

“NVIDIA REINVENTS ITSELF EVERY SINGLE YEAR.
WE ARE GOING TO CALL NVIDIA ‘THE GOAT,’
THAT IS, THE GREATEST OF ALL TIME.”

MAD MONEY

NVIDIA pioneered accelerated computing to tackle challenges ordinary computers cannot. We make computers for the da Vincis and Einsteins of our time so that they can see and create the future.



“NVIDIA PROVIDES THE CORE TECHNOLOGY RESHAPING INDUSTRY AND SOCIETY”

DIGINOMICA



GAMING




HPC



HEALTHCARE



SMART CITY




ROBOTICS




AUTOMOTIVE




vGPU



BASE COMMAND



FLEET COMMAND





NVIDIA AI



OMNIVERSE

Accelerated computing requires more than just powerful chips. We achieve incredible speedups through full-stack invention, from the chips and systems to the algorithms and apps they run.



“NVIDIA NOW ENJOYS A ROBUST
AND SELF-SUSTAINING ECOSYSTEM
OF SOFTWARE, UNIVERSITIES,
STARTUPS, AND PARTNERS”

FORBES

Scientists, researchers, developers, and creators are
using NVIDIA to do amazing things. More than 2.5
million developers and 8,500 startups create thousands
of applications for accelerated computing. We’ve
shipped more than a billion CUDA®-based GPUs.



MILLION
DEVELOPERS

NVIDIA REINVENTS MODERN GRAPHICS

We invented the programmable shading GPUs 20 years ago, defining modern real-time computer graphics.

With NVIDIA RTX™, we have reinvented computer graphics, again. This new rendering approach fuses rasterization and programmable shading with ray tracing and AI to make PC games look much more beautiful and realistic—almost cinematic.



“GAMECHANGER”

DIGITAL FOUNDRY



RTX has come to the world’s best-selling video game: *Minecraft*. Before RTX, game developers painstakingly pre-rendered lighting and shadow effects to make their worlds more realistic. But in user-created virtual worlds like *Minecraft*, only real-time ray tracing can accomplish these beautiful effects.

“WITH ITS SERIES OF RTX GRAPHICS CARDS, NVIDIA IS WOWING THE GAMING WORLD”

TECHRADAR

NVIDIA Ampere architecture-powered devices deliver stunning ray-traced graphics with industry-best performance. RTX is a home run with developers—supported by more than 150 top games and creative apps.

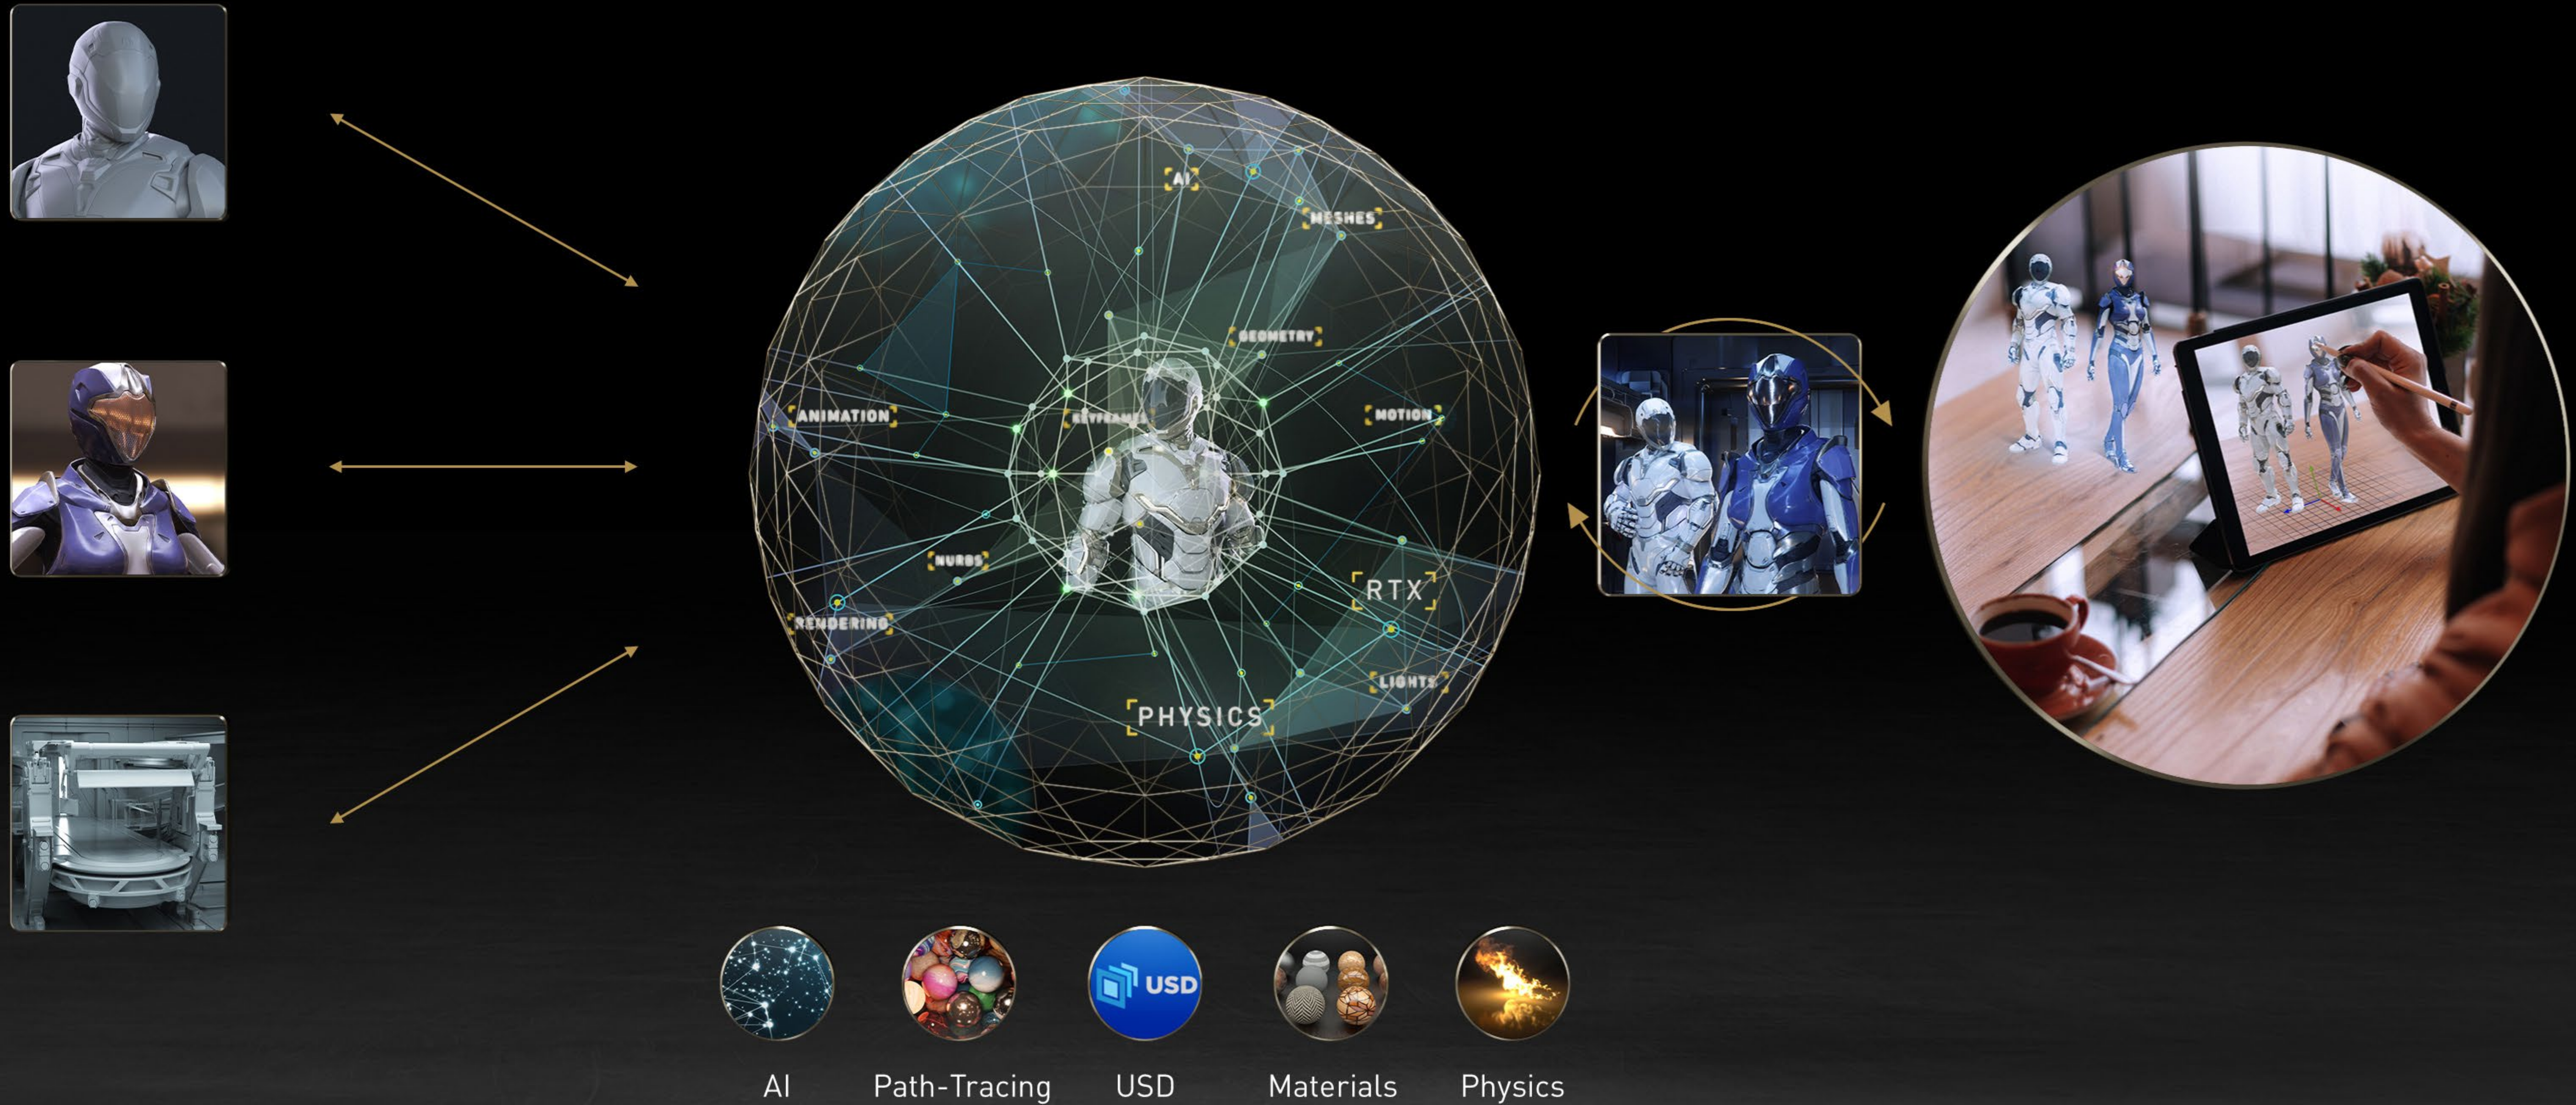


“GEFORCE NOW SHOWCASES THE LIMITLESS POTENTIAL OF CLOUD GAMING”

TECHRADAR

Billions of gamers don't have GeForce® PCs. With our cloud gaming service GeForce NOW™, players can experience the power of a GeForce PC in the cloud, using nearly any device they own. The service hosts 11 million members and is available in 70 countries.





“OMNIVERSE ALLOWS THE ENTIRE PRODUCTION PROCESS TO BE SIMULATED WITH PHOTO-REALISTIC DETAILS, AND WITH PHYSICAL PROPERTIES”

WIRED

NVIDIA Omniverse™ is a platform for simulating and connecting to virtual worlds. In the world of Omniverse, digital content designers can meet virtually to develop complex 3D content in real time. Omniverse obeys the laws of physics. It can simulate particles, fluids, materials, springs, and cables—making it perfect for training robots, designing products, or creating digital twins of buildings, factories, and even cities.

“BMW AND NVIDIA DEVELOP ‘NEXT LEVEL’ VIRTUAL FACTORY PLANNING”

ROBOTICS AND AUTOMATION NEWS

In the future, factories themselves will be robots. BMW is using Omniverse to design and create its own factory of the future. Built in digital and simulated from beginning to end, the factory is a digital twin to the one BMW will build.



FLOOR
PLANNING



ERGONOMICS
TRAINING




ROBOT
ORCHESTRATION



NVIDIA IS POWERING THE NEXT ERA OF COMPUTING

In 2006, the creation of our CUDA programming model and data center GPU platform brought parallel processing to general-purpose computing. A powerful new approach to high performance computing was born.

Today, the universe of supercomputing has expanded rapidly to incorporate AI, advanced data analytics, and cloud computing. The era of the CPU-centered monolithic supercomputer is coming to a close. The next era has begun.



“THE UK’S MOST POWERFUL
SUPERCOMPUTER HAS GONE
LIVE, WITH POTENTIAL TO
TRANSFORM AI-BASED
HEALTHCARE RESEARCH”

ZDNET

With 400 petaflops of AI performance, the Cambridge-1 supercomputer gives healthcare researchers in the U.K. a powerful new tool to take on medicine’s toughest problems. Like NVIDIA’s Selene, the world’s 6th most powerful computer, Cambridge-1 is based on the first turnkey AI supercomputing infrastructure: NVIDIA DGX™ SuperPOD.




“HIGH PERFORMANCE COMPUTING
WORKLOADS GET TO PIGGYBACK
ON ALL OF THIS INNOVATION”

THE NEXT PLATFORM

NVIDIA accelerates more than 700 applications today,
including the top 15 in scientific computing.

By addressing the entire computing stack, we can
drive continuous speed improvements on these
applications even without releasing new GPUs.

We’ve accelerated core HPC applications—used
to discover new drugs, explore the cosmos,
and predict the weather—by 13X in just five years.




“NVIDIA AIMS CLARA HEALTHCARE AT DRUG DISCOVERY, IMAGING”

HPC WIRE

NVIDIA Clara™ Discovery is our suite of acceleration libraries created for computational drug discovery—from imaging, to quantum chemistry, to gene variant-calling, to using Natural Language Processing to understand genetics, to using AI to generate new drug compounds.

Welcome to the digital biology revolution.



NVIDIA IS REARCHITECTING THE DATA CENTER FOR AI

We used to think of a CPU server as the basic unit of computing. But to meet the demands of today's machine learning and AI workloads, we must optimize the entire data center, from end to end.


The new unit of computing is the data center itself.




“NVIDIA IS NOT JUST A GRAPHICS CHIP COMPANY ANYMORE”

MOTLEY FOOL


Next-generation data centers are an orchestration of three pillars: the GPU for accelerated computing, the CPU for general-purpose computing, and the DPU, which processes and moves data in the data center. The introduction of the NVIDIA Grace™ CPU and the NVIDIA BlueField® DPU make NVIDIA a three-chip company, aimed at rearchitecting the data center for AI.



GPU
Ampere A100



CPU
Grace



DPU
BlueField-3

“NVIDIA CONTINUES TO EXPAND ITS LINE OF DGX APPLIANCES AND CLUSTERS FOR AI COMPUTING”

CRN



At the beginning of the big bang of modern AI, we recognized the need to create a new kind of computer for a new way of developing software. This computer would need new chips, new system architecture, new ways to network, new software, and new methodologies and tools. It all comes together as DGX—a computer for AI, supercharged by the NVIDIA A100 GPU.



DGX SUPERPOD
AI Data Center-as-a-Product



DGX A100
AI Data Center Building Block

DGX STATION A100
AI Data Center-in-a-Box

NVIDIA IS BRINGING AI TO INDUSTRY

AI began in research labs and was then adopted by cloud computing providers. Now we stand at the cusp of the next wave of AI adoption: AI automation at enterprises.

The next wave of AI is at the edge, and it will revolutionize the world's largest industries.



“NVIDIA IS HELPING PARTNERS ‘DEMOCRATIZE AI’ FOR ENTERPRISES”


CRN

Companies are integrating 5G, AI, and autonomous machines to create smart services. The opportunity to automate is endless: retail checkout, autonomous forklifts and tractors, pick-and-place robots, and intelligent hospital rooms. Until now, without deep learning AI, no computer software was able to handle the diverse and seemingly infinite conditions in the real world.


METROPOLIS




CLARA



ISAAC



AERIAL



“SPEED UP AI DEVELOPMENT
BY OVER 10X WITH NVIDIA
PRE-TRAINED MODELS”

SMART INDUSTRY

NVIDIA has invested billions in developing AI, but it goes well beyond systems and apps. We've packaged up pre-trained AI models that developers can easily integrate with their own apps. They're production quality, trained by experts, and will continue to improve over time.

TRAFFIC CAM NET

Application

4 class object detection network to detect cars in an image.

Popularity

Domain

Computer Vision

Usage

Unrestricted

License

TLT Licence

Training Dataset

Proprietary dataset from US traffic cameras with more than 3 million objects for car class.

Inference Performance

T4 - 1579 FPS

AGX Xavier - 656 FPS

Xavier NX - 340 FPS

Nano - 17.7 FPS

Computer Vision Deep Learning

TLT Edge Object Detection

Robotics Intelligent Video Analytics

Expand Credentials

TrafficCamNet

Model Overview

The model described in this card detects one or more physical objects from four categories within an image and returns a box around each object, as well as a category label for each object. The four categories of objects detected by this model are – car, persons, road signs and two-wheelers.

Model Architecture

The model is based on NVIDIA DetectNet_v2 detector with ResNet18 as a feature extractor. This architecture, also known as GridBox object detection, uses bounding-box regression on a uniform grid on the input image. Gridbox system divides an input image into a grid which predicts four normalized bounding-box parameters (xc, yc, w, h) and confidence value per output class.

The raw normalized bounding-box and confidence detections needs to be post-processed by a clustering algorithm such as DBSCAN or NMS to produce final bounding-box coordinates and category labels.

Training Algorithm

The training algorithm optimizes the network to minimize the localization and confidence loss for the objects. The training is carried out in two phases. In the first phase, the network is trained with regularization to facilitate pruning. Following the first phase, we prune the network removing channels whose kernel norms are below the pruning threshold. In the second phase the pruned network is retrained. Regularization is not included during the second phase.

How to use this model

Primary use case intended for this model is detecting cars in a color (RGB) image. The model can be used to detect cars from photos and videos by using appropriate video or image decoding and pre-processing. As a secondary use case the model can also be used to detect persons, road signs and two-wheelers from images or videos. However, these additional classes are not the main intended use for this model. This model can be used in smart city applications such as traffic pattern analysis, determining anomalous trajectories, etc.

Input

RGB Image of dimensions: 960 X 544 X 3 (W x H x C) Channel Ordering of the Input: NCHW, where N = Batch Size, C = number of channels (3), H = Height of images (544), W = Width of the images (960) Input scale: 1/255.0 Mean subtraction: None

Output



Category labels and bounding-box coordinates for each detected object in the input image.

Download

“NVIDIA JETSON NANO IS
RESPONSIBLE FOR THE
BIGGEST INDUSTRIAL IOT
REVOLUTION THESE DAYS”

UPSWIFT


NVIDIA Jetson AGX Xavier™ delivers the energy-efficient computational power needed for embedded systems like robots, drones, and smart cities. From the edge to the data center, all of NVIDIA's AI computers run on the same CUDA-X AI™ software stack.



“MERCEDES-BENZ AND NVIDIA SIGN A DEAL TO MAKE CARS MORE LIKE IPHONES”

BLOOMBERG

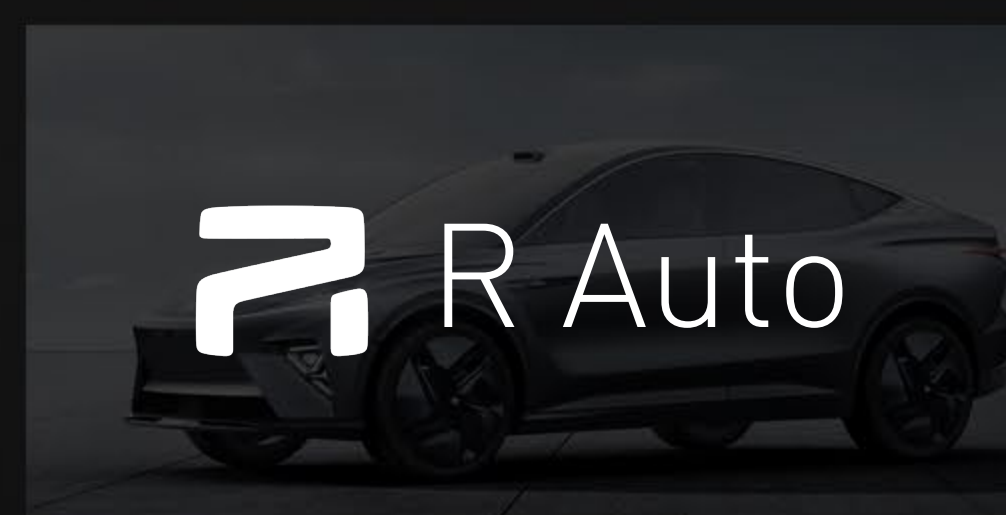
Future cars are going to be completely programmable computers. They will be operated by a supercomputer in the data center that learns from road data and trains the AI networks to drive, and the car's AV computer to perceive and drive—always learning and improving for the life of the fleet. NVIDIA DRIVE™ is our end-to-end platform for automakers to bring these autonomous vehicles to life.



“FROM AUTOMAKERS TO RESEARCH TEAMS AND STARTUPS, ALL DEPEND ON NVIDIA
FOR HARDWARE AND SOFTWARE SOLUTIONS FOR SELF-DRIVING VEHICLES”

ZACK'S

NEXT-GEN EV



MAJOR AUTOMAKERS



ROBOTAXIS




TRUCKING



NVIDIA IS A LEARNING MACHINE

NVIDIA is united by a unique culture—the operating system of our company. We dream big, take risks, and learn from our mistakes together. Speed is the key to our success. Craftsmanship is a passion. There are no org charts—the project is the boss.

These beliefs inform everything we do, from designing amazing products to building one of the world's great companies—a place where people can do their life's work.





“NVIDIA’S CULTURE IS THE PRODUCT OF A FOUNDER CEO WHO EMBRACES COMMUNITY”

FORTUNE

NVIDIA employees are dedicated to building technology that moves humanity forward and supporting the communities in which they work and live.

We’ve been recognized as a top company in social responsibility, and our employees are passionate donors to hundreds of charities around the globe. Employees joined the company in contributing \$16M in 2020 to support COVID-19 response.





“Huang has risen to the elite among Silicon Valley’s visionary leaders. Scores of reports show NVIDIA employees love working for him and his addresses are often technical yet accessible. He commands an audience through his passion for the technology his company is creating.

He’s been at the helm of NVIDIA since co-founding the company at age 30 in 1993 and has led NVIDIA from the maker of computer graphics cards to become the premier platform for artificial intelligence and machine learning. This positions NVIDIA at the forefront as the computing industry contemplates a fundamental shift in processing.

NVIDIA saw it coming.”

TECHCRUNCH

“BEST PLACES TO
WORK IN 2021”

GLASSDOOR

“MOST INNOVATIVE
COMPANIES”

FAST COMPANY

“WORLD’S BEST
PERFORMING CEO”

HARVARD BUSINESS REVIEW

“WORLD’S BEST
CEOS”

BARRON’S

“100 BEST COMPANIES
TO WORK FOR”

FORTUNE

“50 SMARTEST
COMPANIES”

MIT TECH REVIEW

Founded in 1993

Jensen Huang, Founder & CEO

20,000 Employees

\$16.7B in FY21

