cloud and at the Edge

Now Available on NGC

&	
	NV
NVIDIA	dem
MERLIN	optir
Application Accelerated framework for building large-scale deep learning recommenders	
Popularity	•
Domain	
Recommendation Engine	
License Apache License 2.0	
4.0.1-19.11 (Latest) Scan Results	
Linux / x86	
Recommendation Systems Deep Learning	Get
HugeCTR Machine Learning NVTabular	You
Triton Merlin	d
Expand Details 🗸 🗸	/t
	The
	rc
	Activ

I CATALOG

Resources Merlin Overview Merlin Documentation Merlin Developer Forum

DIA Merlin is a framework for accelerating the entire recommender systems pipeline on the GPU: from data ingestion and training to deployment. Merlin powers data scientists, machine learning engineers, and researchers to build high-performing recommenders at scale. Merlin includes tools that nocratize building deep learning recommenders by addressing common ETL, training, and inference challenges. Each stage of the Merlin pipeline is mized to support hundreds of terabytes of data, all accessible through easy-to-use APIs. With Merlin, better predictions than traditional methods and والمرجمين والملايين ومرجم ومحمد والمرجم والمراجع المرجم



ng Started

an pull the training containers with the following command:

ker run --runtime=nvidia --rm -it -p 8888:8888 -p 8797:8787 -p 8796:8786 --ipc=host nvcr.io/nvidia/merlin/merlin-training:0.4 🗌 bash

container will open a shell when the run command execution is completed. It should look similar to this:

oot@02d56ff0738f:/opt/tritonserver#	L.
vate the rapids conda environment by running the following command:	-E
oot@02d56ff0738f:/opt/tritonserver# source activate rapids	
	면



NVIDIA MAXINE

SOTA AI to Reinvent Virtual Collaboration

SDKs for Video, Audio, and Augmented Reality AI-Face Codec 10x Lower Bandwidth vs H.264 Jarvis for Conversational AI Deploy on Client, in Data Center and Every Cloud

Download Today: developer.nvidia.com/maxine

📀 NVIDIA. NGC CATALOG	
NAXINE	NVID NVIDI as vid Maxir
Application GPU-accelerated SDK with state-of-the-art Al features for developers to build virtual collaboration and content creation applications such as video conferencing and live streaming	
Popularity	Y L
Domain Virtual Collaboration License Maxine License	
1.0-21.04 (Latest) Scan Results	
Linux / x86	
Deep Learning Augmented Reality GAN Real-time Audio Real Time Video Al Multimedia Expand Details v	Gett You c doc -p
	The c
	roc
	Activa
	roc

Resources

Maxine Overview **Maxine Documentation** Maxine Developer Forum

NA Maxine[™] is a GPU-accelerated SDK with state-of-the-art AI features for developers to build virtual collaboration and content creation applications such deo conferencing and live streaming.

ine's AI SDKs—Video Effects, Audio Effects, and Augmented Reality (AR)—are highly optimized and include modular features that can be chained into



ing Started

an pull the Maxine container w ith the following command:

cker run -it --name maxine --gpus=all --shm-size=1g --ulimit memlock=-1 --ulimit stack=67108864 -v \${PWD}:/models -p 8000:8000 8001:8001 -p 8002:8002 nvcr.io/nvidia/maxine:1.0-21.04

container will open a shell when the run command execution is completed. It should look similar to this:

oot@02d56ff0738f:/opt/maxine#	
vate the Max ine conda environment by running the following command:	난
oot@02d56ff0738f:/opt/maxine# source activate maxine	
	Ð





that's so much better thanks to NVIDIA Maxine. No es ideal. Ahora vamos a volver a activar todas las características de Maxine.

AI FACE CODEC

EYE CONTACT

MACHINE TRANSLATION SPEECH RECOGNITION





ANNOUNCING NVIDIA TRITON INFERENCE SERVER

MODELSInfinite AI Models
CV, ASR, NLU, TTS, RecSysMultiple Frameworks
TensorFlow, Pytorch,
ONNX, TensorRTVarying Services
Batch, Real-Time,
Streaming

GPU CPU



TRITON INFERENCE SERVER

Any Model Any Framework Any Service Requirement CPU or GPU





1200	6 2 C	Sec. 300	12	men St	12 mg 14	b ₁ , b	1. 1	mar 2	No. Phy	¥ [No.	1 2 3		2		100	1. 2	and the	8 - P 34	14 A	S	11	18 N.	2 14	1000	10 1	1 2 3	2 Sund	-3, %	2 5	100
1. P	4 20	a de		100. 20		3.		24-4	march 1	5. 2	1	1	44	3		- K.	100	197 m	5 m 20	Trie 3th	20 ⁴⁷ 2 1		a	78 10	2 3	S. K.	3 and when	- 99 6	4 1.	50	5
8	Se al	5. 10	4 4	2 3.	and a	6 100	22	and the state	mart	1	1. S. 16	par 2	a have	23	3	in sur	10	1 mar 1		Page 2-4	may way		1 34	2 10	time to the	S. may	no b	12 8	-20 -2	3 3	pre-
2	der.		Nº 2	500 m	6 3 .	2	20 K.	mar &	1 × 1	23	2.5%	10		100	8. J 4	d may	50000	19	1010	1	5 3		5 10	1 set	40 31	Sec. Sec. of	بعثر مسر المج	1 the second	4	1. 18	1.20
2.4	Sand	19 2	and 5	18. 7	Y Lyon 6	and a	phi the	3.00	40 1	1 and	1		60 ⁰⁰ -8,154		1 8 3	20	Section	m 1	199 Mar.				70	1. Sec.	1. 20	2. 8 1	2. 20 m - 20	10 47	18 19	and at	alary.
5. 5		and the	L. A	24. 20	. A. 1	ie. Z	1	Test and	1 12	1	14	197	1 30	-	1 3mm	8 14	~ 1	3	112	17	S	÷ 4.	8	100	4.58	1.0	14 - 14 - A	Sec. 2	100 100	Anne arts	×
Louis C	1 33	1 33	25	199	14.2	Y 14	* Y	A Aug	Sec. any	15 7	18	and 1	Y	4 -	14 190	2 14	90 40	C Quer		1ª m	1 20	8-15-	1	1 440	2.1	Designer 1	a start	the work	\$ 5	the of	6
2	3 34	1 Sec.	19		9. A 3	1.0	- 18 Ba	· 16. Page	1 13	5 3	2. 2	2 3 9	4. AL	1	1. Real	er _ 20	10 %		4 4 84	rest of	2 Care	77	1	1. W	1 11	1 6 1	· · · · · ·	See Se	1 30	Page and	1
1	1	18	5. 3	1	2 1 1	4 h.,	4 2	1 1 mar	and the	146 8	They We	in Dog	Frees	Ser.	5 - Sec. 1	\$ e	85	200	1 2 /	2 2	65	< 3	¥5 - 4	87	1 100	1	you the	and a second	A	No.3	1
140 V	1.	Se al	1 30	12.19	and an an	1	78 - Y	1 40	1 64	40 mg		S pure 3	and P	2	p. ma	10	2 have	and g	* J		5° V2	or 👔	12	XY	12	and a	126 ju	189- 38	a 1 3 4	12. 8	2
N ^{SC} S	13 C.	# 1	1.	143	der and	14	1. 2.	Sec. Sugar	2 -	-4 K.	- way	1 Sec. 1	× 1.	10-	1	1	\$ m	* e** 3	- + ¥	Are 100	100	~ ¥8	5 6	- 100	See Mar	1 - 8 -	1 6 3	See Me	1 2	Jugar	1
1	4 May	A.	W	200 24	×	a foreig	in the	19. 14.	3 3	4.34	Saw 3	17	3 (b.	Ser.	1	1	1. 5	- 1	7 ing 244	- 14	1. D	£ \$	4	3.5	8.00	- 1g a	1 4 m	100 m.	14	1 1	1
ale A	5 A.	in Tre-	1 3	1. 1			20	1 13	Real Property	2.	18 1	1		4	A deres	4 Si	¥**	1/4	÷∕ %	1º at	5 52	6 64		1.00	2.5 20	1 13	der to	- and a first	$\sim <$		20
1	1 Street		ية فمر	12 3		3. L.	des (12)	See 3th	Į X	Series .	Sec. 10	1.1	£ (;	here	1 8	\sim	4	<u>k</u> 1	e bul big	~5°~~20	191	4 - S	÷ 8.	12 3	Sec. Part	- Very		100		1. 3	1
	deland.	1 south	and in	port an	0° - 0 -	-	2 5	1.	1945 - 146 1		23		200	-7	x 24	14 (A	1-4	\$ \$1.04	- X X	151	19 2	ř. %.	See, 7	1	8 28	and a	1 6 40	18 mar	m f	1	
Ş.4	L. Sa	n 62	10 2	100	AX	* <			12.8	- 1ª	1.0	100	Jan	10.0	2 5	le the	* *	3	1. 10	18 Mary	8 2	a. Ca	1 27	toy per			1 m 3	1 4 m	in the	10	3
100	<		10 N	140	1 1	1	19 A. 18	1 10	100 To 10	1	1 -	1.5	> *	1.	5 200 3	~ ~	80			- 21	3 -18	\$ 5.C	1 3	m /	\$ m.	1 any	0-1 34	6 4	1. A.	1910 34	***
	1	all star	W. 12	1 and	754	18	110	100	5	Day 9		15 - A - A	1. 14	S. 1	6 al 1	1	4-88 S.	2.4	y 🚧 👌	1 5	1 X	1 2		100 A	A. Beek	- A.	1 an 11	and the second	14 . 28 m		1
8. 2	p ^{in t} ure	N 64		1	6 10	1 47	and the second	1	1 ×	8-8	Page 2	5-X-5		de al		5 /	20	· · · ·	2 S 😽	1 west		10 at	17	7 2	al and	- 2	y mit w	3 sta	See the	1. de 1900	1
1	210 May	89 ¹⁶ 2	1 22	15 100	1 100	1 1	* *	2. 8	and the	$P_{n_i} = \sum_{j=1}^{n_i} \sum_{j=1}^{n_j} \sum_{j$	Se - 4			1000	- 1		1 6	· /** **	2 S. 5		$\lambda \Sigma$	12	2 No.	No at	Sey &	1 2 4	1 8 2	1 300	18	South 30	3
1	1.11	14 Jan	1 4	137 3	0	1	8 3.	a mit	saids at	The state	and se	1987 -	a della	4	, 13 0 ,	3 20	26.00	12	2 5	wi 48	3 9.	A 42	a 1	~~ ?		and the T	to the second	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	194	1 3	- (
Y.	1 14	alle de	14 E	$ S_{i}^{*} \leq$	"out"		1	3 3	10	1	10	the most of	1 Y	114	Sant 1	2 2		er 1	1 . ×	1. mar	χ_{2}		2. ²⁸ - 24	1 8		200° 34 8	* 7		* *	2 2005	2
24	1 3.4) C3	125-1		3 7	in the	att for	100	1 . 10	pr 3	* 3		(* 44 1	de 7	a same	61	3.1	*	2.0	2:40	3 ×	8		Sec. 1	3.34 . 3m	and the second	a ang a sa	and the	* 3	10-al	de .
2	2 1		5 100	See. 3	1 - Mar	- Agent	£ {	W 4	2 -	1.5 Mar.	S		12 -21	1 3	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	\$ 7"		- 8	1. 1. 1. 1. 1. 1. 1. 1. 1. 1. 1. 1. 1. 1		4 1	32	1	2,8 2.2		5" " " P	1 2 2	Sugar Strate	8 20-0	2.64	1
1	in the		-	-100 4	1 16	an shere	\$. t	100 B	18 and		13.9		- B	<i>b</i> .	a der	\$ 7	1	18	× * ~	7 %	1 56	2		300 3	5 1	193	2 7 BV	3 Back	1 600		Ľ.
4	2 ×	S - 5	2 3	3.2	14.6	5 × 1	5	1 10 A	No. Car		1.65	ha 1.		× 10	131	3 and	3 2	100		1 3	1 200		 aris 	- 4 K	Sec. and				A an	1	
1.	· 8	1	3.50	100 ² 1		1 14	10 mm		1		- 277 - 1	1	4 (°		Super-	1	1	48 4	~~~~				T Co	See /	, atta 1	1 × 1	-	Sec. Br	7 2		
1 A	-		1				1. 1.		83	3 4	and the	1 8 m 3	a	1	¥	A 90	0				12 C.	2-2	fuit of	1	they from		1.12	37	173	her-	
2. 8	2 mg 1	100	10		1 X_ 1	8 T.	with the say	and the second	maple Pools	100		1 18	74	90 A		40 X	1. A	100	8		11	-213	1	\mathbb{C}^{\times}		2 3	and some		ange ange	4.3	4
1	1		3 1	249		-	9. Per	a shore a	7 30	20	100	× 2	(1		1	100				1 4 0	- <u> </u>	He (1800		18		198 9 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	14 J	1.09	1
184.01		34. A		·*****			3. 18	1 5	The short of	13.14	1.20			~?	<u>کا ک</u>	40 2	100	1			1 200		1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	16		- 21 Car	5 m 5 1	1.31		- 4 ² - 104	
1		1 20	11	× .		10 B.	3 23	14 -	26 6	18 20	1.	1.11	11 20	and a	5 mm	11	1 4	X×		2 4 X.	14 Marga		<u>&</u>	YE	19 Ar	an set	2 6 2 2	12	and the second		44
inter 1	1	42			2 6 8	Ser.	25. 20			S	and a	and the second	1 2	* -	and a free	B Mil	100 A	9.4	200		phe they	5 2	14.1	and the second s	26	10		A6- 133	1 3	1 20	
		10 Mar.	11- L.	10 6		a second	A Pro-	20 2	-	1		20		The second		~	6 8	1000	1	14 2			a min	4	100	19 J I	7 1 7			*	156
		2	64	Alter of	1.1	17		384 A.	3.17	Carl S		2 2	<u>,</u>	19 al		2 22	1	o Ma	76 9	5 59	3		3	Sec. 1	Star of	A 494 W	23 3	2 3	Sec.	14 M	÷.,
4.4	10 10	2	1	See.	1 200	2 3	and the	8 3	1 Part	and a	1	Se - 1	5 1			ar 57	2. 3			12.24	6-00 C	and a second	1	2 34	State 1	1 1 1	to and go			10	



THOUSANDS OF COMPANIES DOING COOL THINGS WITH NVIDIA AI





ANNOUNCING NVIDIA AI ENTERPRISE PLATFORM







NVIDIA DRIVE AV

TOPS IS THE NEW HORSEPOWER

250





2022

254 TOPS 25 SPECint



CLUSTER





ANNOUNCING ORIN CENTRAL COMPUTER 1 Chip – 4 Domains

INFOTAINMENT

PASSENGER INTERACTION & MONITORING



ANNOUNCING HYPERION 8 AV PLATFORM

State-of-the-Art Advances for Data Collection, Development and Testing

2x Orin AV Computer

1x Orin IX Computer

4x Orin + 4x MLNX 3D GT Data Recorder

Sensor Suite: 8 Cameras (8MP), 4 Fisheyes (3MP), 3 In-Cabin, 9 Radar, 2 Lidar

Source Access to AV & IX Software Repository

OTA Ready



TOPS IS THE NEW HORSEPOWER







ANNOUNCING NVIDIA DRIVE ATLAN The Next Level – Same Programmable Architecture

Fusing Next Generation AI and BlueField

Industry's First 1,000 TOPS SoC

400 Gbps Networking with Secure Gateway

ASIL-D Safety Island





A NEW BREED OF TECHNOLOGY COMPANIES

























THE WORLD'S BIG BRANDS HAVE GIANT OPPORTUNITIES







TRANSFORMING 10 TRILLION MILES A YEAR INTO A SERVICE









THE INTERNET OF ATOMS





NVIDIA DRIVE DIGITAL TWIN IN OMNIVERSE













ANNOUNCING NVIDIA DRIVE SIM POWERED BY OMNIVERSE







NEW NVIDIA TECHNOLOGIES



Megatron | Drug Discovery Quantum Computing DGX | Grace BlueField-3 | DOCA 1.0



NVIDIA AI EGX Aerial 5G







