# The Next Wave of Edge AI and Robotics

Ettikan Kandasamy Karuppiah (Ph.D), Director/Technologist, ROAP Region
November 2023

# Agenda

- Edge AI and Robotics Industry Outlook and Trend

---

- Unlocking New Applications; GenAI, LLM and Simulation with ISAAC ROS Platform

---

- Use Cases

---

# Edge AI Transforms Nearly All Industries

AI could potentially deliver an additional economic output of around US $13 trillion by 2030

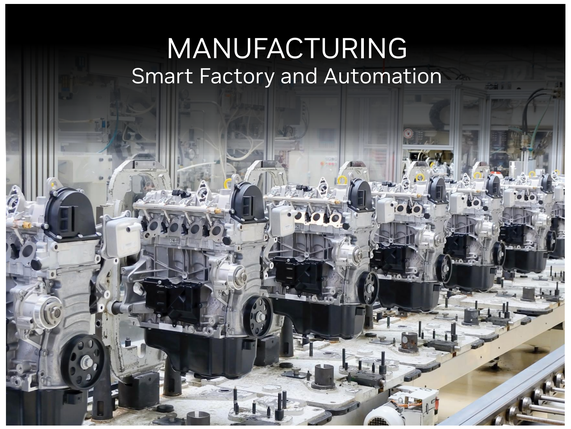


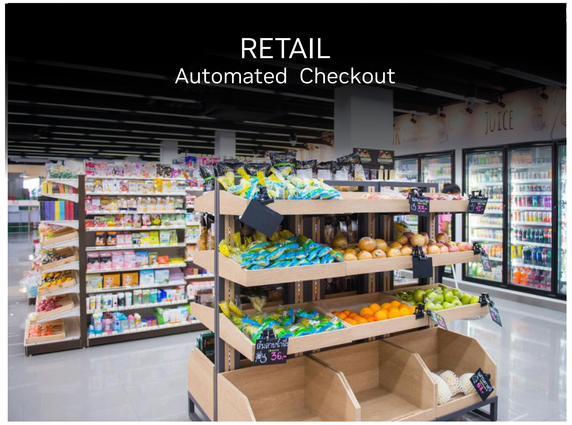



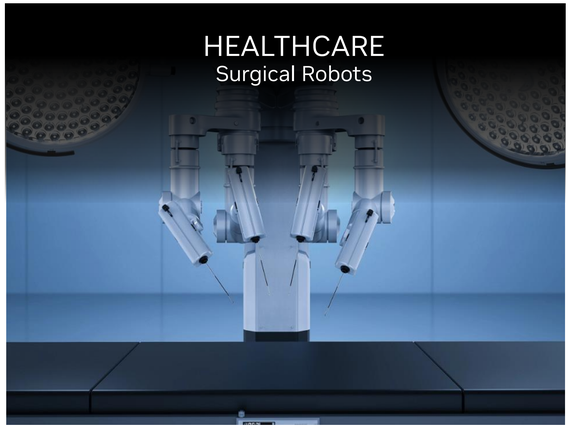





# Why Edge AI and Robotics? Why Now?

Operate with the “intelligence” of human cognition


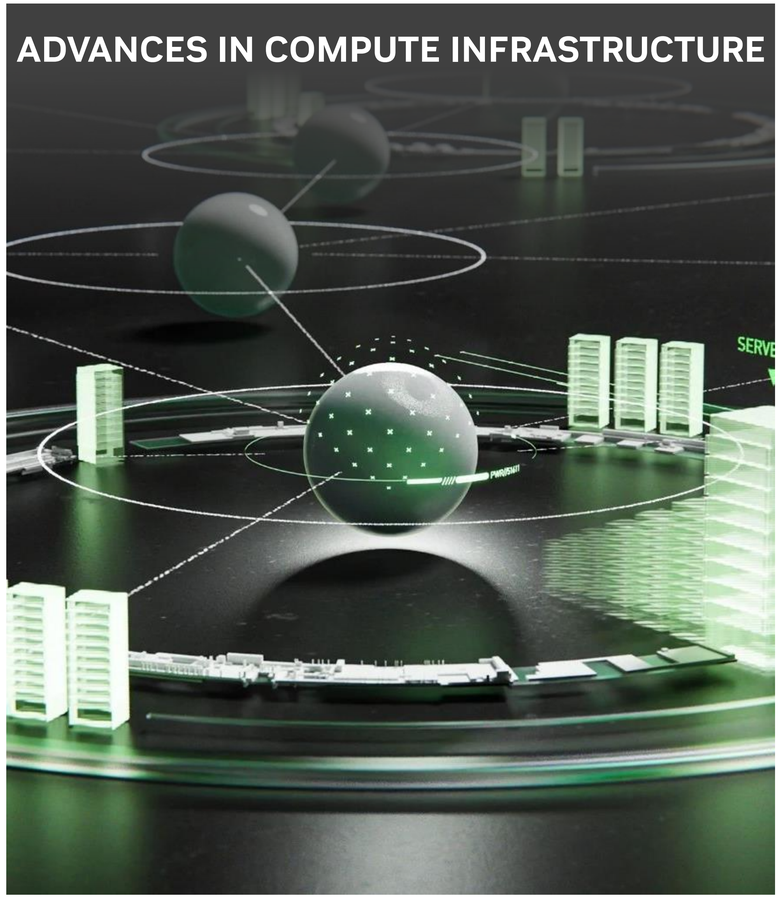


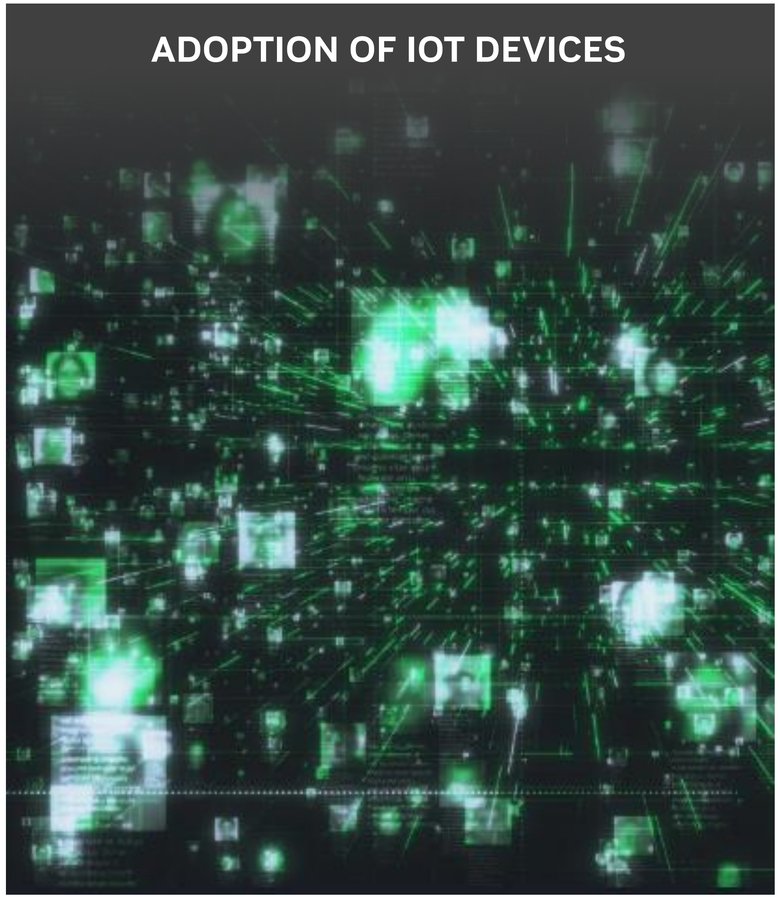

# Macrotrends

More compute is needed for real-time insights and autonomous actions

Smart Cameras

Drones

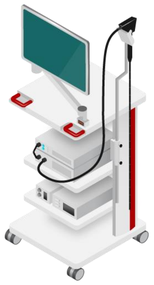

Handheld Medical Instruments

Delivery Robot

Warehouse logistics

AMR /UGV

Single Fixed Function Device

<10 TOPS

40 TOPS

200 TOPS

>275 TOPS

Autonomous Machines

AI Compute Power

# Common Challenges


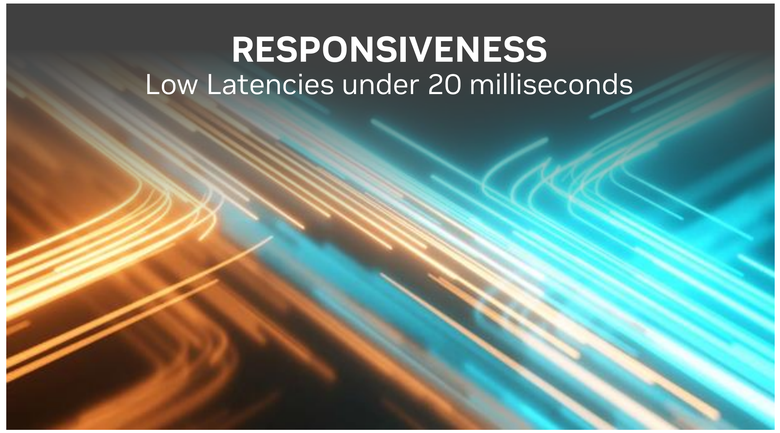




# Giant Leap in Performance for Next Gen AI

## Vision and conversational AI pretrained models

| | Jetson AGX Xavier | Jetson AGX Orin |
|---|---|---|
| PeopleNet (V2.3) | 418 | 1294 |
| Action Recognition 2D | 471 | 1577 |
| Action Recognition 3D | 32 | 105 |
| LPR Net | 1190 | 4118 |
| Dashcam Net | 670 | 1895 |
| BodyPose Net | 172 | 559 |
| ASR: Citrinet 1024 | 19 | 44 |
| NLP: BERT-base | 271 | 780 |
| TTS: Fastpitch-HifiGAN | 36 | 95 |

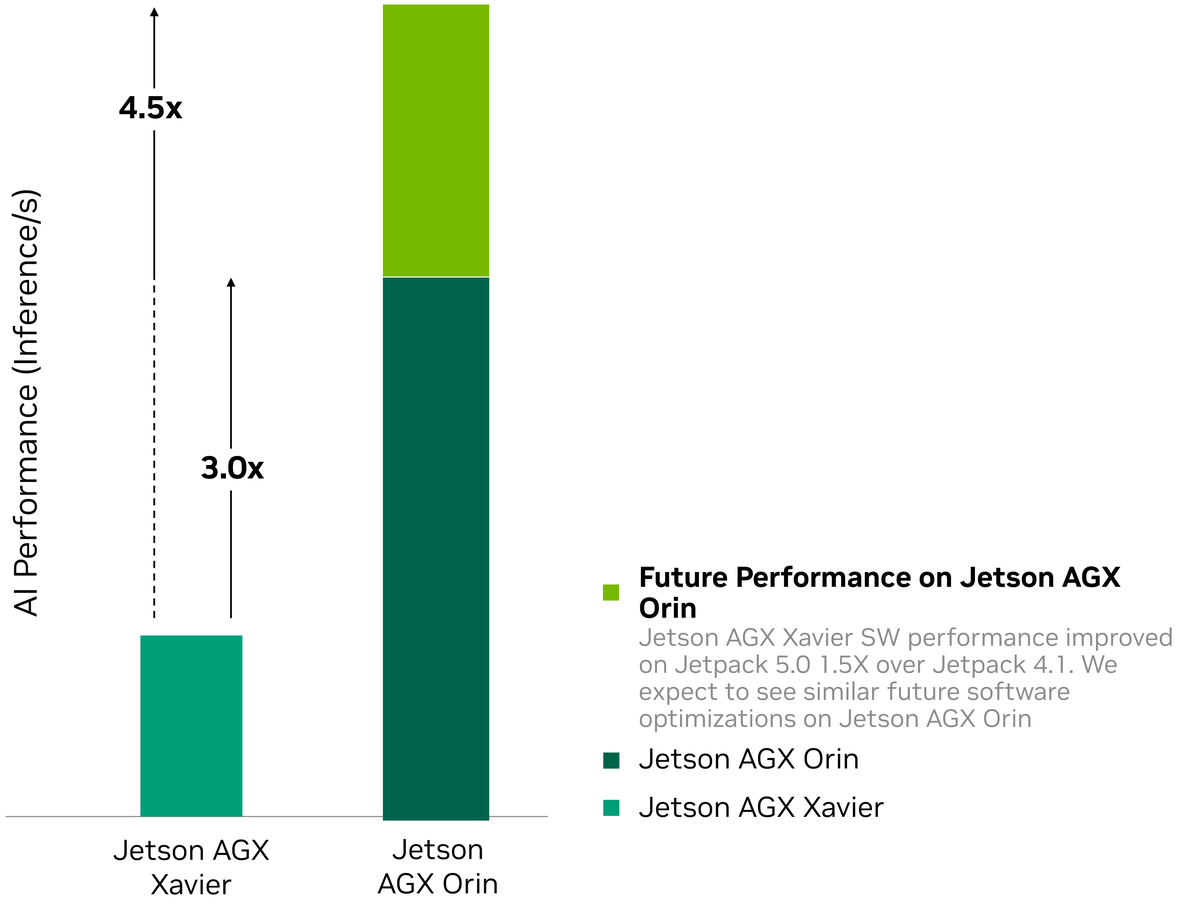


** These are Dense Models | * PeopleNet used here is v2.3 with pruned performance. Previous results used are v2.5 unpruned data.*

NVIDIA

# JETSON AGX ORIN SERIES

## Server Class AI Performance at the Edge

- Up to **275 INT8 TOPS** powered by Ampere GPU +DLA
- Up to 12x A78 ARM CPU
- Up to 64 GB memory, 204 GB/s
- Production Modules are available **now**
- Get started developing today for the entire Jetson Orin family with the Jetson AGX Orin developer kit

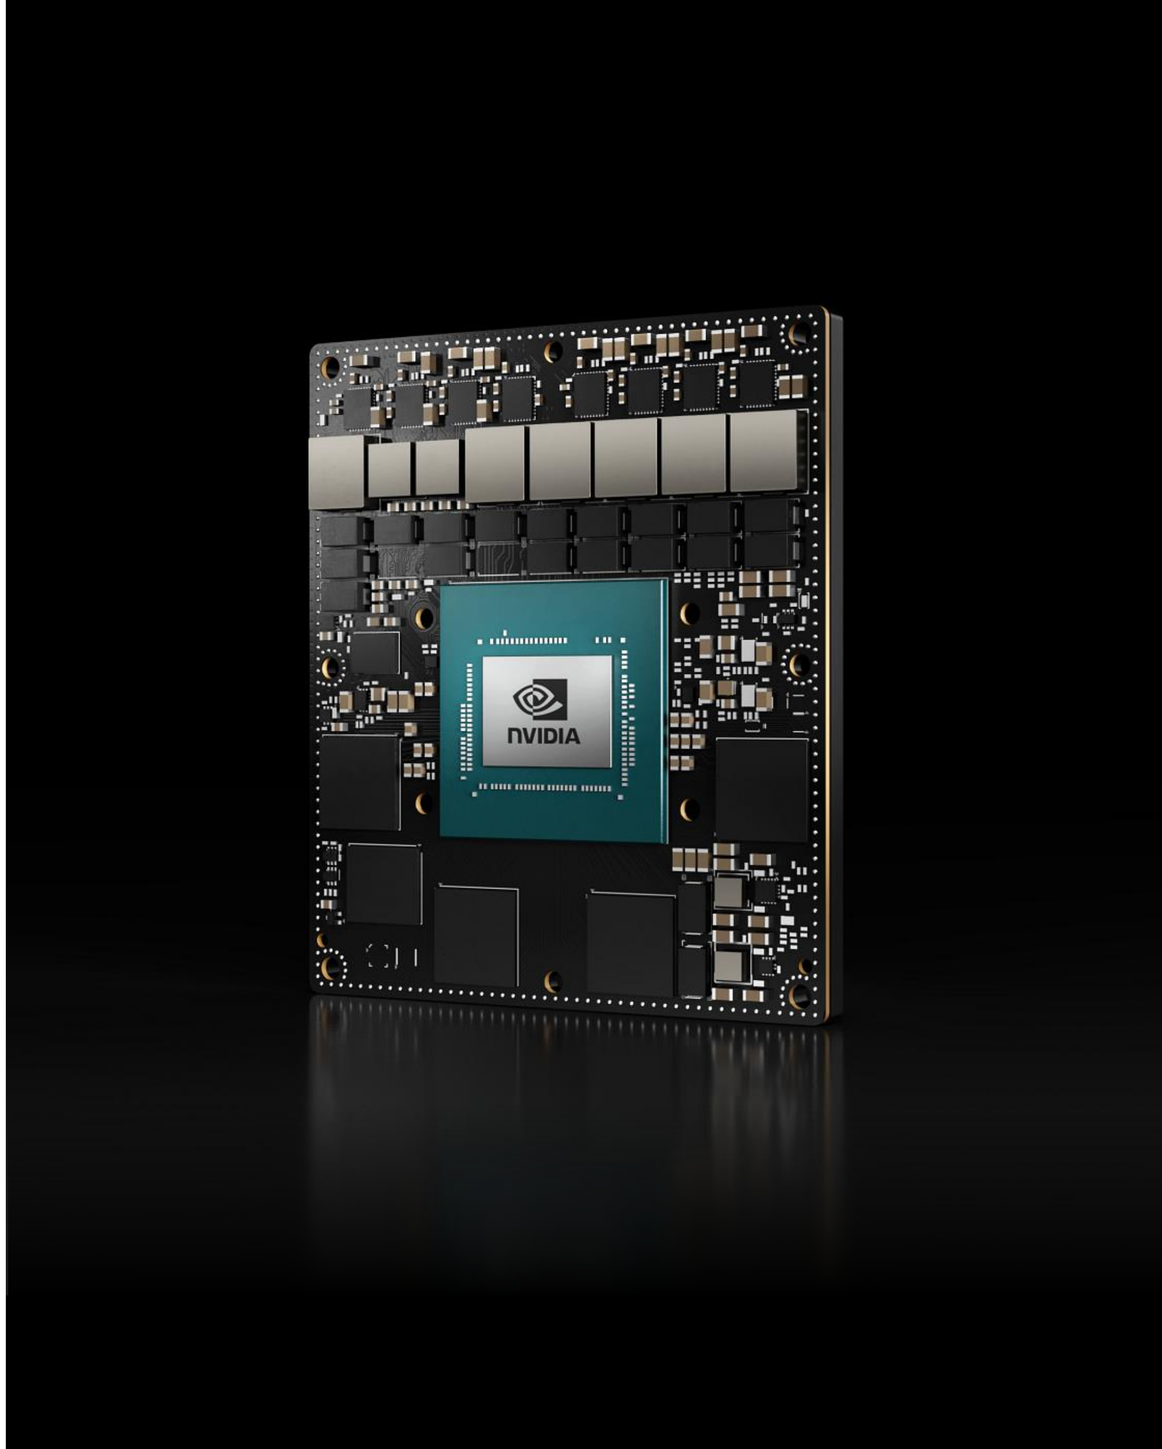

# ISAAC ROS for JETSON Orin

## Increasing Performance on Orin and ROS DP

| | Isaac ROS EA3 Jetson AGX Xavier | Isaac ROS DP Jetson AGX Xavier | Isaac ROS DP Jetson AGX Orin |
|---|---|---|---|
| AprilTag* (720p) | 101fps | 150fps (1.5X) | 260fps (1.7X) |
| DOPE (VGA) | 12.5fps | 12.5fps (1X) | 43fps (3.4X) |
| Image Segmentation* (544p) | 30fps | 208fps (6.9X) | 325fps (1.5X) |
| Proximity Segmentation* (576p) on DLA | N/A | 33fps | 62fps (1.9X) |
| Stereo Disparity* (ESS)(1080p) | N/A | 24fps | 51fps (2.2X) |
| Stereo Disparity (SGM)(540P) | 60fps | 80fps (1.3X) | 166fps (2X) |

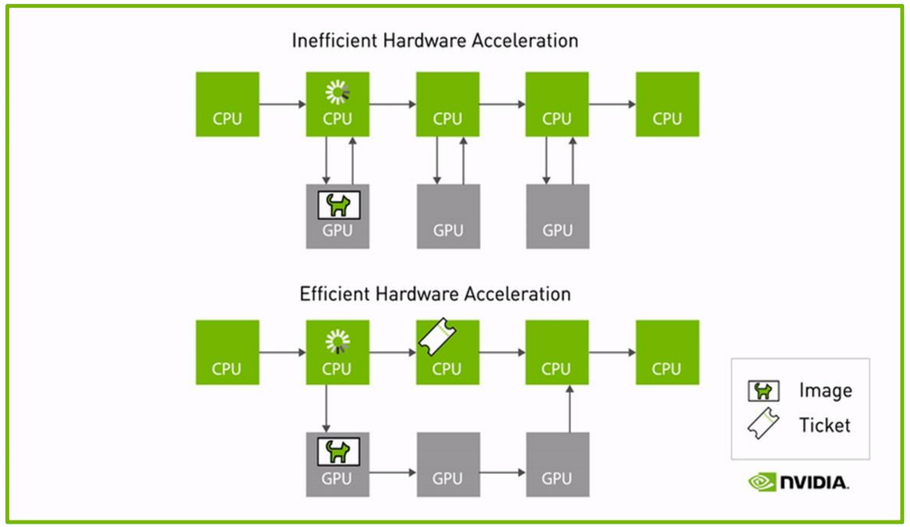


** NITROS accelerated models*
*** The AI performance is calculated by taking the geomean of all tested models*

https://github.com/NVIDIA-ISAAC-ROS

NVIDIA

# Agenda

- **Edge AI and Robotics Industry Outlook and Trend**

---

- **Unlocking New Applications; GenAI, LLM and Simulation with ISAAC ROS Platform**

---

- **Use Cases**

---

# End to End Robotics With NVIDIA Isaac

Focus of this Presentation


Isaac Sim

Topics Covered Today :
Demo + Handson Homework


Isaac Replicator
TAO


Isaac ROS GEMs
Jetson Orin

Topics Covered Today :
Presentation


cuOpt

# NVIDIA ISAAC ROS

# Building Smarter Robots is Challenging

## More Autonomy Driving Compute Requirements

- Real-time planning, perception, and control require more than CPU-only compute. Solutions need to take advantage of modern SOC architecture that have multiple compute engines (GPU, NN Accelerators, DSPs, etc.)
- Fast-moving innovations in AI and computer vision are difficult to keep up with.  Lots of advances in DNNs and CV technologies that are relevant to robotics but are almost impossible to track.
- Many packages available but are not built for production. Testing and improving quality of demoware software costly and time-consuming.

# Isaac ROS – Building Smarter Robots, Faster

## Accelerated Computing for ROS 2 and Jetson

ROS native packages for bringing AI, CV and camera processing to ROS 2 graphs.

***NOTE**: Packages can be modified or bridged to interface to ROS 1 or other frameworks*

Leverages large collection of pre-trained models from NVIDIA repository

Seamless integration with open source/custom ROS tools and packages.

Hardware accelerated libraries and engines enable highest performance and efficient resource utilization.

Optimized to run on all NVIDIA compute platform.

Highest performance edge AI platform available (Up to 275 TOPS with Jetson Orin)

### ISAAC ROS PACKAGES/GEMS

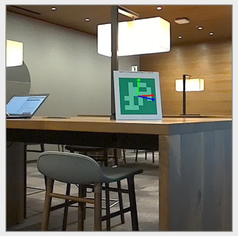
APRILTAG DETECTION


STEREO DEPTH

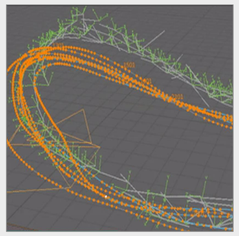
VISUAL SLAM

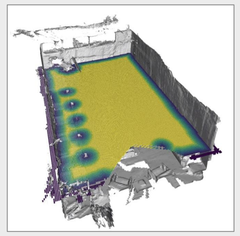
NVBLOX


IMAGE SEGMENTATION DNN

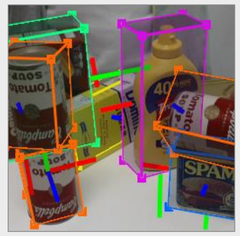
OBJECT DETECTION and 6D POSE DNN

COMMON

TRITON SERVER | NVENGINE | ROS MIDDLEWARE

### HARDWARE ACCELERATED LIBRARIES

TensorRT | VPI | CUDA | ARGUS

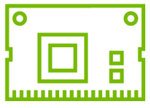
**JETSON COMPUTER**


**GPU WORKSTATION**


**GPU SERVER**

https://github.com/NVIDIA-ISAAC-ROS

NVIDIA

# Isaac ROS Value Proposition

## NVIDIA Delivers High-Performance, Cutting-Edge, AI and CV to Robotics

### Cutting-Edge NVIDIA Technology

Best In Class VSLAM, 3-D Reconstruction, New Depth Perception DNNS, and more


### Power Efficient, High-Performance Graphs

Running the right algorithm on the right HW core yield lower power and preserves CPU MIPs

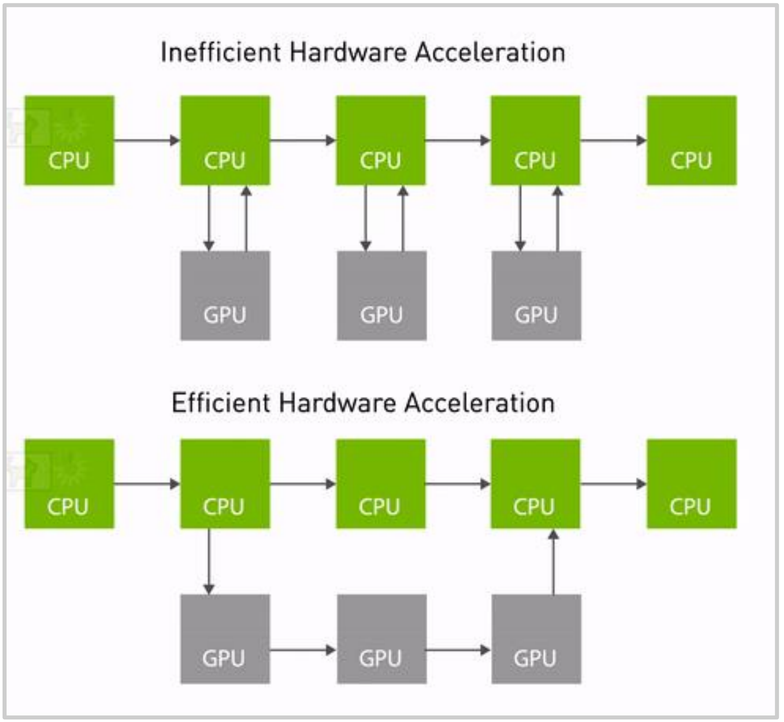


### Time-saving, Ready to Integrate Modules

Qualified and tested on reference robots

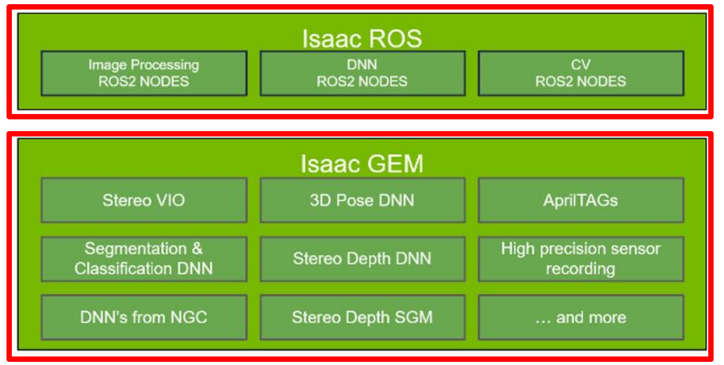


All Isaac ROS packages can be used with ROS2, ROS1, or custom framework.

# Visual SLAM

## Optimizing ROS for accelerated computing

- Vision based method for visually estimating the position of a robot relative to its start position
- High performance low latency multi stereo-camera solution *(250fps | 3.1ms at 720p on Jetson AGX Orin)*
- Best in class visual odometry based on KITTI Visual Odometry metrics *(https://www.cvlibs.net/datasets/kitti/eval_odometry.php)*

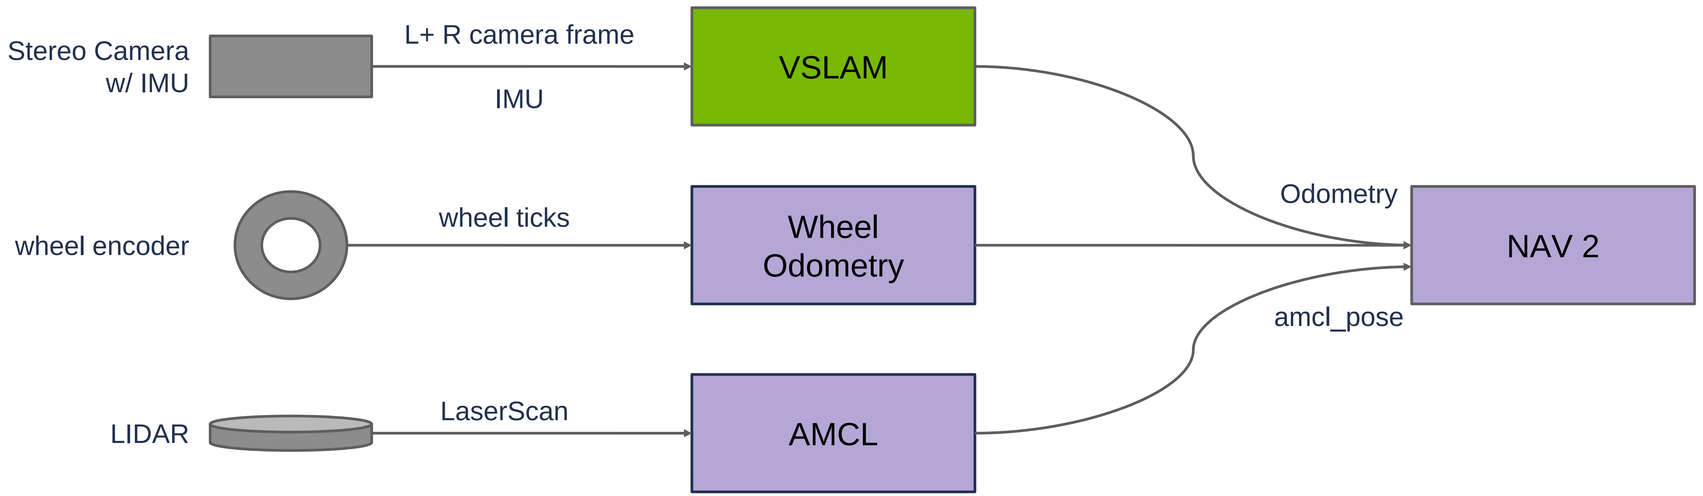

# 3D Scene Reconstruction

## Cost maps for navigation

- Real-time scene 3D reconstruction to create cost maps for navigation
- Vision based obstacle avoidance from a stereo camera
- Plugins to work natively with Nav2 in ROS

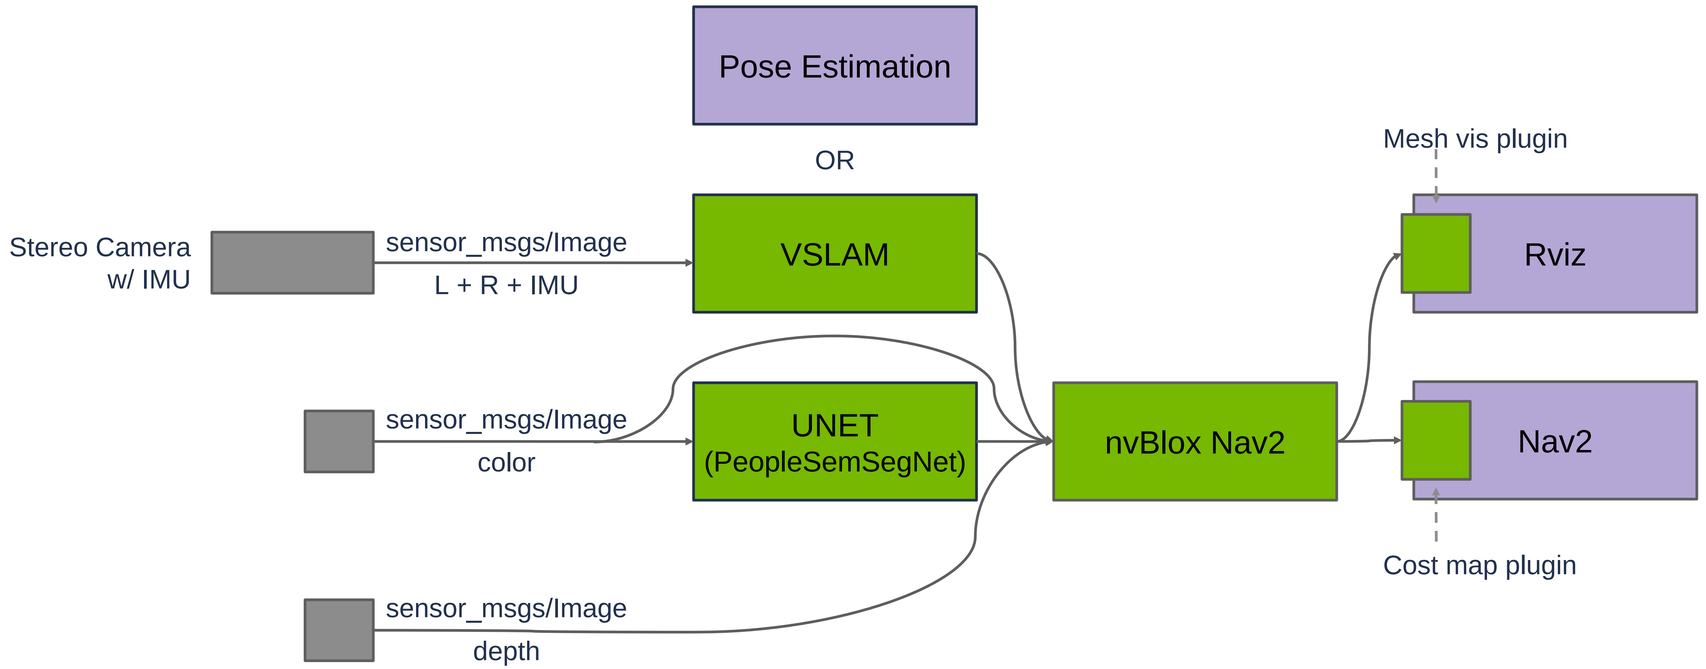

# Map Localization

## Automatically finding where the robot is on initialization

- Automates previously manual process to specify the location of the robot in RVIZ
- Uses LIDAR scans to estimate pose relative to a map
- Designed for compatibility with planar and 3D LIDARs
  Additional packages to process LIDAR format conversion form LaserScan to flatscan

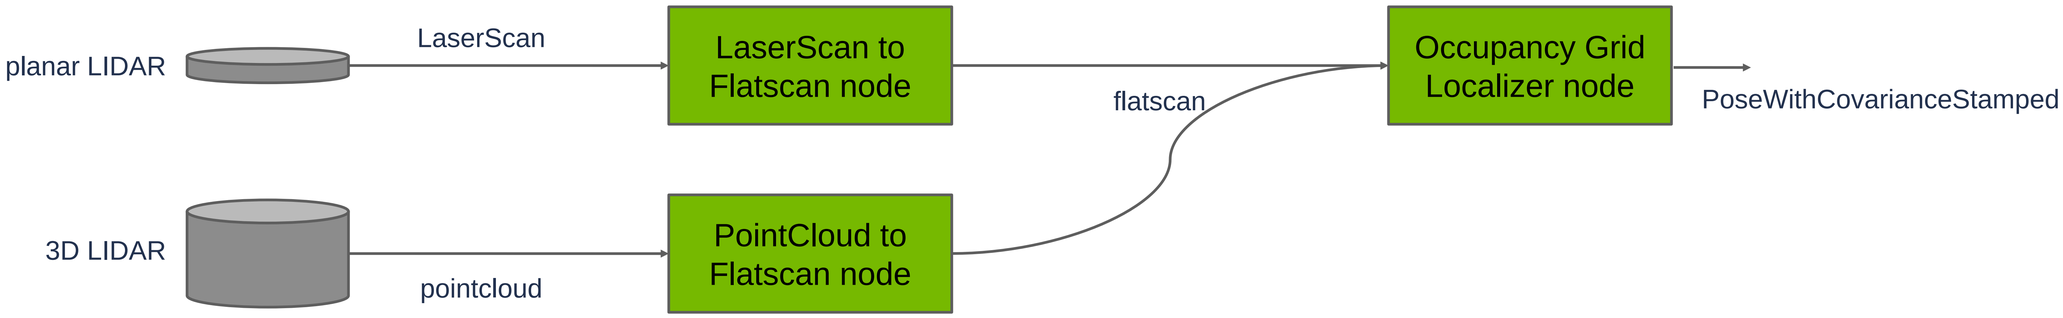

# Fiducials

## AprilTag

- Vision based detection for a point of reference
  - Identifying location in places where LIDAR localization fails due to a lack on unique features
- Calculates orientation and pose relative to the pose of the AprilTag
  - Helps in precise local planning to an object


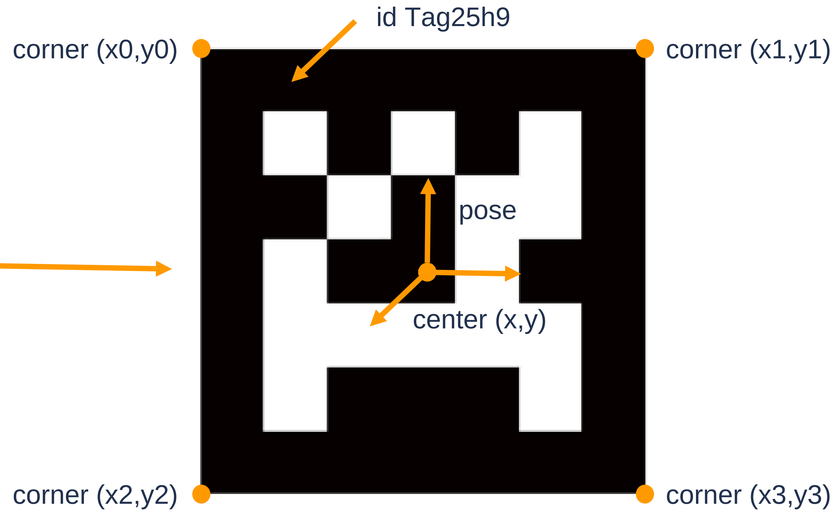

# Depth Perception

## Stereo vision redundancy and diversity

- Vision based depth perception to help create occupancy map of obstacles around the robot

- 3 hardware accelerated perception functions providing diversity and redundancy
  - SGM is a classic stereo matching approach on OFA hardware in Orin
  - ESS is an AI model to predict disparity on GPU hardware
  - Bi3D is an AI model to predict safety fields on DLA hardware

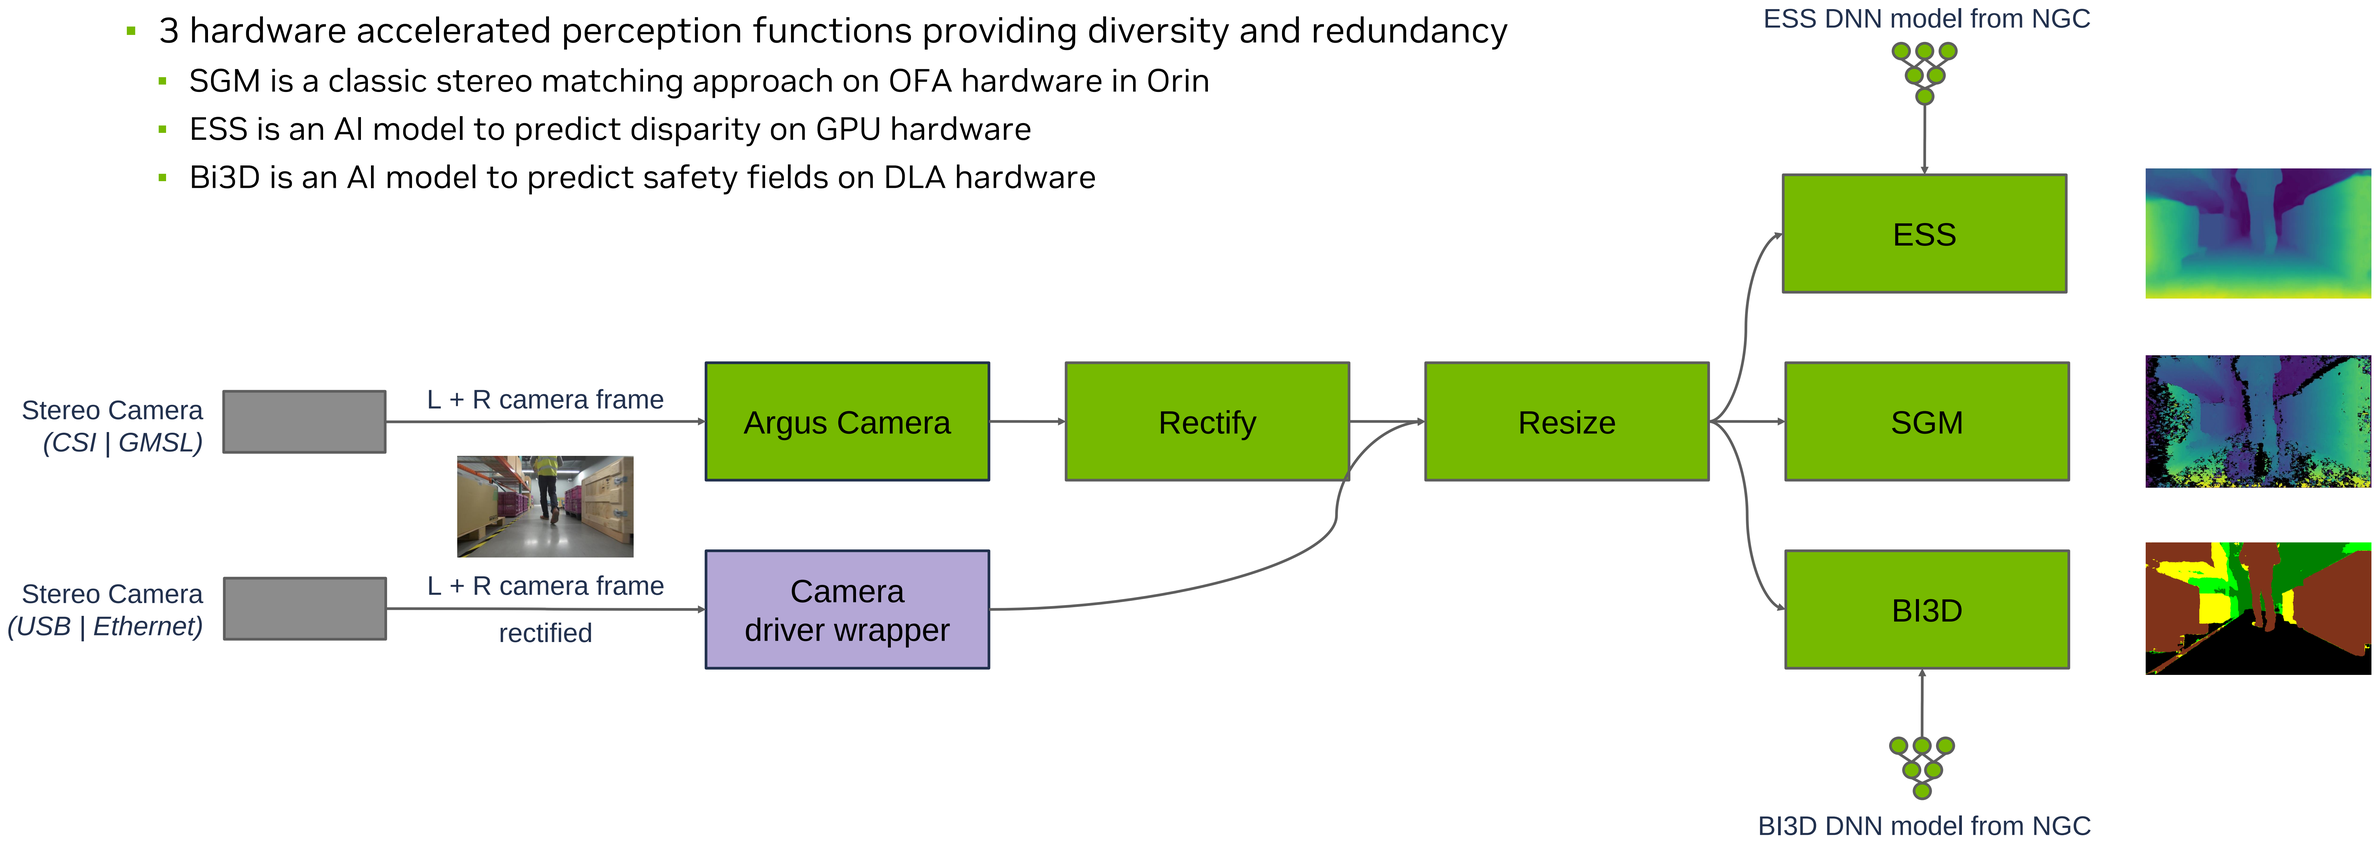

# Freespace Perception

## Diverse vision-based function for obstacle avoidance

- Freespace perception uses vision to detect absence of `obstacles`
  Perception is never perfect, and requires multiple methods to detect obstacles
- Diverse solution to detection of obstacles
- Uses segmentation mask from BI3D and camera extrinsic to compute freespace relative to the ground plane

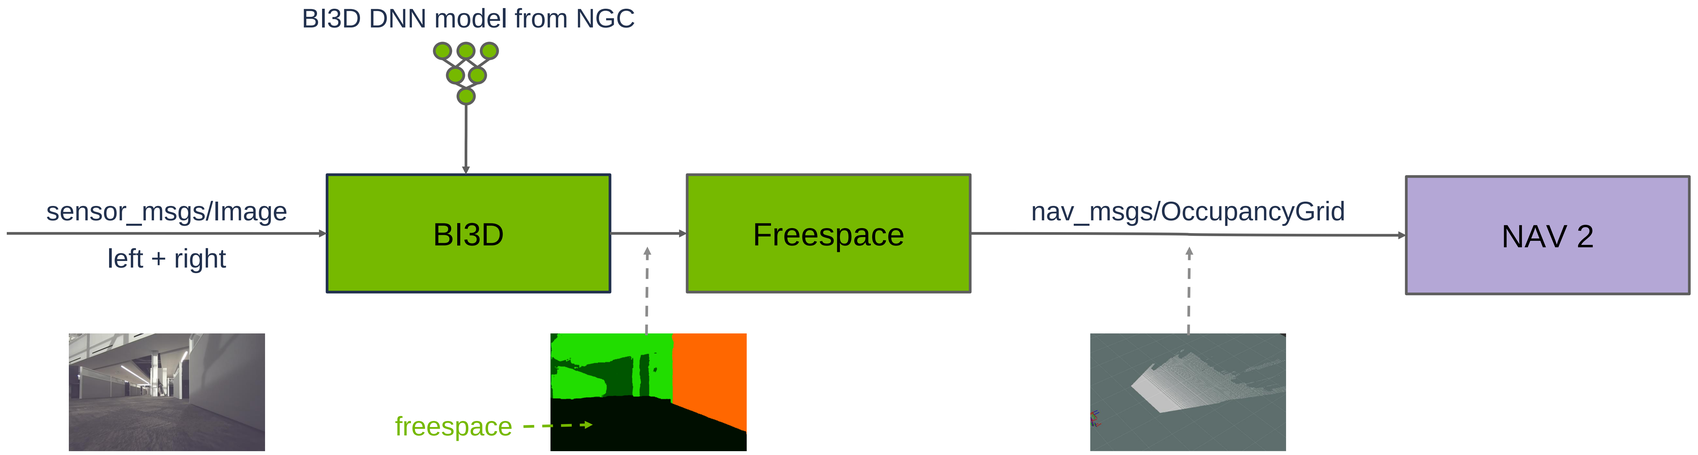

# Image Segmentation

## Classifying pixels in an image

- AI based perception function to classify an image at the pixel level
  - Uses a pre-trained model from NGC (i.e. *PeopleSemSegNet*)
  - Provides pre-processing from image to Tensors, inference via TensorRT, and output segmentation image
- `Model trained` with a mobile robots point-of-view closer to the ground
  - Leverages TAO to fine tune models specific to the robotics application
- Best suited for robot applications where spatial location of an object in a scene informs planning
  - When fused with a depth image for range from the camera in a scene

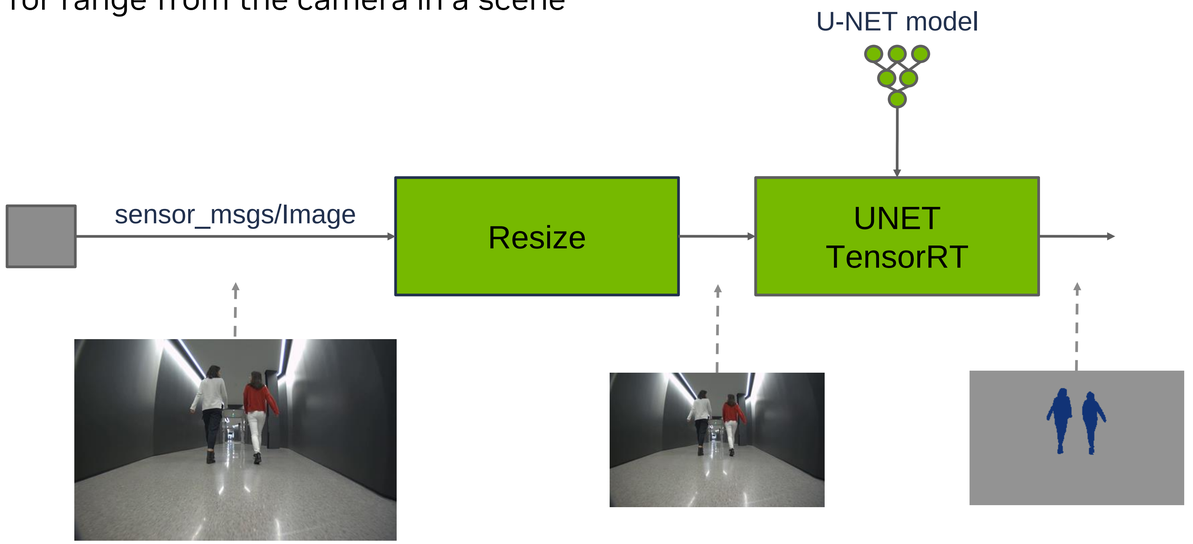

# Object Detection

## Classifying obstacles spatially in an image

- AI based perception function to detect classes of objects
  - Uses a pre-trained model from NGC (*i.e. PeopleNET*)
- Model trained with a mobile robots point-of-view closer to the ground
  - Leverages TAO to fine tune models specific to the robotics application
- Best suited for robot applications where presence and general location of an object in a scene informs behavior or additional processing

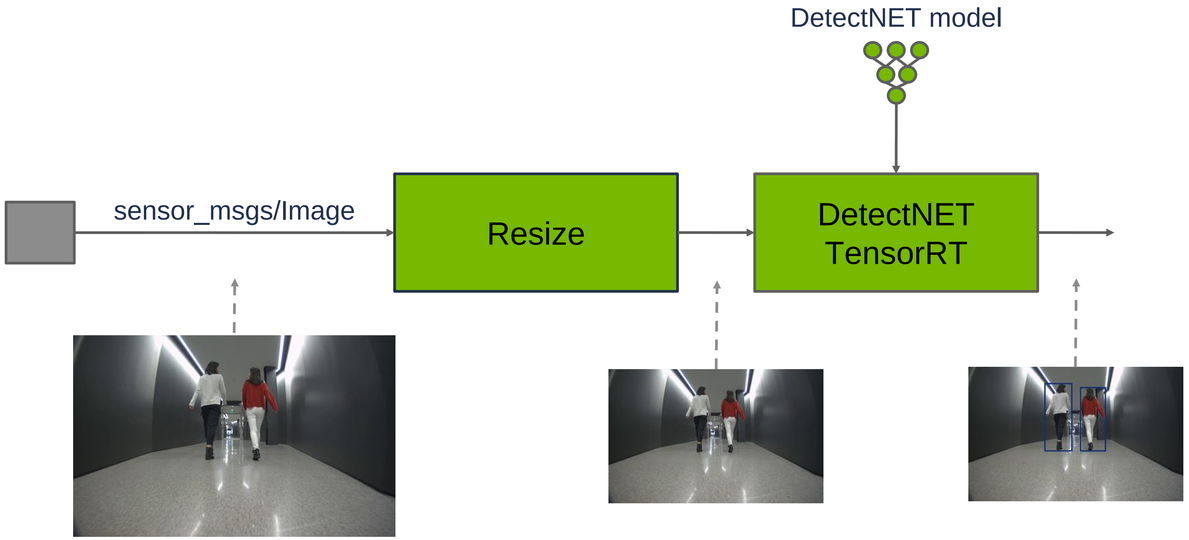

# Image Compression

## Reducing image memory footprint in rosbags

- Accelerated H.264 encode for image real-time data recording
  - Data campaigns to develop, train, and test AI and CV perception functions
  - Event recorders to capture incidents and events of interest for offline analysis

- Accelerated H.264 decode for open-loop re-simulation for analysis, debug, and testing of perception functions

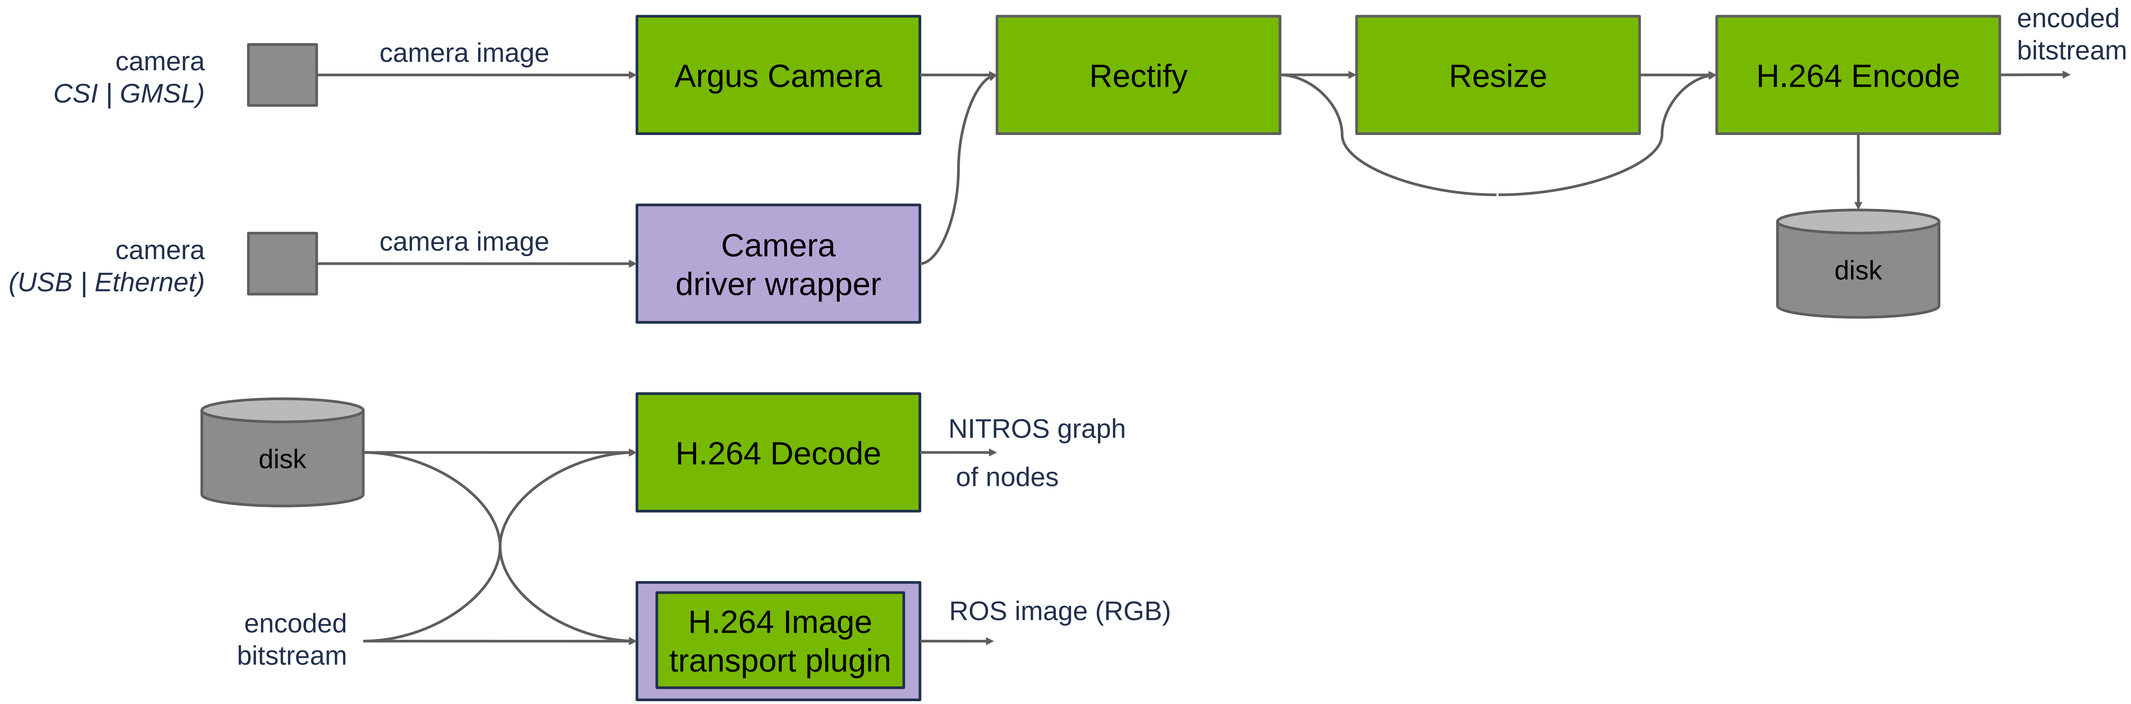

# Argus Camera

## Sensor processing for vision-based perception

- CSI or GMSL camera with ISP processing in Jetson native hardware
- Leverages hardware time synchronization across cameras
- Uses hardware timestamping for acquisition time of camera frames and for synchronization with LIDAR
- Reduced memory copies and latency compared to USB or Ethernet cameras

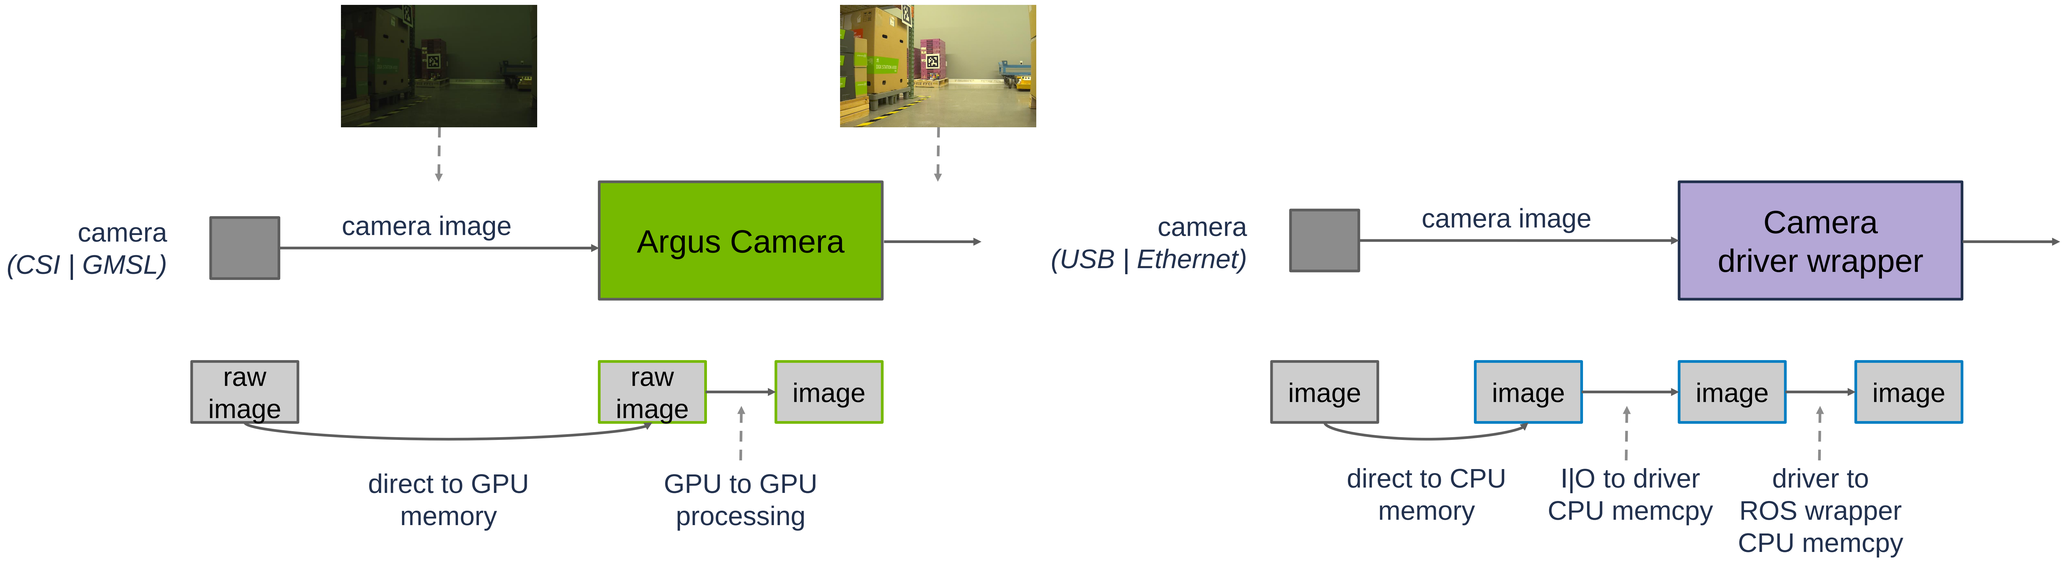

# NVIDIA ISAAC ROS – More Than Just Modules

# Nitros

## Optimizing ROS for accelerated computing

- Open Robotics and NVIDIA jointly developed and finished changes for accelerated commuting in ROS 2 Humble

  REP-2007 for accelerated message passing

  REP-2009 for optimized message formats

- NITROS (Nvidia Isaac Transport For Ros) implements REP-2007 & REP-2009 for Isaac

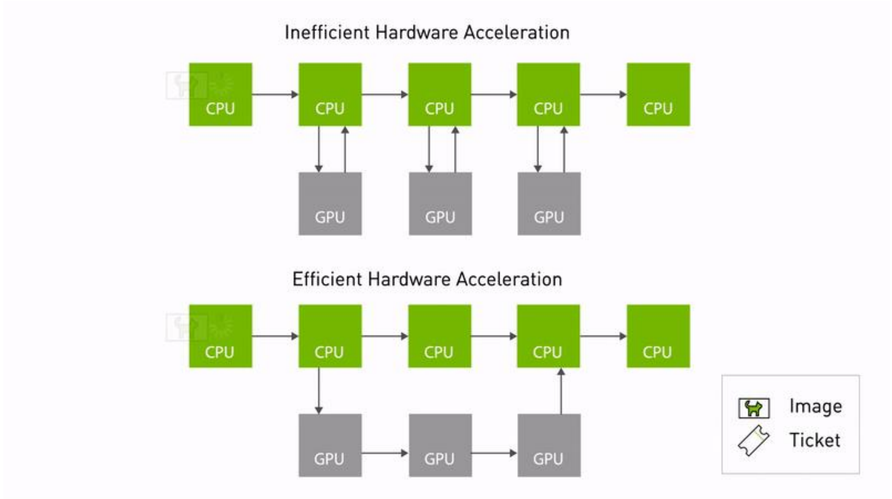


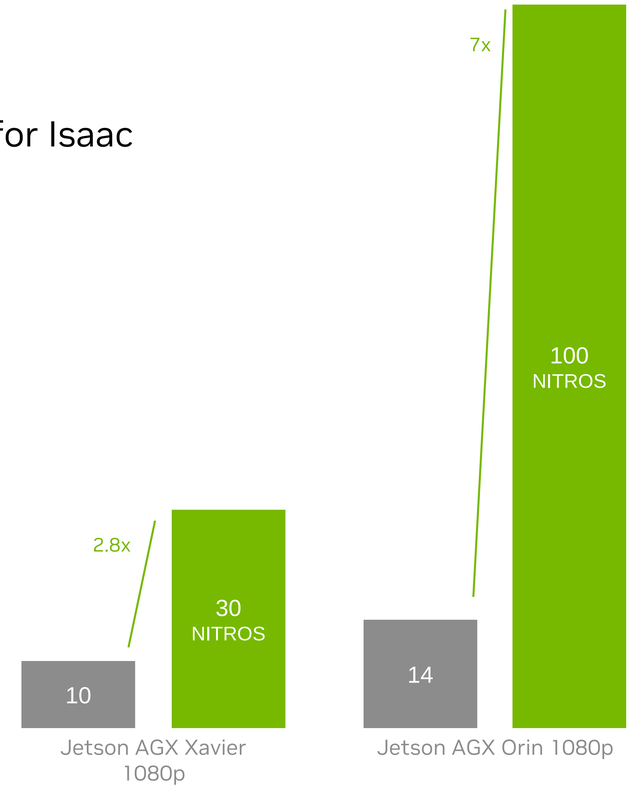


https://discourse.ros.org/t/ros-2-humble-hawksbill-released/25729/10

# Benchmark Tools

## Measuring performance of graphs of nodes

- Robotics applications require complex graphs of heterogeneous computation to perform perception, planning, and control
  - Computation need to be executed deterministically and with known latency
  - Fixed budget for heterogeneous computation (TOPS) and throughput

- ros2_benchmark provides open source tooling to measure the performance for graphs of nodes in ROS
  - Realistic assessment of robot application performance
  - Informs timeouts for monitors checking robot health at run-time

- r2b dataset provides standard time synchronized full sensor data rate captures to rosbag for evaluation and testing
  - 8 diverse sequences provided under CCv4 Attribution

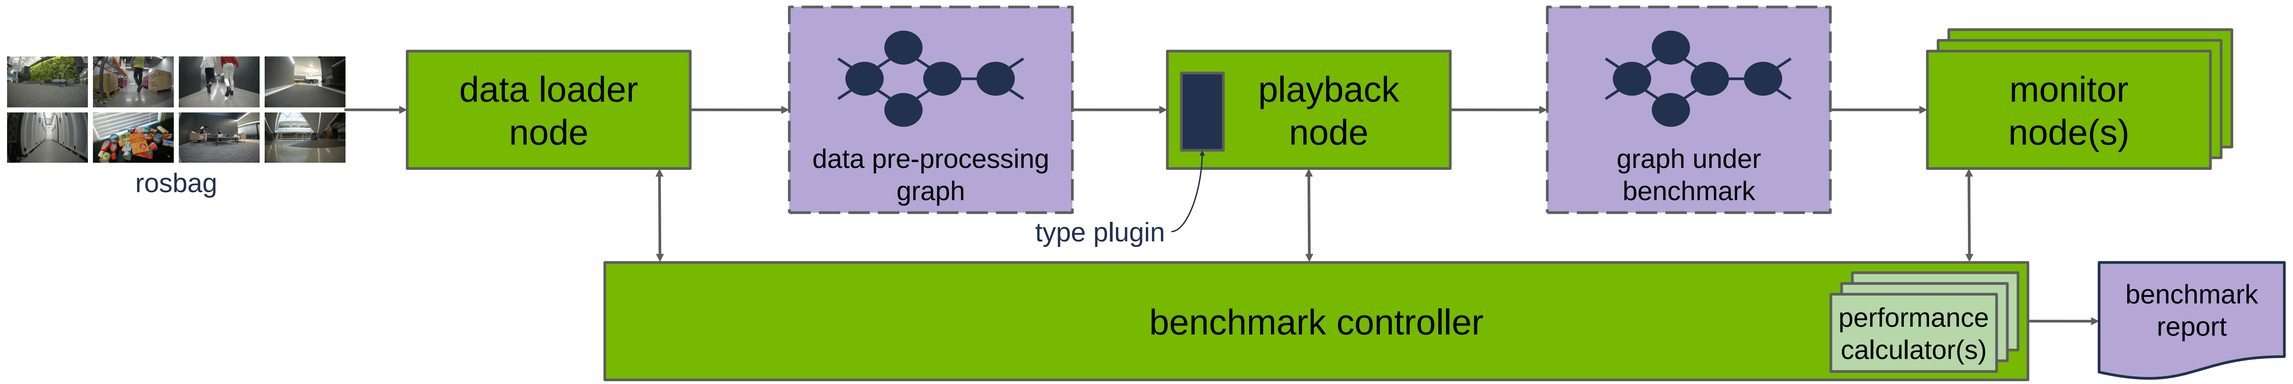

# Software 2.0

## Isaac ROS as part of modern data driven workflow for robotics application development

- Colcon native build produces Docker images from source control managed by Jenkins for modern CI|CD
- Kubernetes native testing for simulation of SIL *(software in the loop)*, HIL *(hardware in the loop)*, and canary robots
- Continuous feedback from deployment with data driven event logging

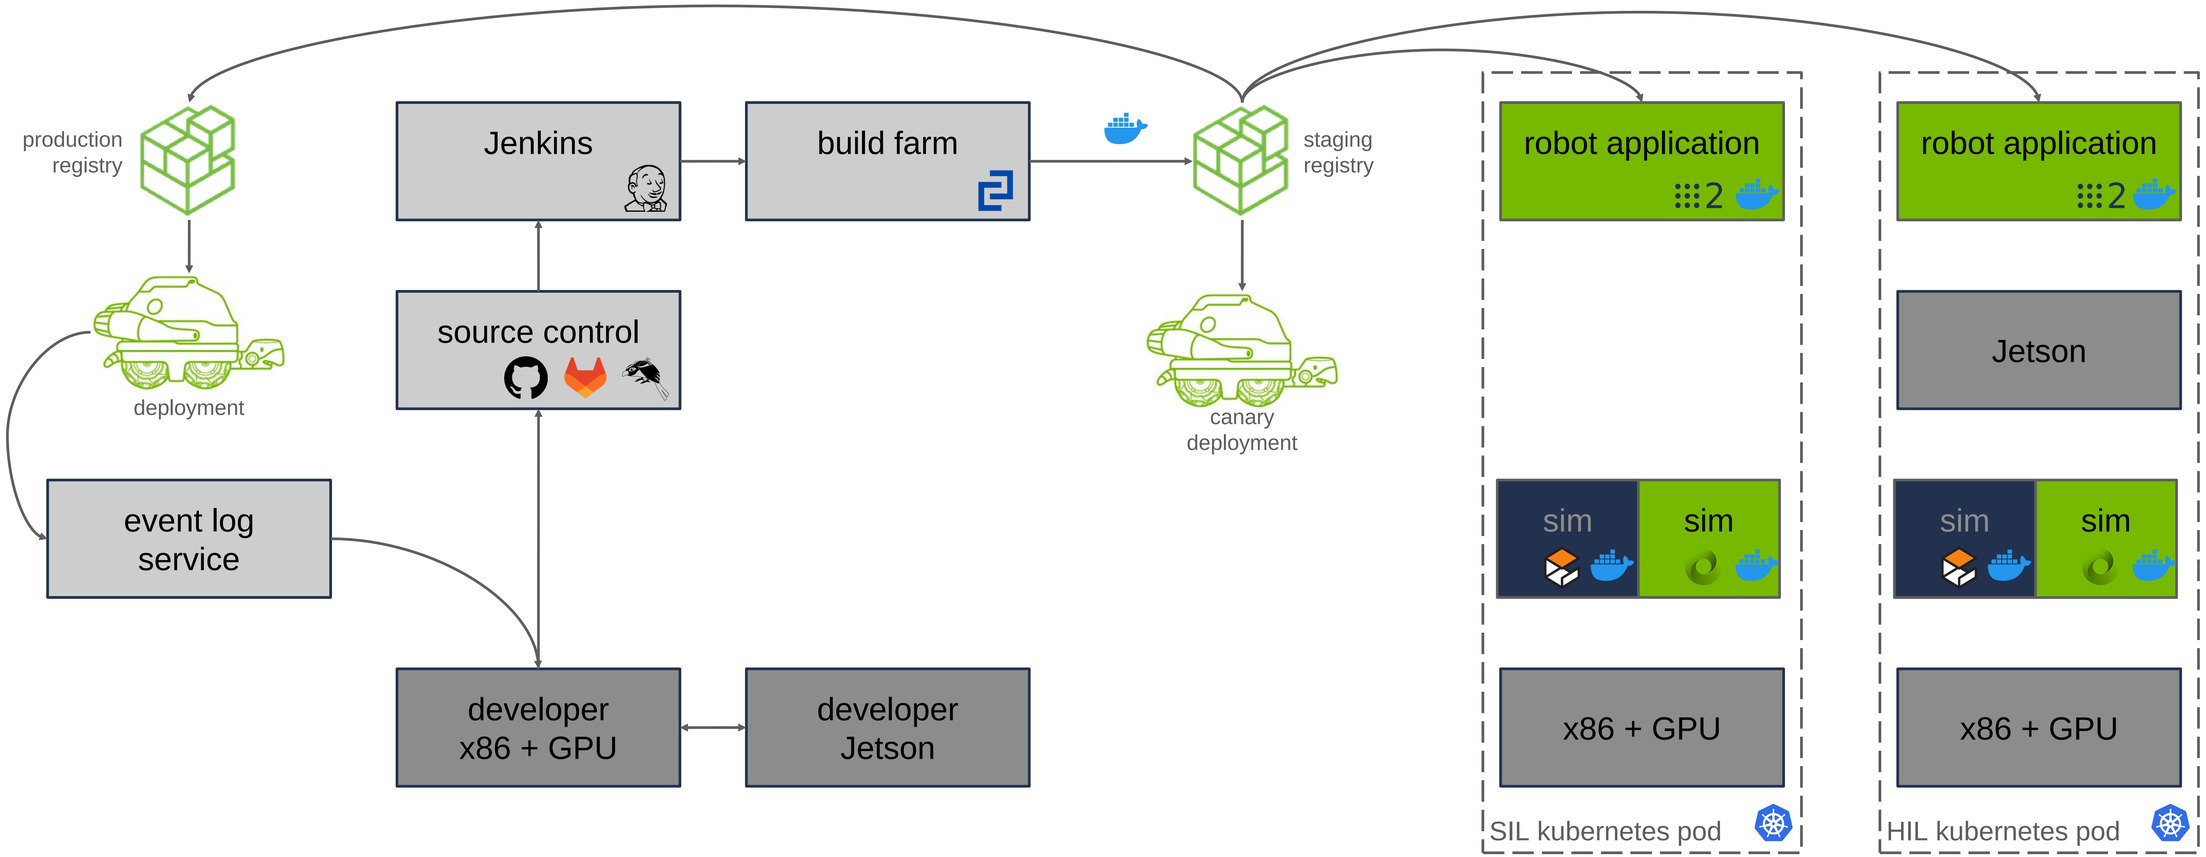

# Generative AI in Simulation

# Generative AI

## The iPhone moment of AI

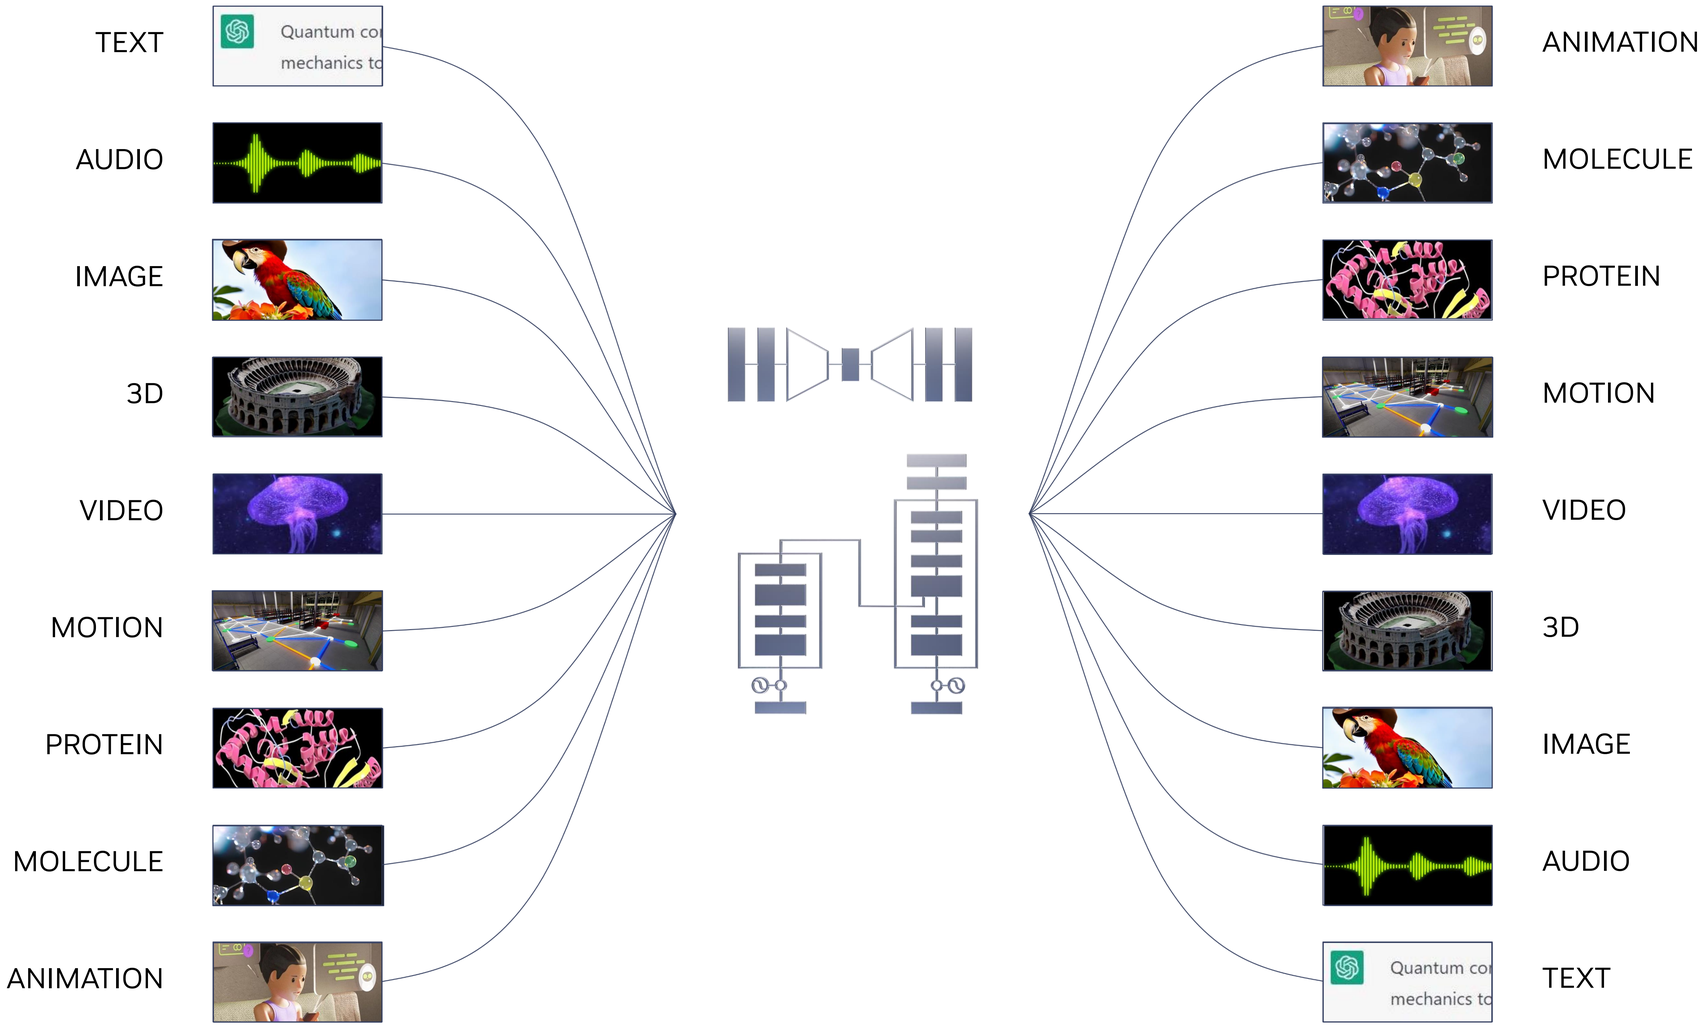

# What is Generative AI?

Let’s ask ChatGPT

Generative AI, short for Generative Artificial Intelligence, refers to a subset of artificial intelligence (AI) techniques and models that are designed to generate new, original data or content that resembles human-created content. Instead of performing tasks like classification or prediction, generative AI models are trained to create something new, whether it’s text, images, music, or other forms of data.

# Generative AI

## Two main model types

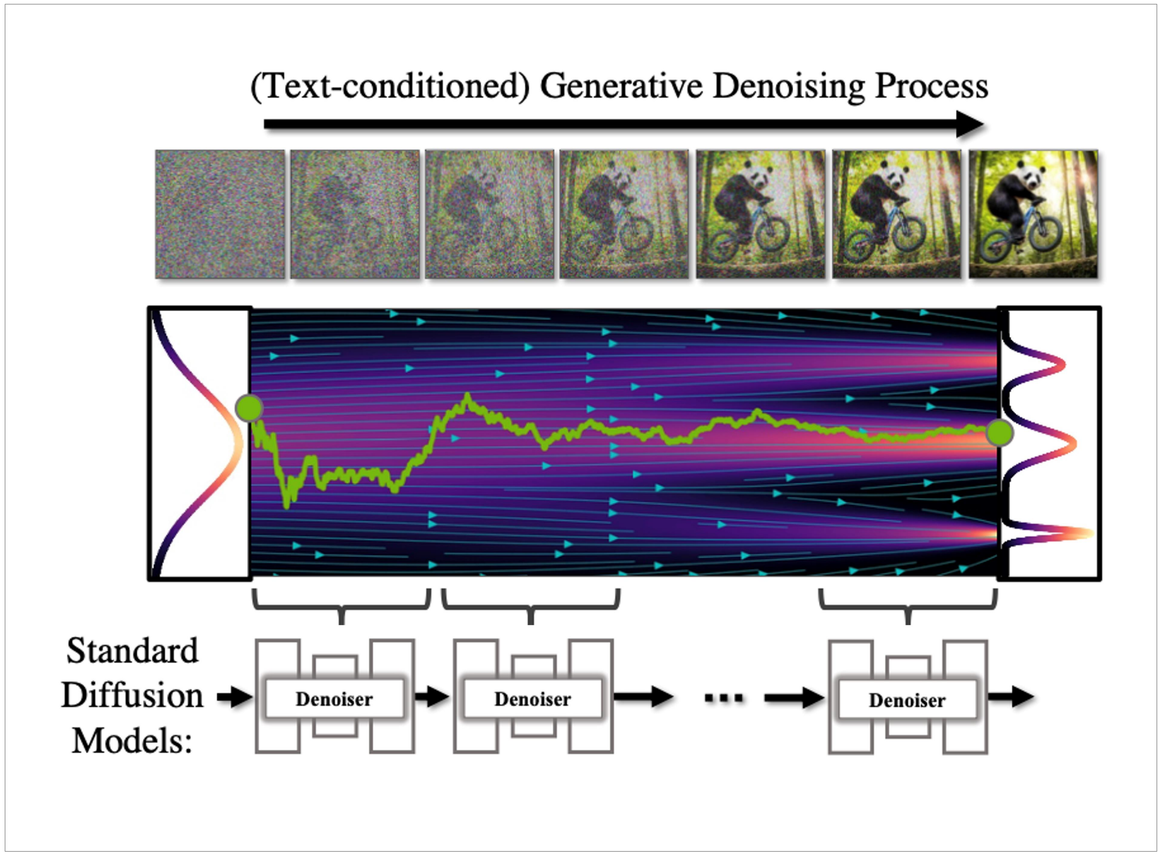


### Diffusion Models

Generate coherent images from noise

Stable Diffusion, Midjourney, DALL-E, NVIDIA Edify

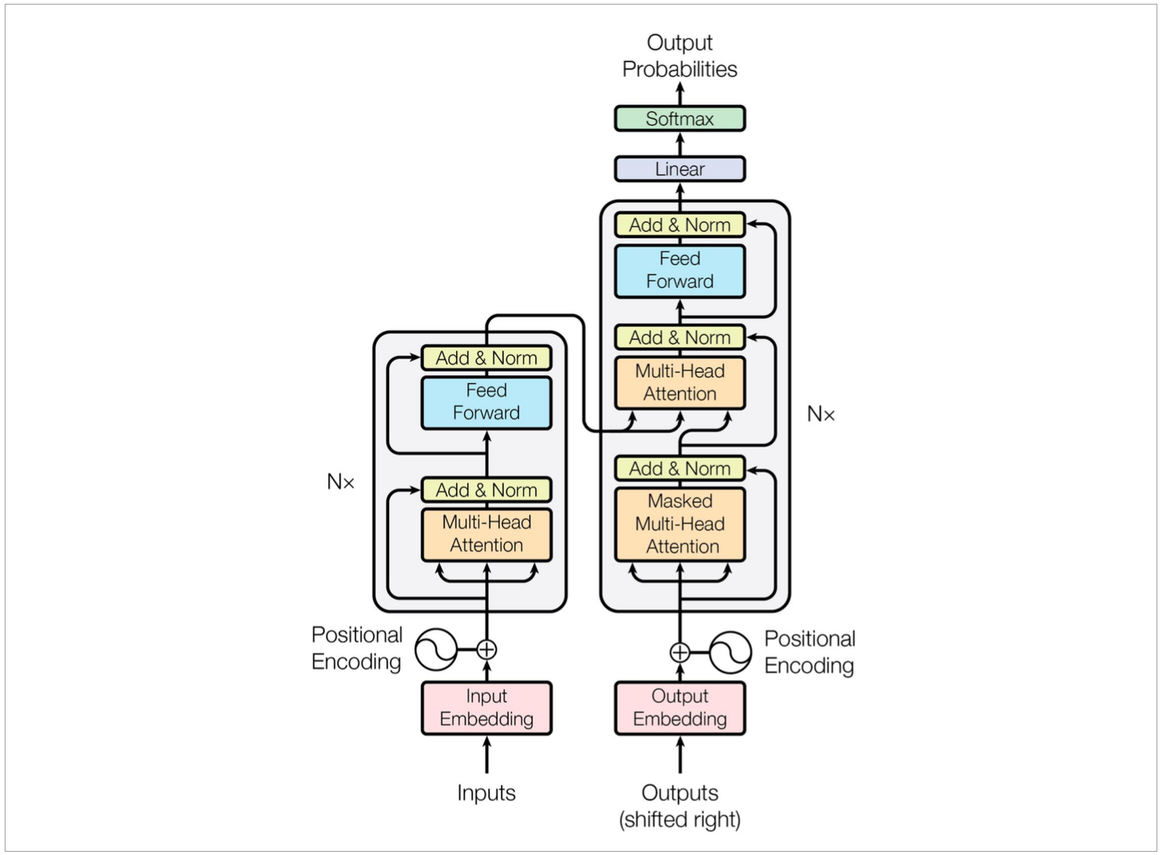


### Transformer Based Models

Unsupervised training on large corpus data, fine-tune for downstream NLP (Natural Language Processing) tasks

BERT, Open AI GPT, BARD, NVIDIA NeMo

*Fig.2: Attention Is All You Need*

Foundation Models

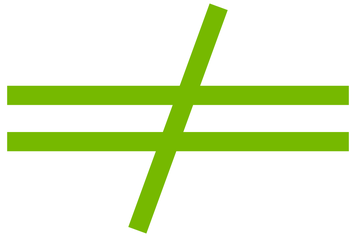

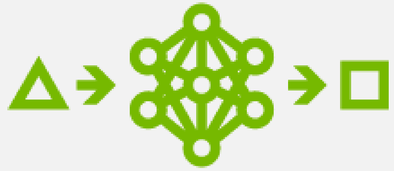

Generative AI Models

# What is the Role of Generative AI in Robotics?

Let’s ask ChatGPT again

Generative AI plays a crucial role in robotics by enabling robots to perform various tasks more effectively and flexibly. Here are some of the key roles and applications of generative AI in robotics:

**Motion Planning and Control**

**Simulation and Training**

**Object Recognition and Scene Understanding**

**Human-Robot Interaction**

# The Path to Generative AI in Robotics

## Transformer based models

### LANGUAGE

Predict the next word in a sentence

Use web-scale text data

### IMAGES/VIDEOS

Predict the next patch in an images

Use web-scale images/videos

### ROBOTICS / CONTROL

Predict the next robot action

Use human expert demonstrations, teleoperation, AI
can leverage web-scale text + visual data
can leverage simulation

# Challenges in the Robotics Industry

## How can Generative AI help?

Lack of enough data
for training

Environment diversity
and variations

Programming robots and
interacting with them

# Generative AI in Robotics

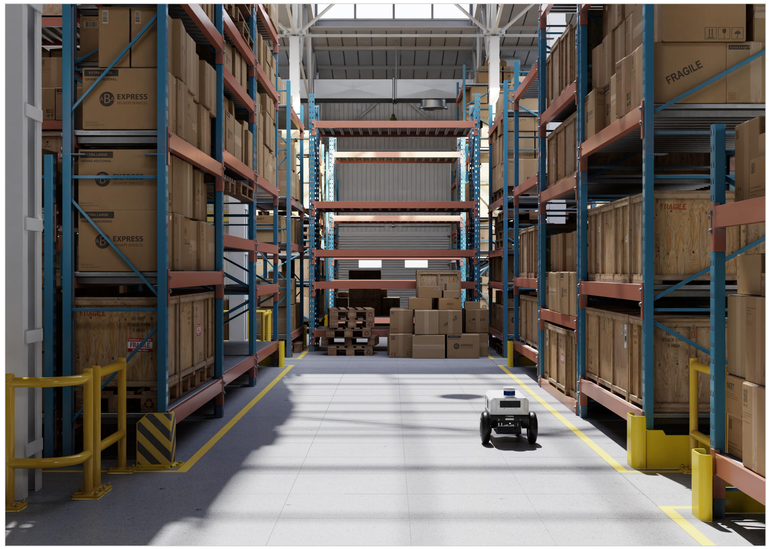


## Simulation

Synthetic data generation and augmentation
Building virtual worlds

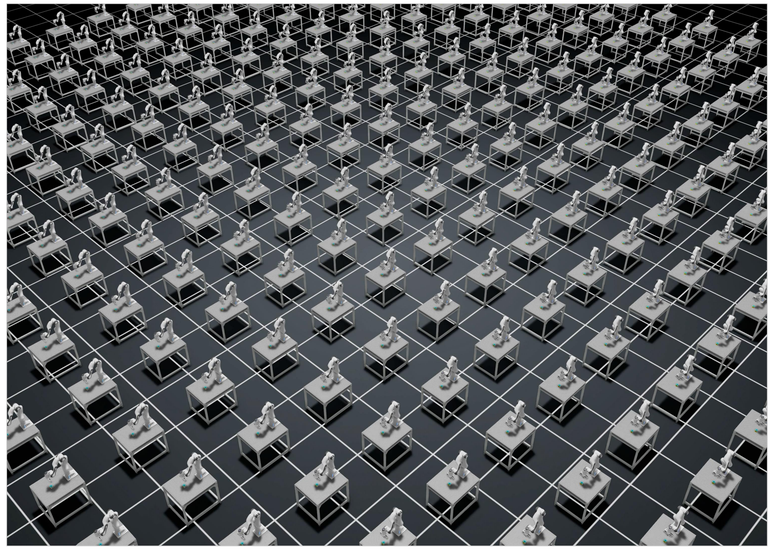

## Learning

Visual understanding
Control/reward code generation
Control policy learning
Interaction/Instruction following
Motion generation

## Deployment

# Generative AI in Synthetic Data Creation (SDG)

# Synthetic Data Creation

## Text-to-image generative models

**Stable Diffusion**
Ludwig Maximilian University of Munich & IWR, Heidelberg University, and Germany Runway ML

**DALL-E 2**
OpenAI

**eDiff-I**
NVIDIA

*"A photo of a plate at a restaurant table with spaghetti and red sauce. There is sushi on top of the spaghetti. The dish is garnished with mint leaves. On the side, there is a glass with a purple drink, photorealistic, dslr."*

# Synthetic Data Augmentation


**CACTI**
CMU and Meta

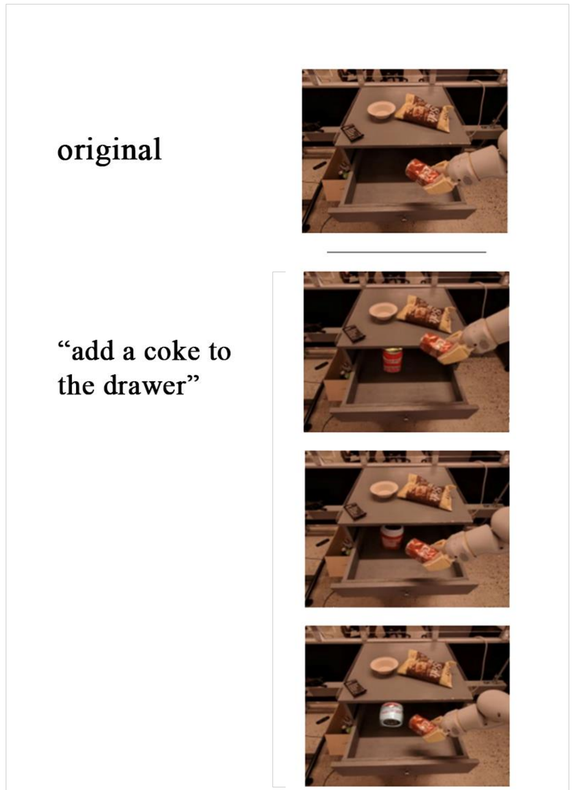


**ROSIE**
Google and Google Research

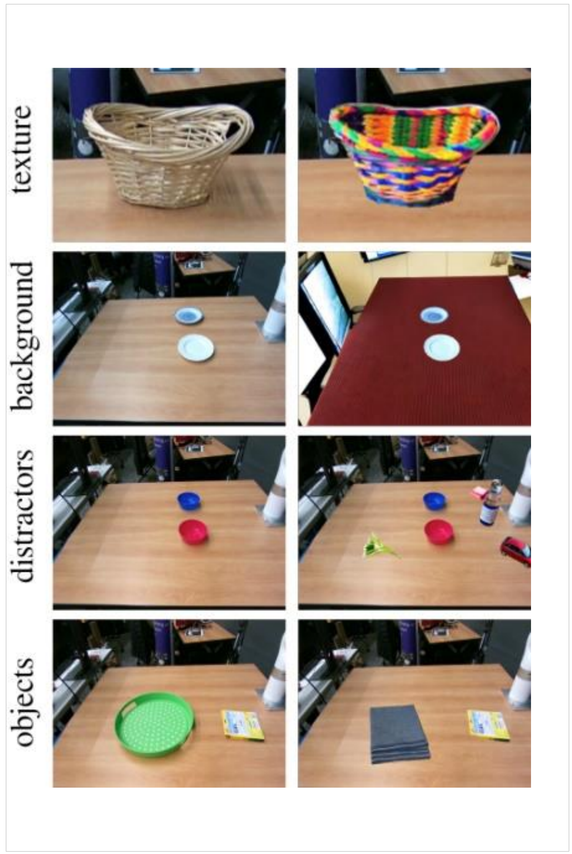


**GenAug**
U. Of Washington and Meta AI

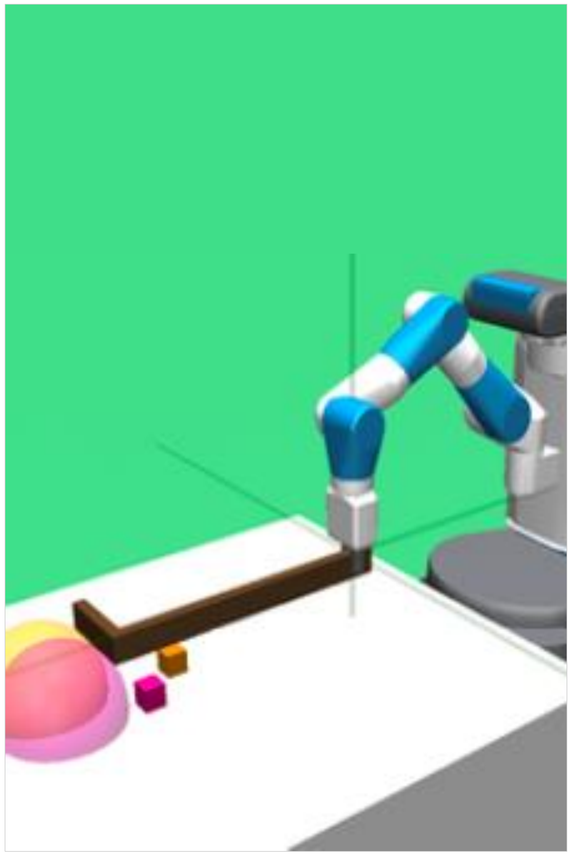

**MoCoDA**
U. of Toronto, Vector Institute, and NVIDIA

# Building Virtual Worlds

## Generative AI can help everywhere

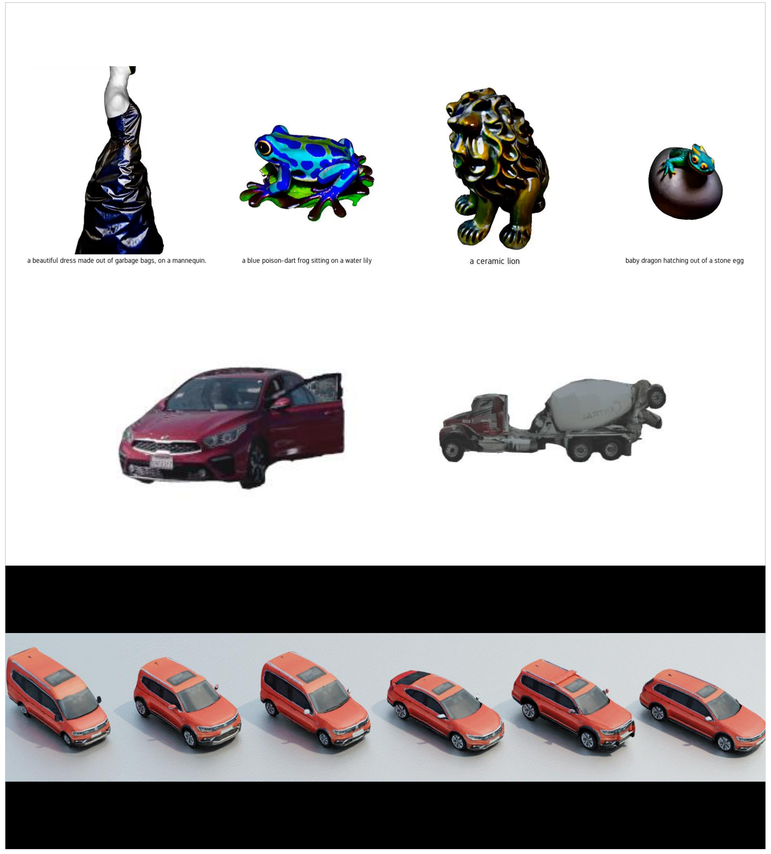


Asset Creation

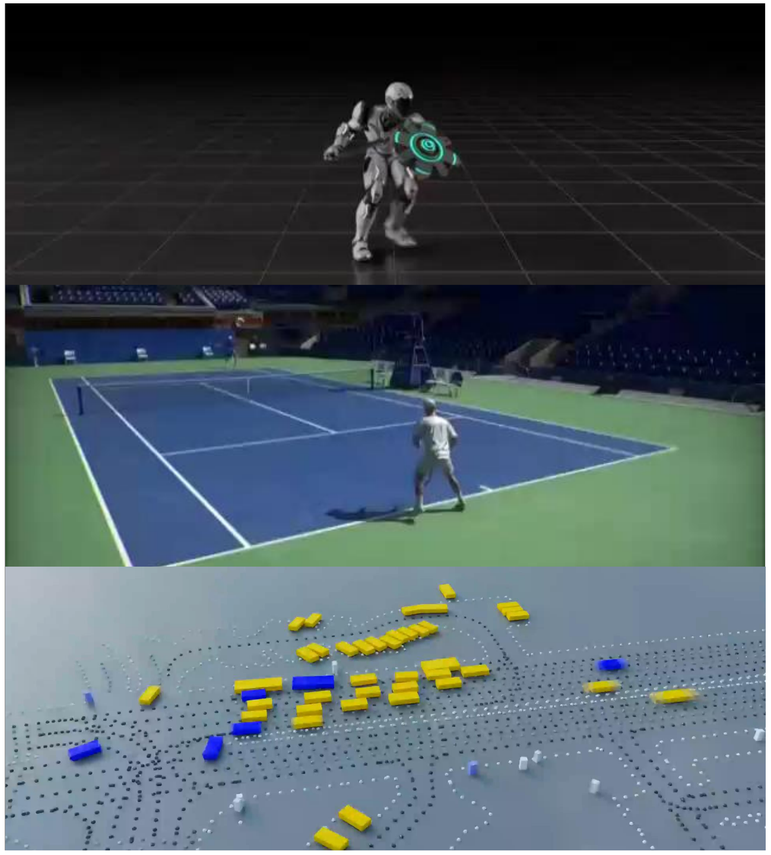

Behavior & Animation

World Capture + Augmentation

# 3D Asset Creation

## Text-to-3D generation

*“Burned Car”*

*“Polygon Car”*

**Get3D**
NVIDIA, U. of Toronto, and Vector Institute

*“a silver platter piled high with fruits”*

**Magic3D**
NVIDIA

# Texture Generation

## Text-Guided stable diffusion

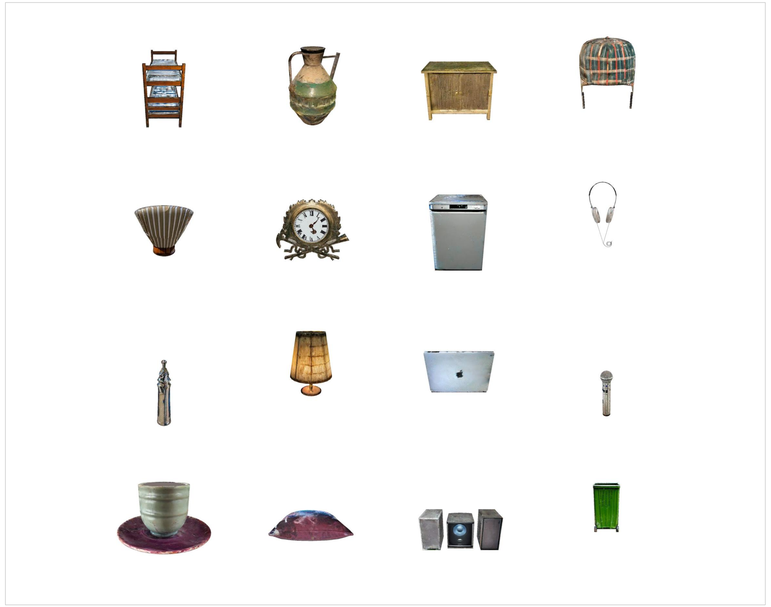

**TEXTure**

Tel Aviv University

**TexFusion**

NVIDIA, U. of Toronto, and Vector Institute

# Building Virtual Worlds

## Using Neural Radiance Fields (NeRFs)

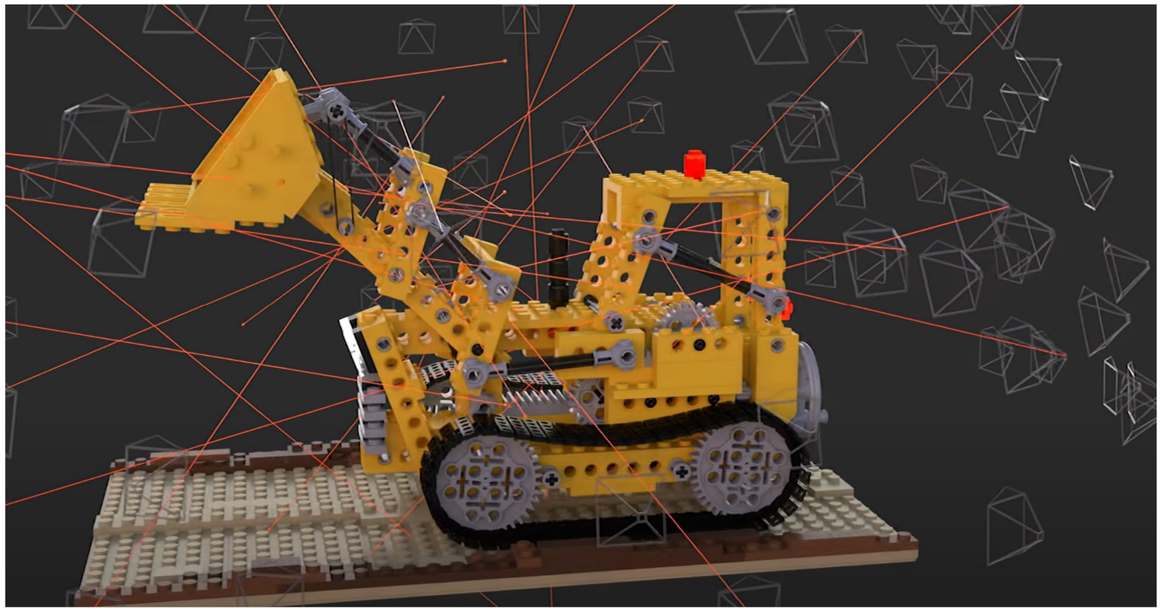

### NeRFs

Introduced in 2020 by UC Berkeley, Google Research, and UC San Diego
Several hours to train even a simple scene
Could render at rates in the seconds-per frame range

### Instant NGP*

Introduced in 2022 by NVIDIA
Train in seconds
Render in real time frame rates

**TIME Magazine Named NVIDIA Instant NeRF a Best Invention of 2022*

# Using LLMs* to Create Virtual Worlds

## Combining ChatGPT with NVIDIA Omniverse

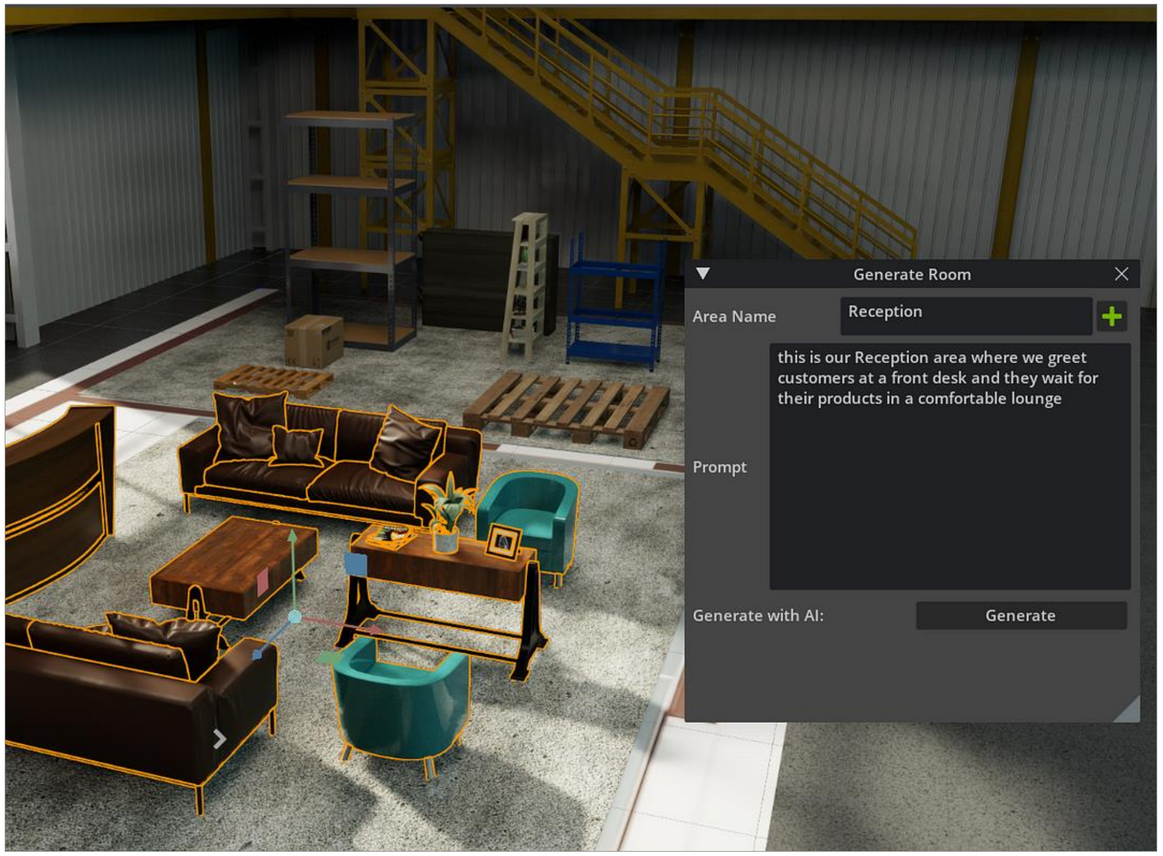


1. User input: *"This is the room where we meet our customers. Make sure there is a set of comfortable armchairs, a sofa and a coffee table."*
2. Create the prompt for ChatGPT
3. Pass the output from ChatGPT to a DeepSearch API which retrieves the best list of assets
4. Place the assets in the scene

**LLMs: Large Language Models*

*Credit:* https://medium.com/@nvidiaomniverse/chatgpt-and-gpt-4-for-3d-content-generation-9cbe5d17ec15

# Using LLMs* to Create Virtual Worlds

## Assisting robot simulation

### Developer Assistance

*“Write a Python function that randomizes lighting in the scene”*

```python
def randomize_sphere_lights():
    lights = rep.create.light(
        light_type="Sphere",
        color=rep.distribution.uniform((0.0, 0.0, 0.0), (1.0, 1.0, 1.0)),
        intensity=rep.distribution.uniform(100000, 3000000),
        position=rep.distribution.uniform((-250, -250, -250), (250, 250, 100)),
        scale=rep.distribution.uniform(1, 20),
        count=NUM_LIGHTS,
    )
    return lights.node
```

### Scene Generation for Simulation

*“Generate a realistic dataset for a warehouse”*

### Scene Editing for Simulation

*“Can I have a more cluttered scene?”*

*LLMs: Large Language Models*

# Generative AI in Robot Learning

# Visual Understanding

## Transformer models

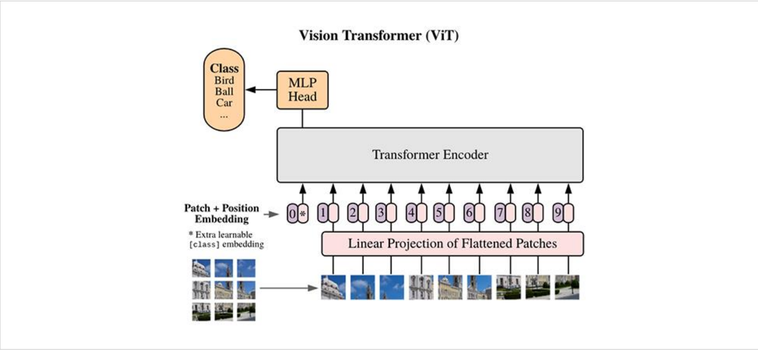


**VIT**
Google Research

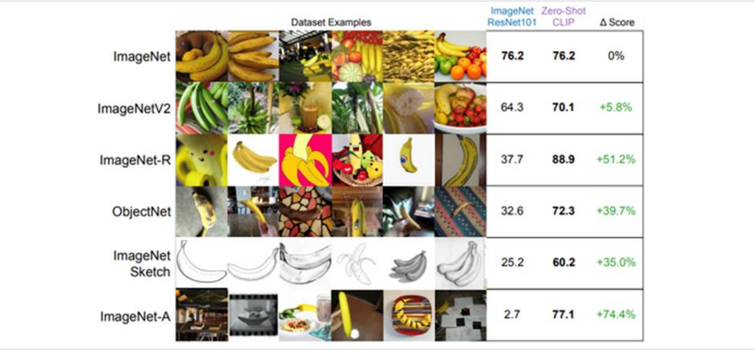


**CLIP**
OpenAI

**SegFormer**
The University of Hong Kong,
Nanjing University, NVIDIA, Caltech

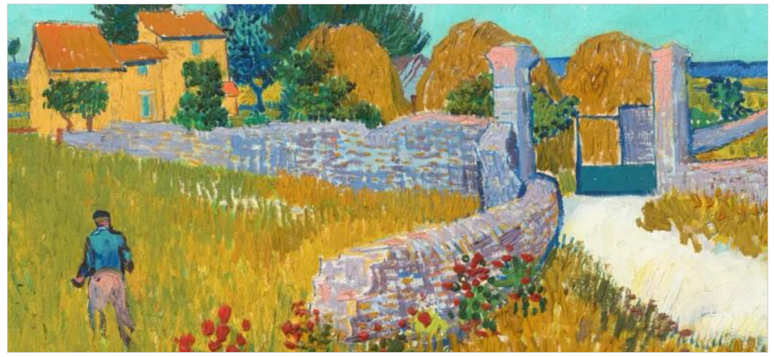

**SAM**
Meta AI

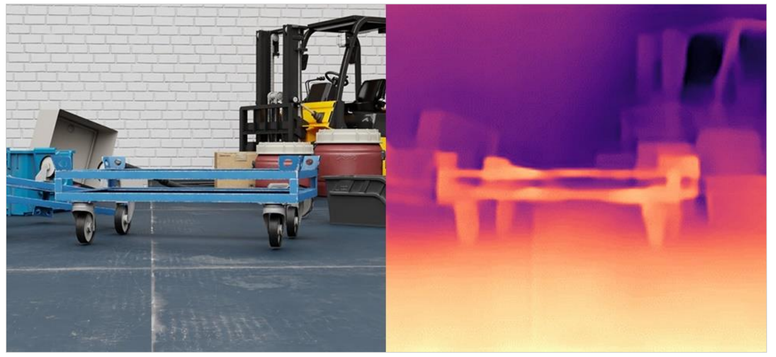

**DINOv2**
Meta AI

# Writing Code for Control/Reward Functions

## Multi-task transformers

**PROGPROMPT: VirtualHome Demo**

```
from actions import walk <obj>,
grab
```

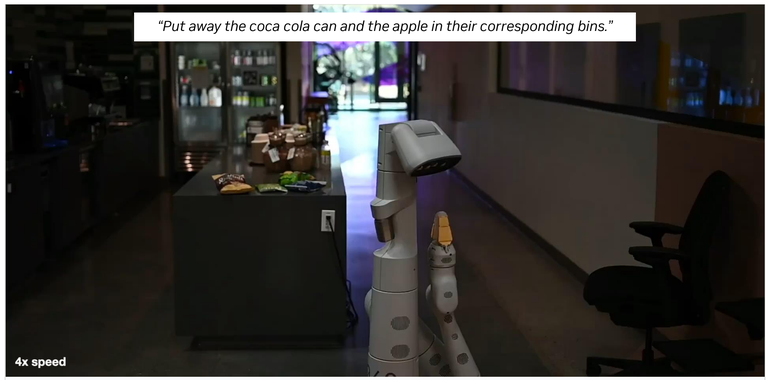


```
# This is the compost bin.
compost_bin_pos, compost_bin_angle = get_robot_pos_and_angle()
compost_bin_name = 'compost bin'
say('ok')

# This is the recycle bin.
recycle_bin_pos, recycle_bin_angle = get_robot_pos_and_angle()
recycle_bin_name = 'recycle bin'
say('ok')

# This is the landfill bin.
landfill_bin_pos, landfill_bin_angle = get_robot_pos_and_angle()
landfill_bin_name = 'landfill bin'
say('ok')
```

User

**ProgPrompt**
USC and NVIDIA

**Code as Policies**
Google

**Language to Rewards**
Google DeepMind

# Control Policy Learning

## Multi-task transformers

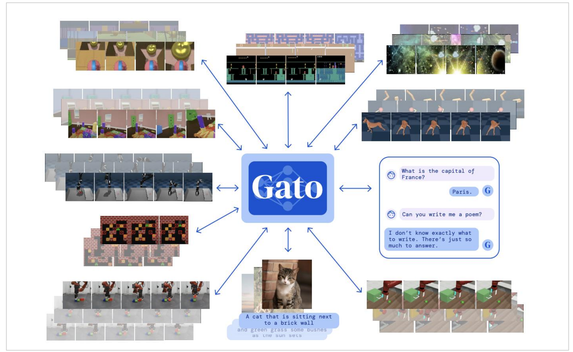


**GATO**
Google DeepMind

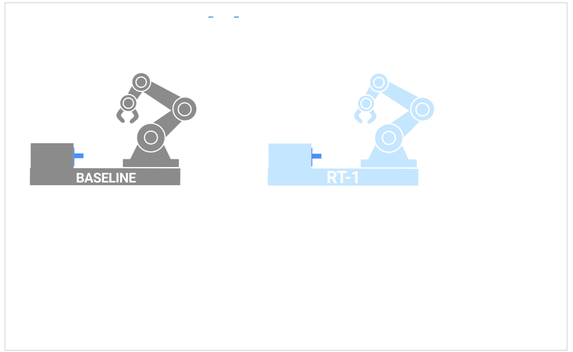


**RT-1**
Google, Everyday Robots, and Google Research

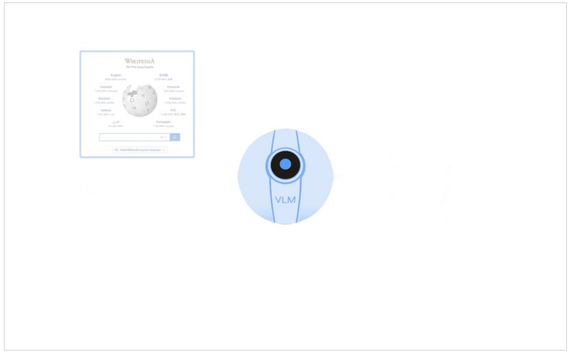


**RT-2**
Google DeepMind

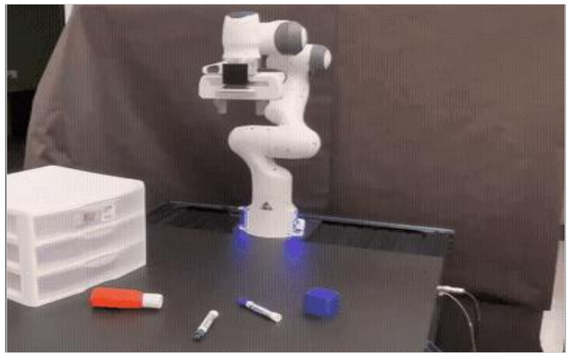

**RVT**
NVIDIA

# Interaction/Instruction Following

## Multi-task transformers


### TidyBot

Princeton University, Stanford University,
The Nueva School, Google, and Columbia University


### SayCan

Google and Everyday Robots


### PaLM-E

Google, TU Berlin, and Google Research

# Motion Generation

## Diffusion models

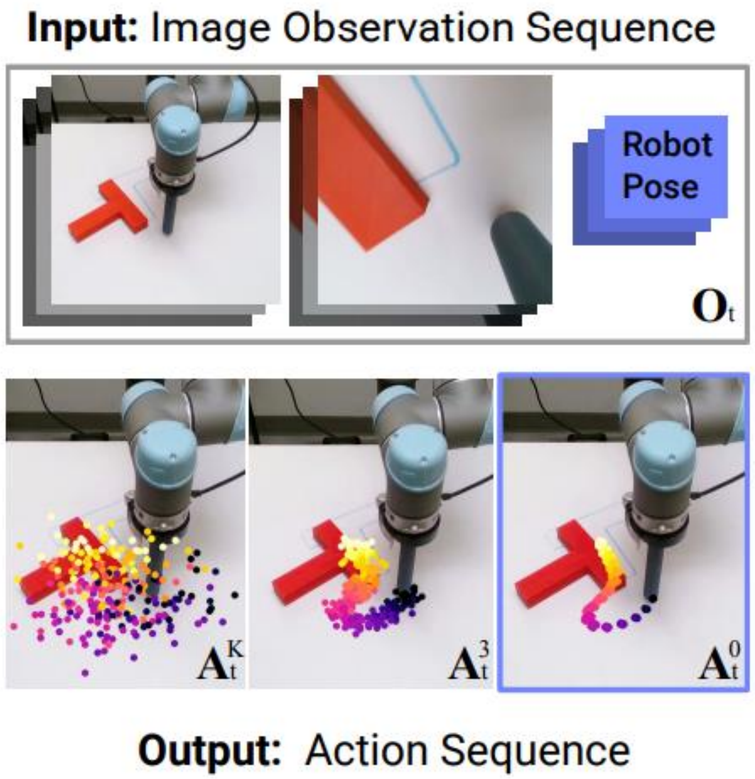


**Diffusion Policy**
Columbia University, TRI, and MIT

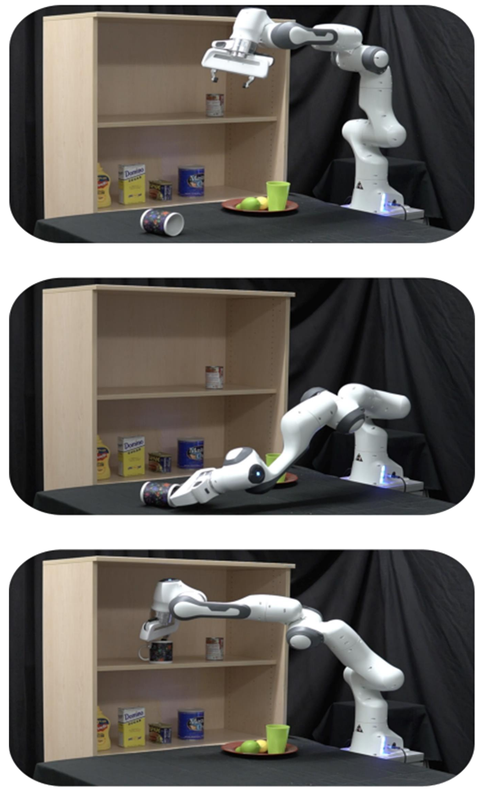

**SE(3)-Diffusion Fields**
TU Darmstadt, DFKI, Hessian.AI, and Centre for Cognitive Science

# Synthetic Data Generation and Training

## Creating a skills foundry

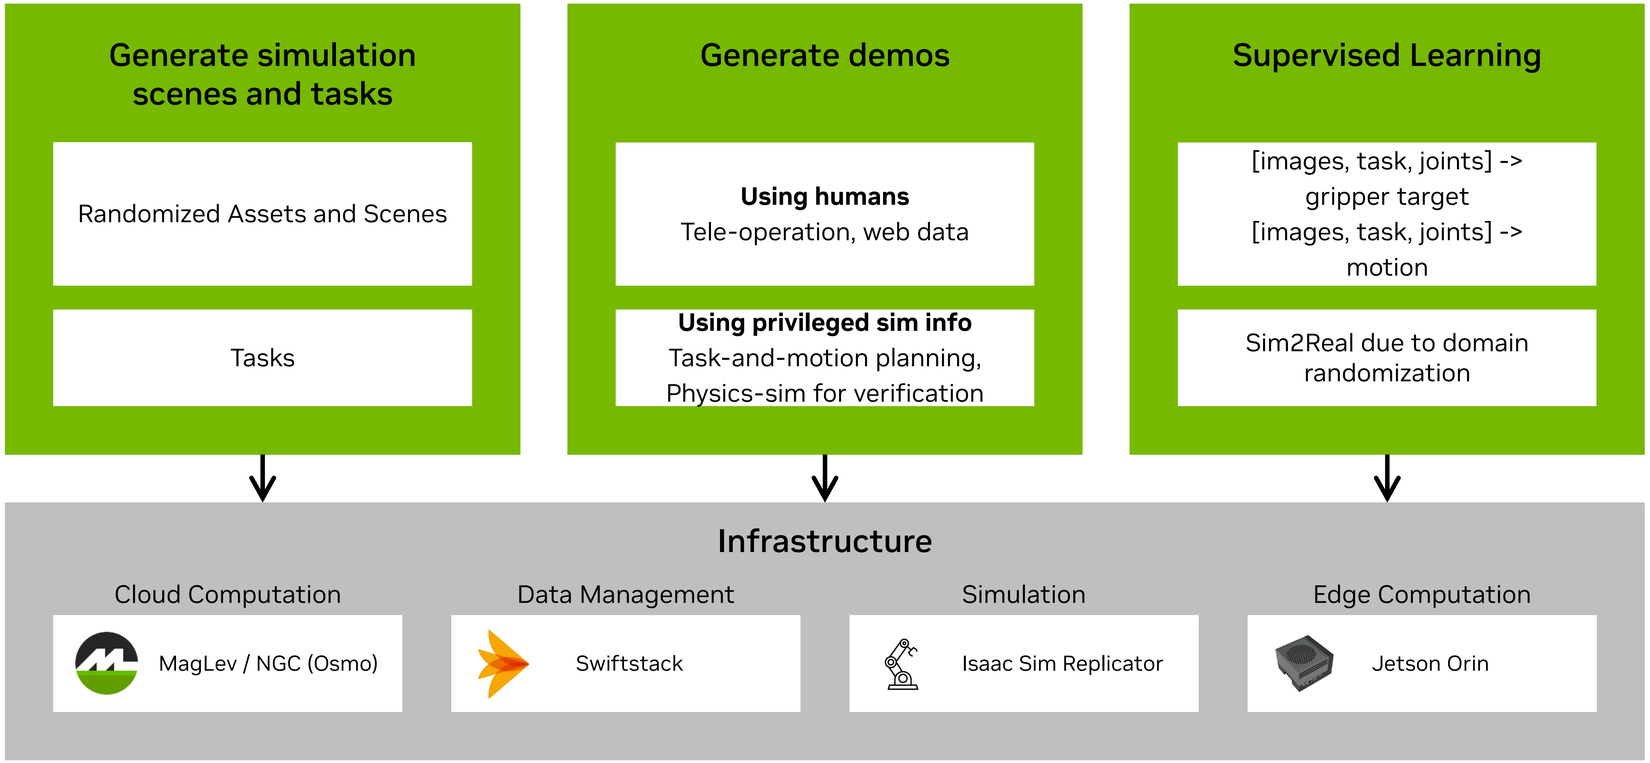

# Generative AI in Deployment

# Generative AI on Robots

## Use cases


Robot Programming

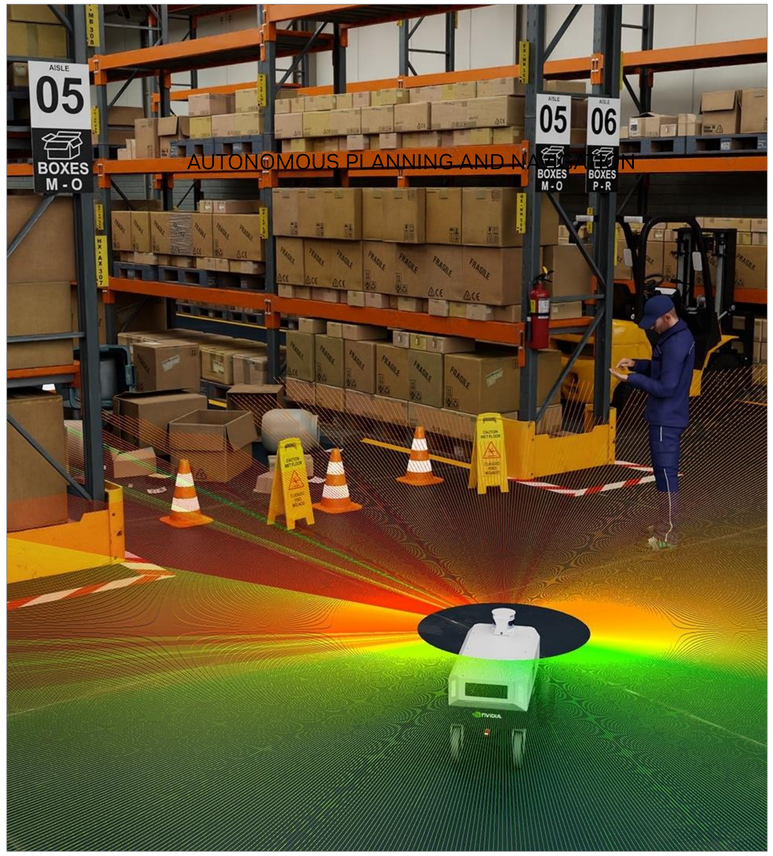


Autonomous Planning and Navigation

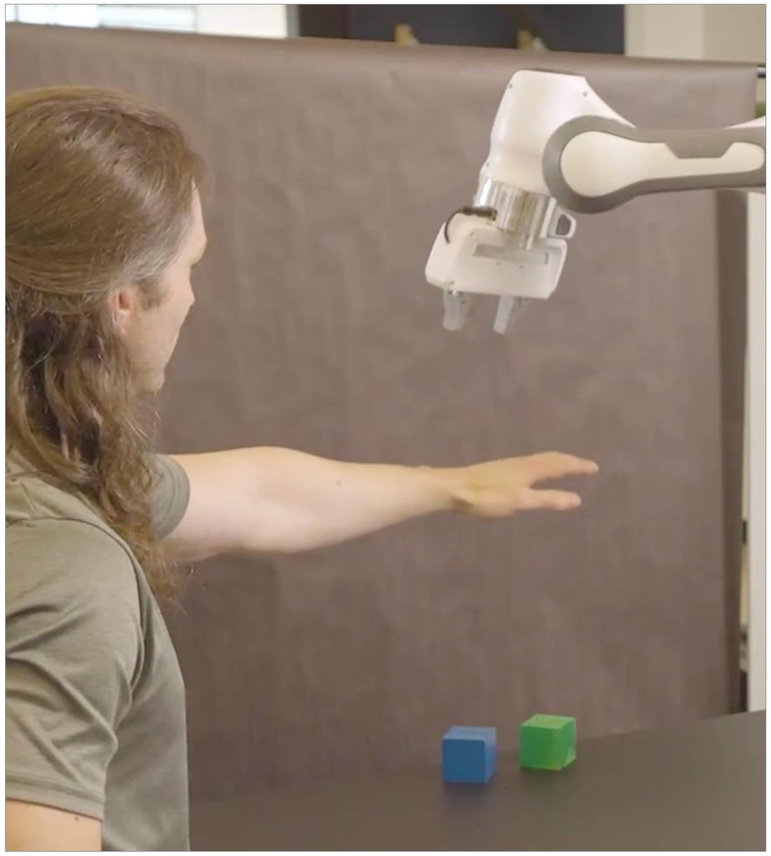

Human-Robot Interaction

# Jetson Generative AI Models

- NEW Transformer PeopleNet model delivers higher accuracy
- NEW Detect Anything model showcases zero-shot inference
- NEW muti-modal AI visual agent helps image/video search

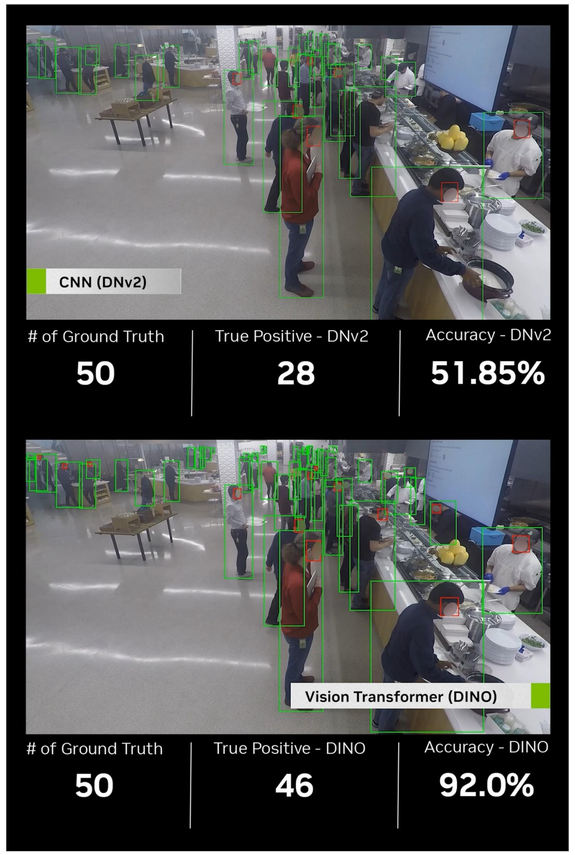


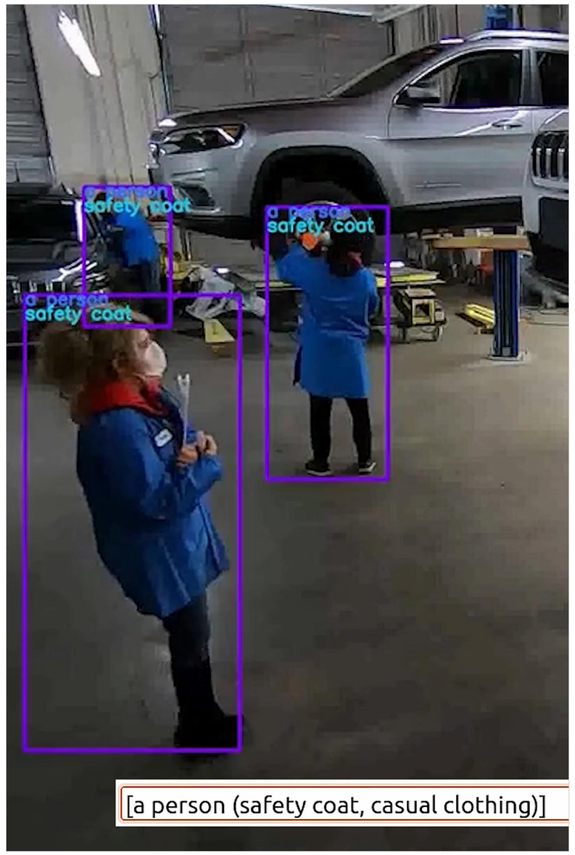


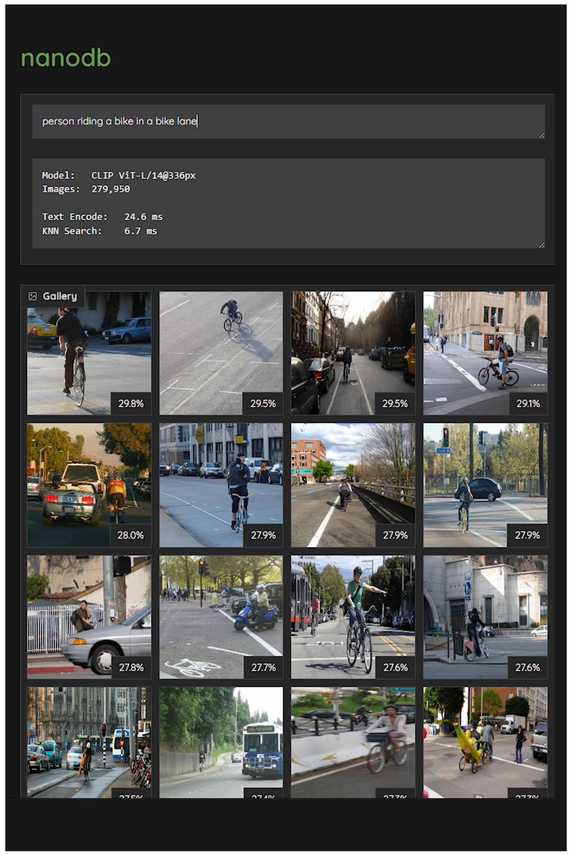


https://www.jetson-ai-lab.com/

# Deploying Generative AI on Robotics Hardware

## Knowledge distillation

MobileSAM — 146ms on Jetson Orin Nano

NanoSAM — 27ms on Jetson Orin Nano

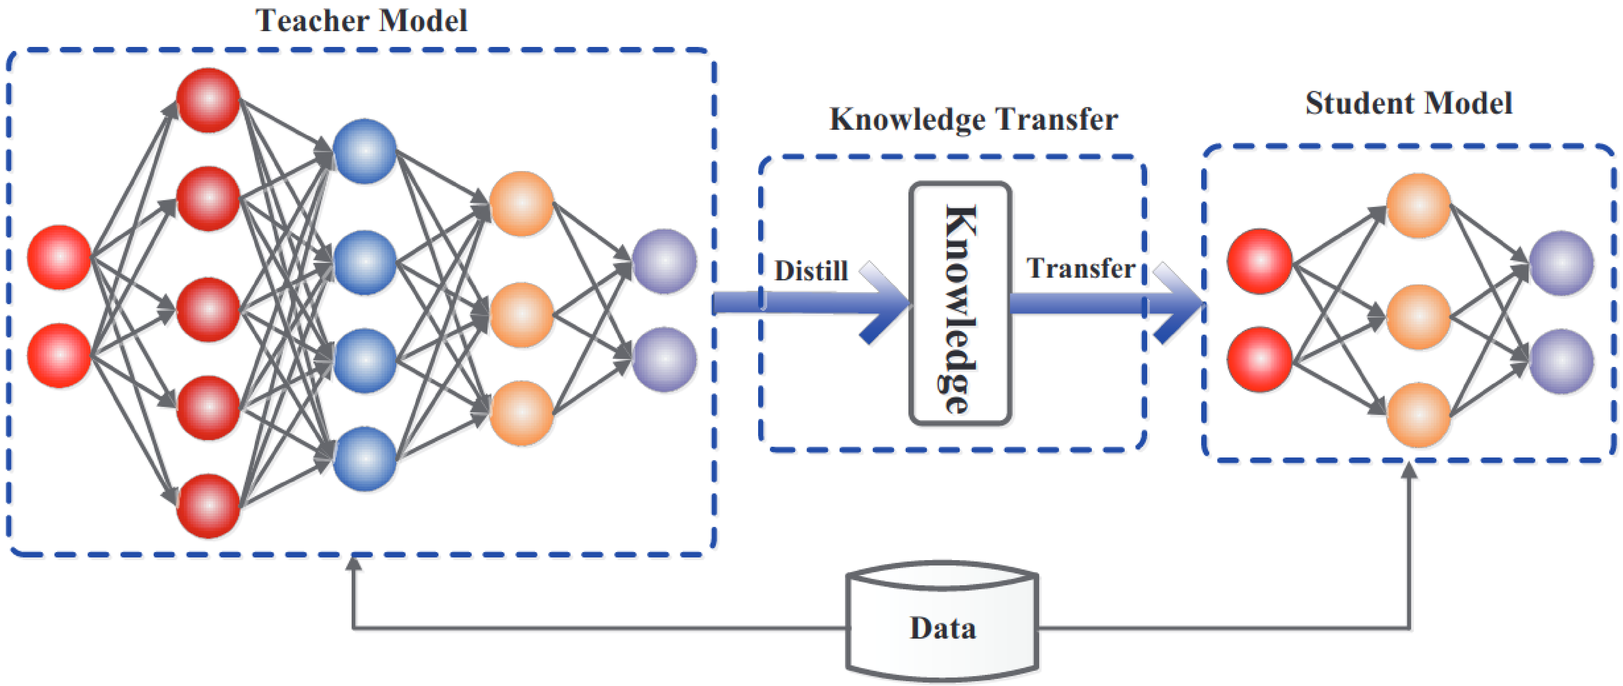


*Image Credit: Knowledge Distillation: A Survey*

# Inception Robotics Startups Using Generative AI

SCALED FOUNDATIONS

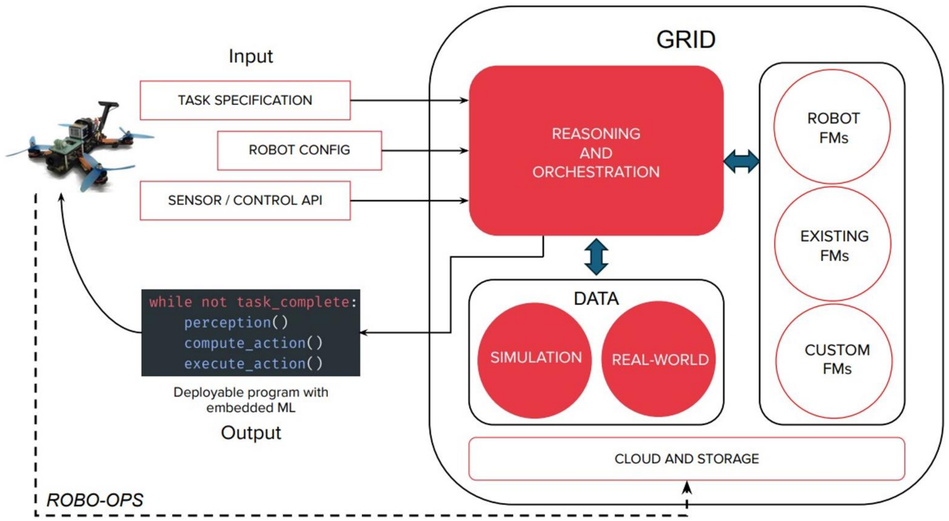


VAYU ROBOTICS

### Biology-Inspired Sensors

Nature provides examples of vision systems that are designed to ensure safety and accuracy for their host. Taking a page from nature, our sensor systems are built to combine optimal detection with cost-effective scalability.

### Scalable Machine Learning

Developments in self-supervised representation learning, generative AI, world modeling, closed-loop training and simulation building point to a new path forward for autonomous systems.

### Purposeful Design

Society requires autonomous solutions that are lightweight, energy-efficient and designed precisely for the task at hand. We use composite materials and leading batteries to maximize efficiency and reduce carbon output.

# Agenda

- Edge AI and Robotics Industry Outlook and Trend
- Unlocking New Applications; GenAI, LLM and Simulation with ISAAC ROS Platform
- Use Cases

# SUPER-HUMAN FARMING

**Verdant Robotics**

Verdant Robotics is building sustainable, high-fidelity farming: spatially, temporally, and physically working the farm at a precision, accuracy, frequency and scale never before possible.

By digitizing the farm at sub-millimeter scale, indexing it, and taking actions that unlock new value, we are helping transform how our food is grown while improving the lives of rural communities.

**Verdant is helping farmers reimagine their entire operation to improve efficiency, land stewardship, crop management and profitability.**

- Creates a digital twin of each plant for real-time decision making
- Reduces chemical usage by up to 95 percent
- Enables technology access to more farmers through RaaS model
- Unlocks the ability to discover transformative growing practices using microarray technology

VERDANT ROBOTICS

## WHY JETSON Orin?

“NVIDIA Jetson delivers all the power, software and versatility of a high-end workstation on a compact, edge-friendly package so we can operate in any field 24/7.”

-Lawrence Ibarria, Co-Founder and Chief Technology Officer at Verdant Robotics

# Edge Computing Helps Accelerate Superhuman Farming

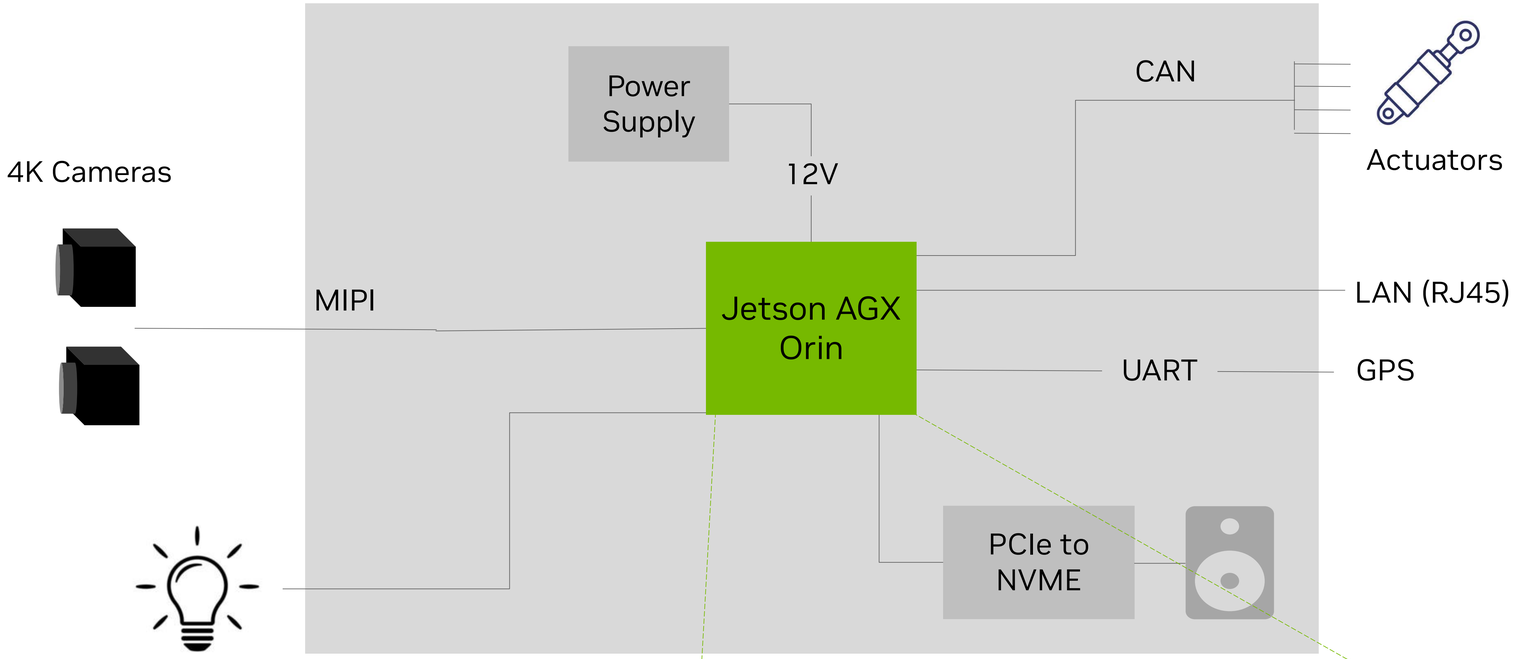


## Jetson Orin Value Prop

- 270+ TOPS to run concurrent ML models (classification and segmentation)
- Being able to process more pixels for each inference compared to Xavier

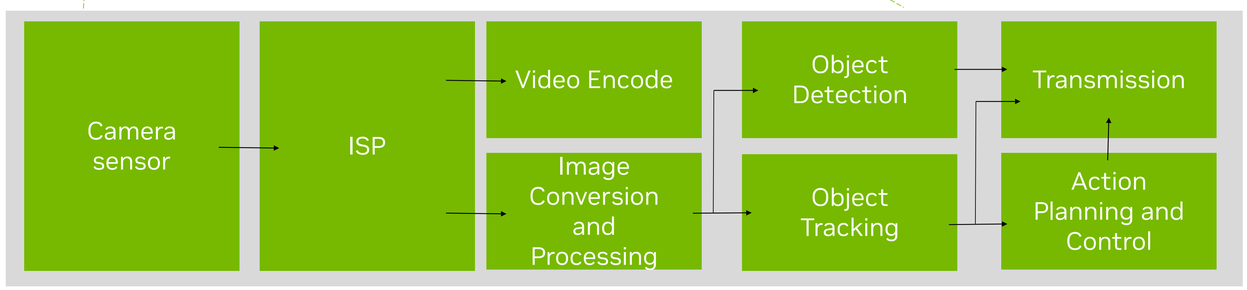


VERDANT ROBOTICS

# LAST-MILE DELIVERY

**Meituan**

The Meituan autonomous UGV safely navigates the neighborhood and delivers your favorite takeout to your home, on-demand with L4 autonomy.

To meet the increasing computing resource requirements for the L4 software stack, Meituan launched in-house development for the next-gen embedded SoC named MADC (Meituan Autonomous Driving Computer). Thanks to NVIDIA Orin modules, MADC currently supports 254~1016 TOPS and 20+Gbps sensor data, and has been tested since July 2022. The solution:

- Aggregates information from a range of sensors, improving safety and efficiency
- Supports the L4 delivery UGV with powerful CPU/GPU and high-fidelity sensors
- Enables next-gen autonomous UGV design with balanced computing performance, reliability, power consumption, and cost
- Prototypes technologies rapidly with extensive Jetson ecosystems and a user-friendly NVIDIA software toolchain

Meituan


## WHY JETSON Orin?

“NVIDIA Orin provides powerful CPU/GPU and integrated ecosystems, which significantly enhances the performance of the Meituan Autonomous Driving Computer.”

——Huaxia Xia, Head of Meituan Autonomous Delivery

# L4 UGV DELIVERS CONVENIENCE TO CONSUMERS

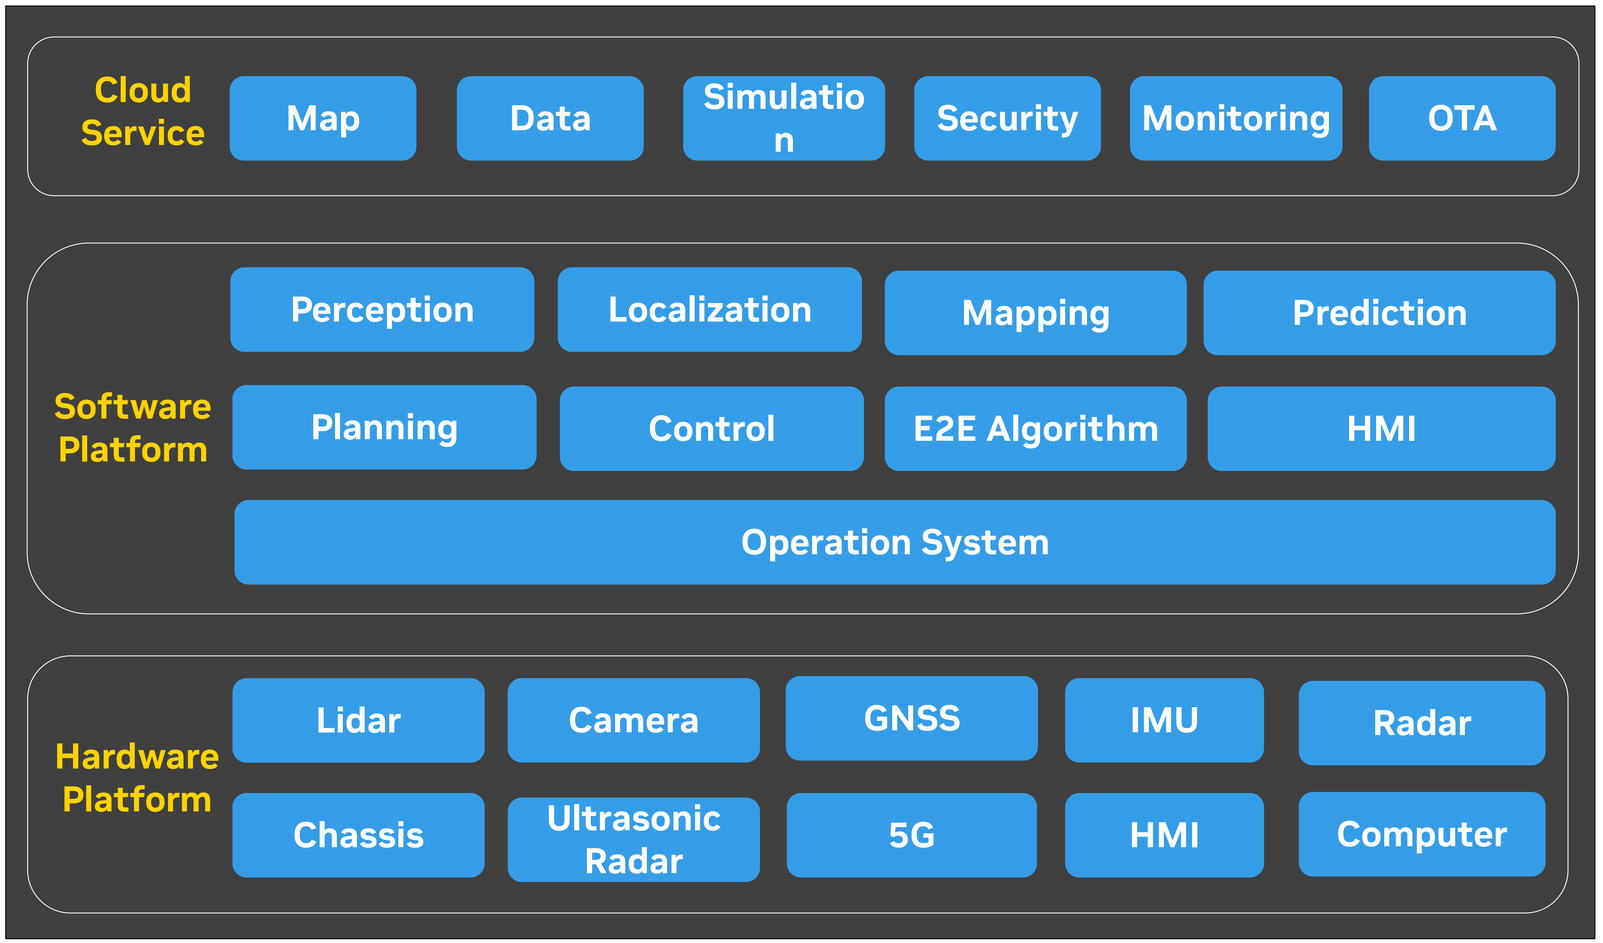


## Jetson Orin Value Prop

- 80% smaller form factor than x86 IPC-based solution
- 40%+ cost and 70%+ power reduction from x86 IPC-based solution, aiming 60watt per Orin module, support for 1~4 modules
- More performance per dollar compared to the last generation of Xavier modules
- Efficient DLA offload from GPU for supported ML operations
- Supports modern AI neural network-Transformer, Quantization Avoidance Training and 3D lidar model.

美团 Meituan

# HOME INTELLIGENT ROBOT

## HACHIBOT

The pace of life is getting faster every year—which can lead to greater stress. People sometimes feel depression and loneliness, and mental health is becoming a serious social problem. An intelligent dog-like robot from HachiBot can meet the various human needs, from entertaining the household to dancing with music.

HachiBot Captain is one-of-a-kind in the home-entertainment industry. The robot moves autonomously in complex indoor and outdoor environments, supporting pathfinding and obstacle avoidance. It can interact with people and make the right decisions to meet people's needs. The robot's various sensors and great computing power lets it concurrently process computer vision and audio data in real time. Thanks to AGX Orin, Hachibot can put super AI compute power in a small robot body. It delivers:

- A super perception planning system, 360-degree all-in-view, to identify the surrounding environment anytime, anywhere
- Intelligent human-computer interaction, able to recognize scenes, different people and objects, various motions, and action commands.
- Indoor and outdoor all-terrain adaptability, stable gait with lower cadence and better reliability in the home environment

HACHIBOT
BRING WORLD INTO THE ROBOT AGE

## WHY JETSON Orin?

"NVIDIA Jetson Orin is a powerful AI computer which provides the ideal solution for a new age of robotics. We are very glad to have Jetson Orin as the core of our home intelligent robot."

- ZhenYu Chen, Chief Executive Officer (CEO) at HachiBot

# RUNNING COMPLEX AI MODELS IN A SMALL HOME INTELLIGENT ROBOT

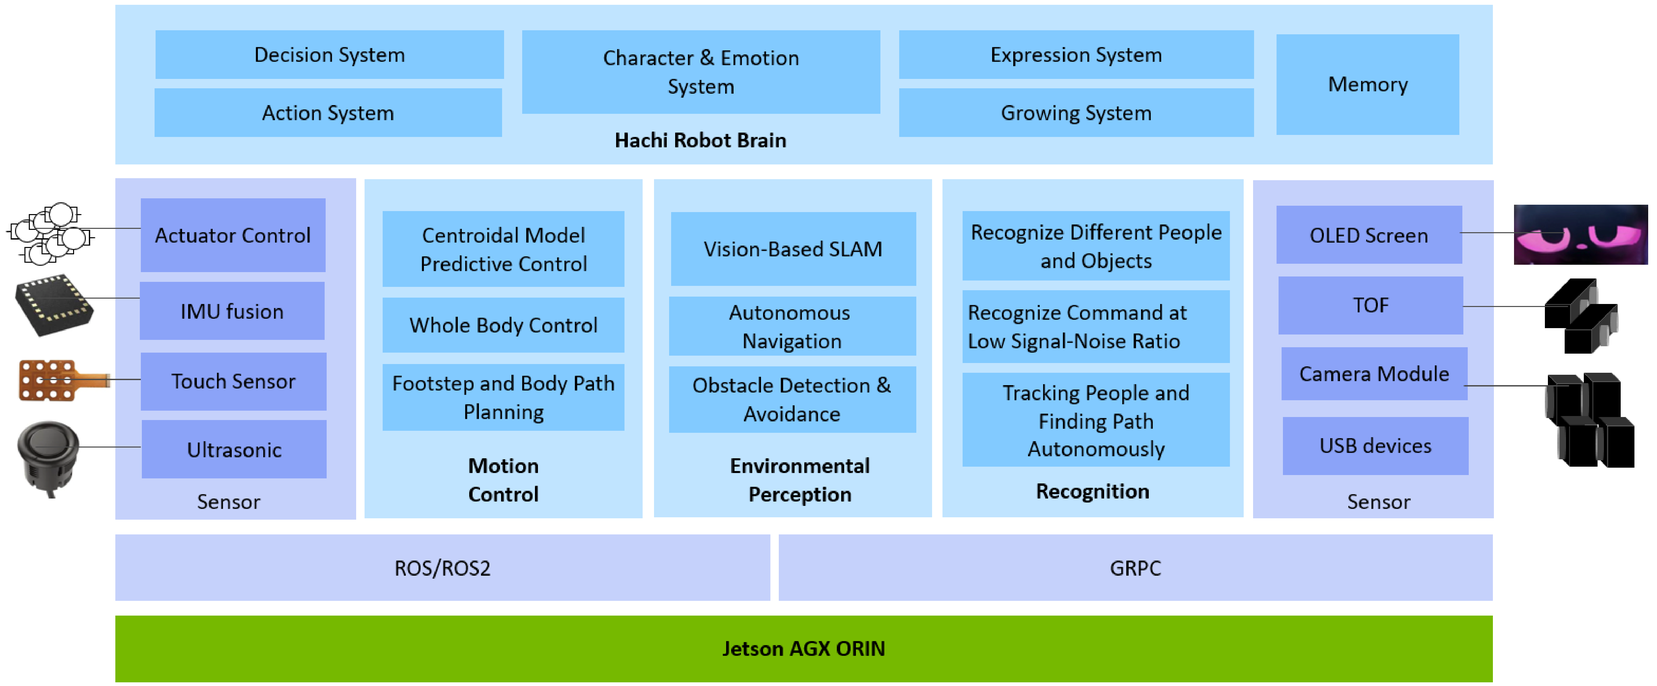


## Jetson Orin Value Prop

- Support for ROS/ROS2 natively with highly optimized and hardware-accelerated packages from Isaac ROS
- The ability to process multi-modal applications without sacrificing performance with a specialized AI processor such as DLA, PVA, VIC, and OFA
- 64GB LPDDR5 memory and high-speed processing power to remove bottlenecks

# ASTRONOMICAL DATA PROCESSING

## OURSKY

Thousands of asteroids and hundreds of millions of undetected pieces of space junk exist in our local space environment, threatening our future in space and on the ground. Single telescopes aren't enough to accurately monitor and track these objects due to weather and cloud coverage.

OurSky is building a global platform to solve these problems by coordinating the space data collection, processing, and applications that will enable the entire sky to be observed globally 24/7/365 with better than human analysis and interpretation of the data in real time.

Building on the Jetson framework, the OurSky platform:

- Provides a next-gen telescope control and data acquisition hardware solution (the OurSky Node Controller) that can be attached to any amateur or professional telescope system
- Performs real-time, AI-driven analysis of the data being ingested at the edge by each node in the telescope network
- Provides a platform for new AI applications used in analysis of this space data, including object detection, object classification, orbital prediction, and real-time calibrated photometry

OURSKY

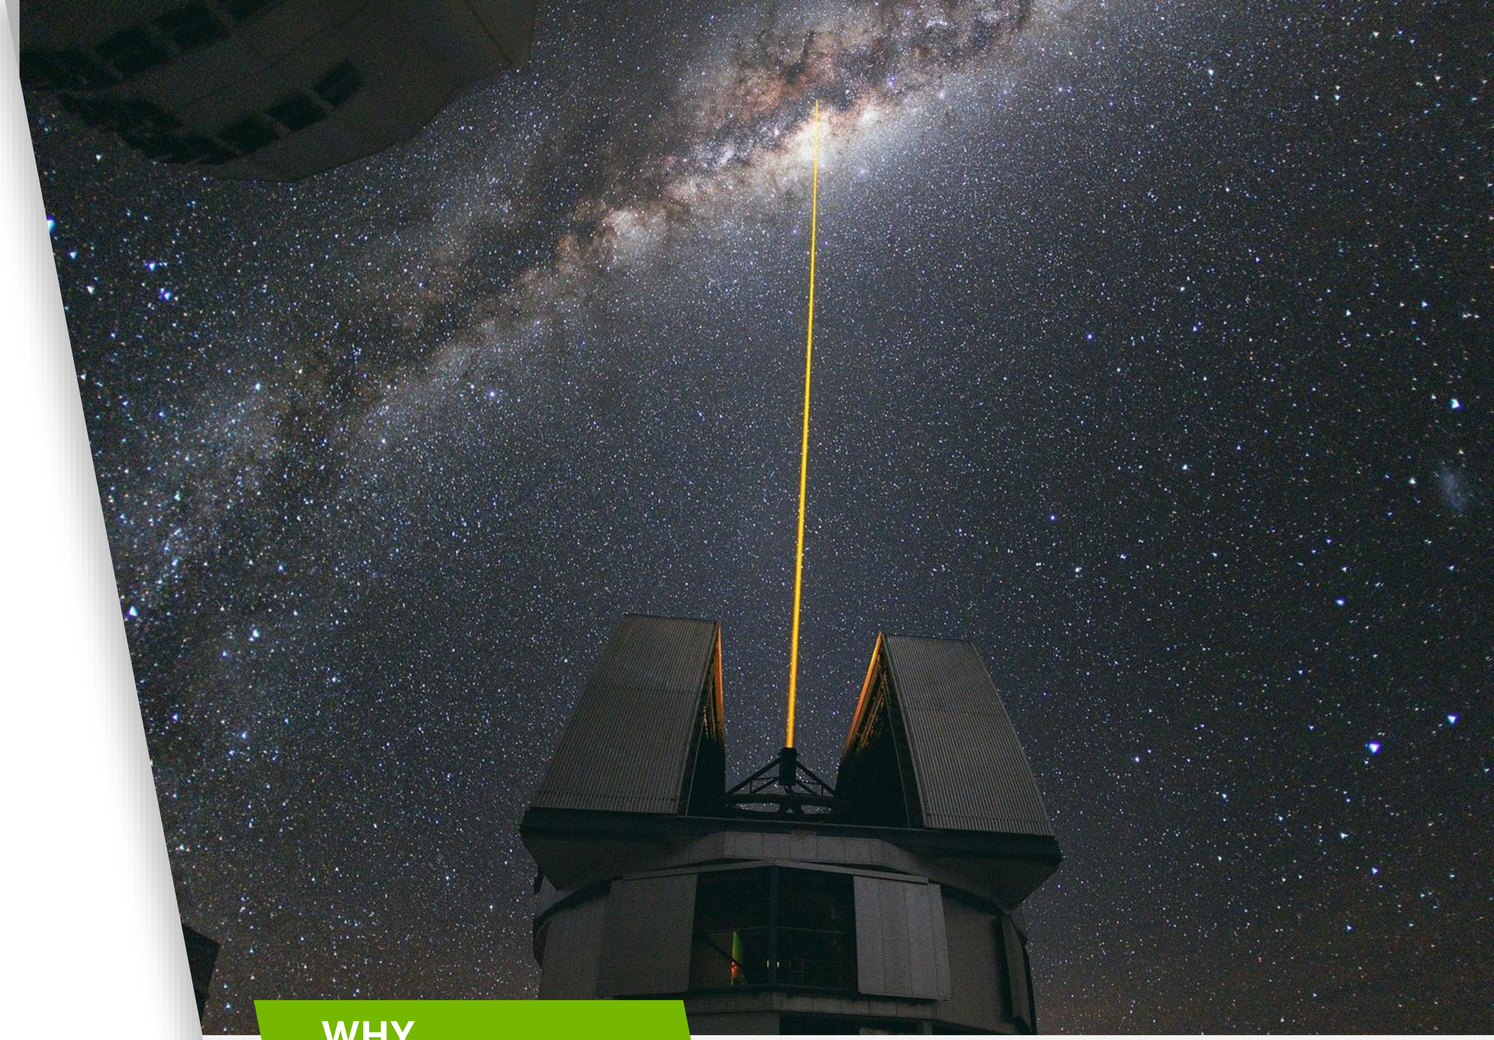

## WHY JETSON?

"The NVIDIA Jetson Orin Platform enables OurSky to leverage powerful AI to process and analyze a monumental amount of space data in real-time, right at the edge as it is being collected. Given the globally distributed nature of our telescope network, it would not be feasible to move all of this data to the cloud. Jetson uniquely enables OurSky to address this existential problem for humanity."

- Alex Hawkinson, Founder of OurSky

# Edge Processing on the OurSky Node Controller

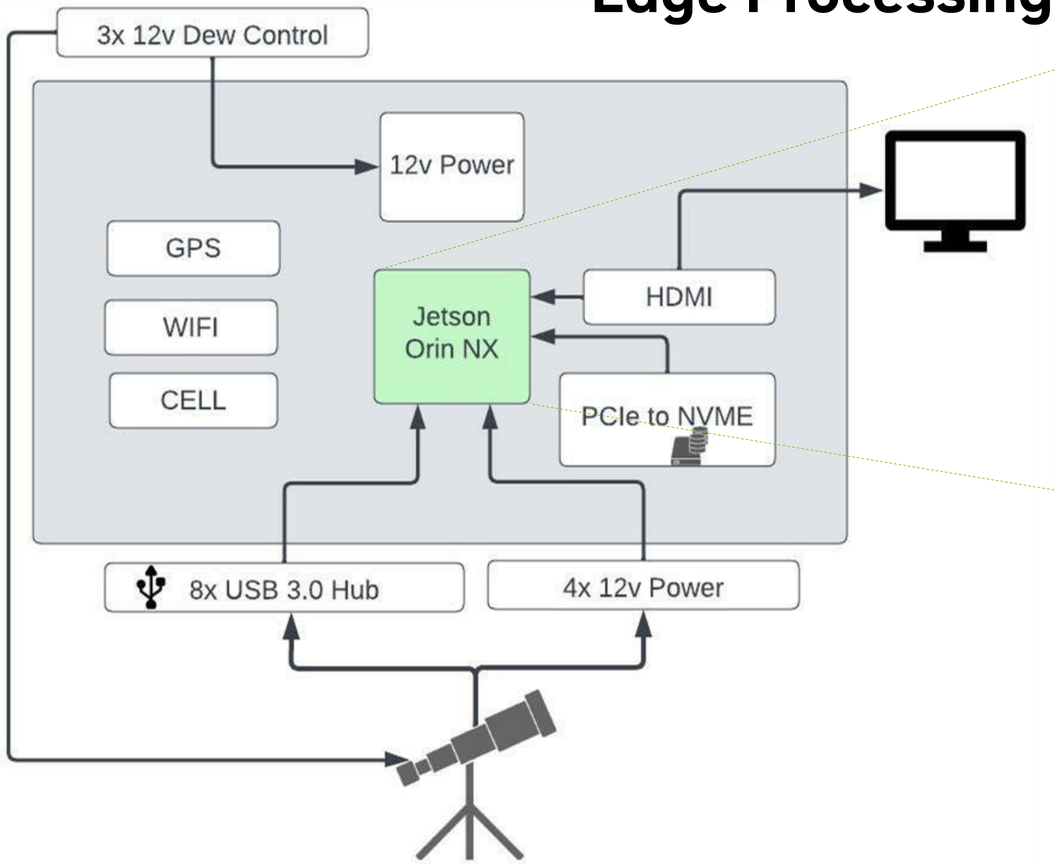


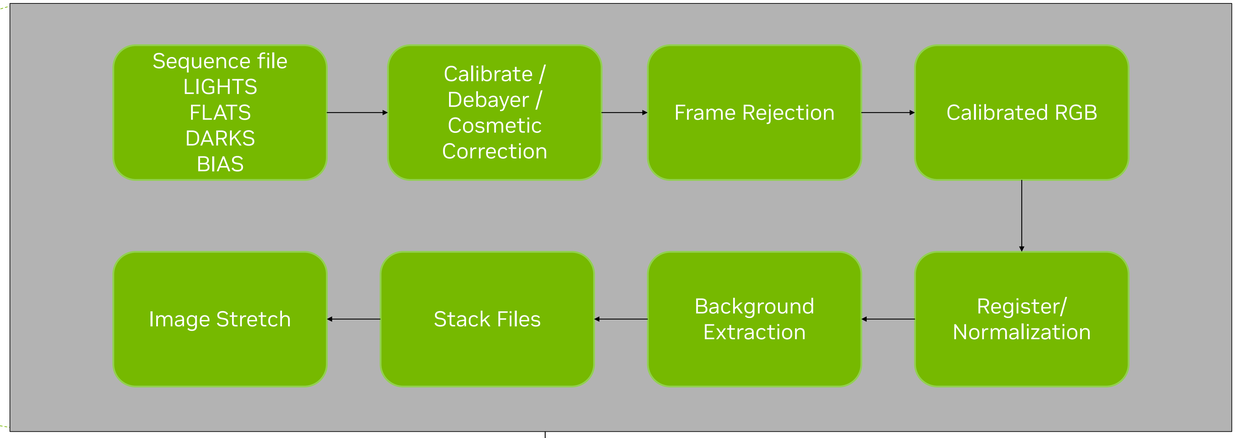


**Extracted Data**

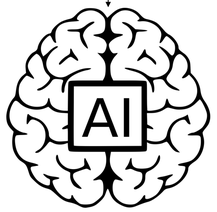


## Jetson Orin Value Prop

- An AGX Orin Developer Kit that can emulate Orin NX and accelerate the time-to-market
- Performance that enables OurSky to use powerful AI to process and analyze data right at the edge, in real-time
- A faster AI pipeline to help analyze surrounding topics such as asteroid detection, space junk detection, orbital prediction, photometry, and astronomical events

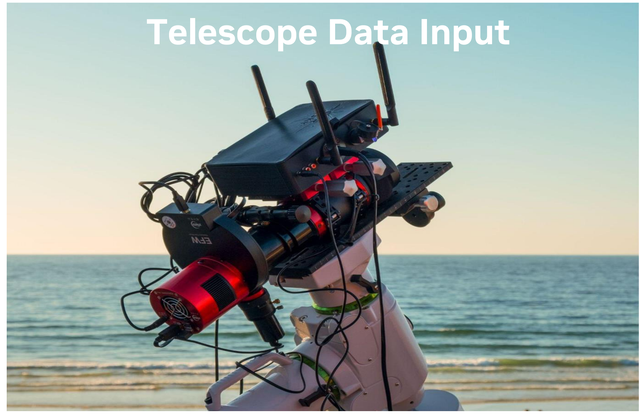


Satellites/ Space Junk

Comets

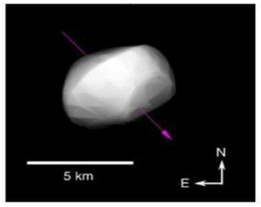


Near-Earth Objects

Emerging Events

Astro photos

# Edge Processing on the OurSky Node Controller

Option #2

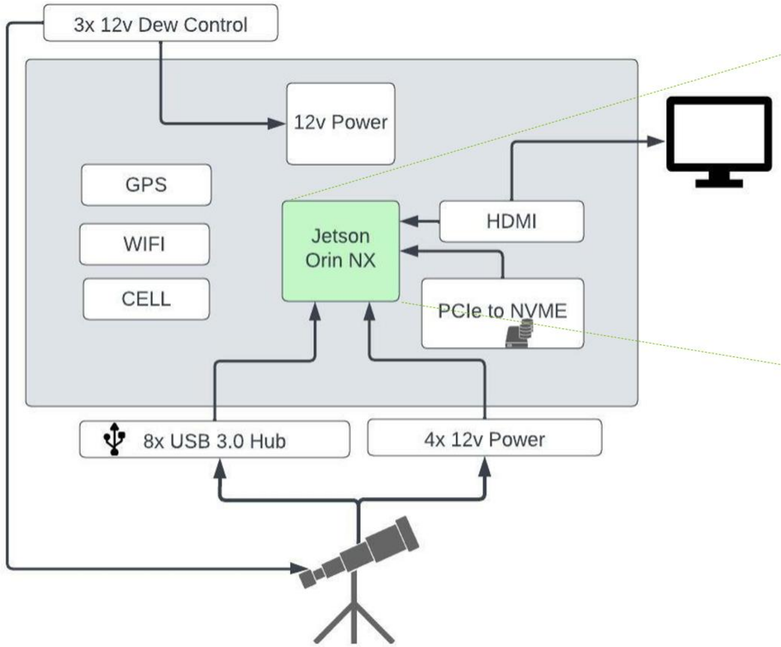


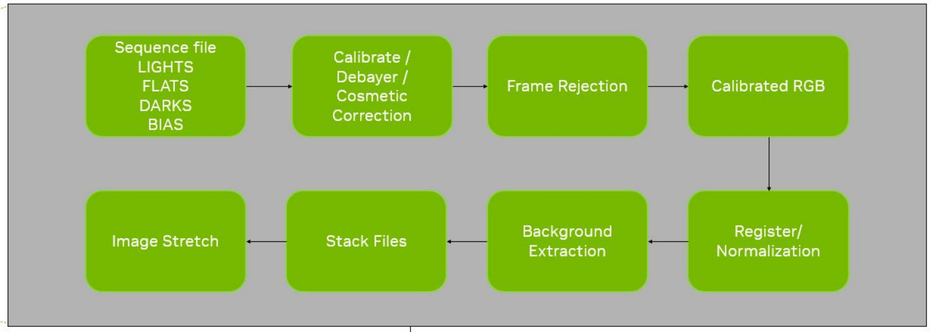


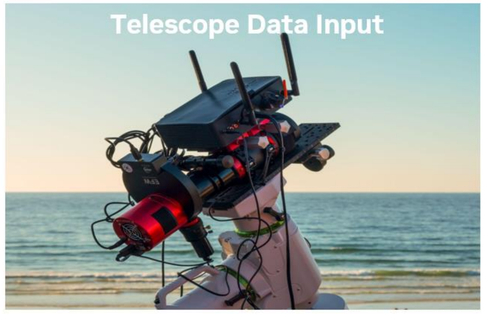


**Extracted Data**

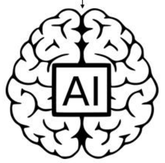


## Jetson Orin Value Prop

- An AGX Orin Developer Kit that can emulate Orin NX and accelerate the time-to-market
- Performance that enables OurSky to use powerful AI to process and analyze data right at the edge, in real-time
- A faster AI pipeline to help analyze surrounding topics such as asteroid detection, space junk detection, orbital prediction, photometry, and astronomical events

Satellites/ Space Junk

Comets

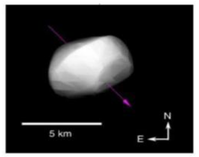


Near-Earth Objects

Emerging Events

Astro photos

# End-to-End Vision AI Development

Fast-track data generation, AI model creation, app development and deployment.

## Generate

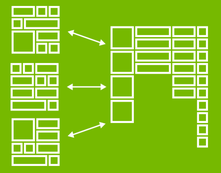

### NVIDIA Omniverse Replicator

generate synthetic training data

## Train

### NVIDIA Pre-Trained Models

use state of the art models out of the box

### NVIDIA TAO Toolkit

for simplified and accelerated model tuning

## Build

### NVIDIA Video Storage Toolkit

easily manage cameras and storage

### NVIDIA Deepstream, TensorRT, and Triton

optimize inference throughput

### Metropolis Microservices

build cloud-native applications

## Deploy

### NVIDIA Fleet Command

Streamline edge deployment and app management

### Any Cloud

deploy on any public and private cloud

# AGX Orin Reinvents the Far Edge

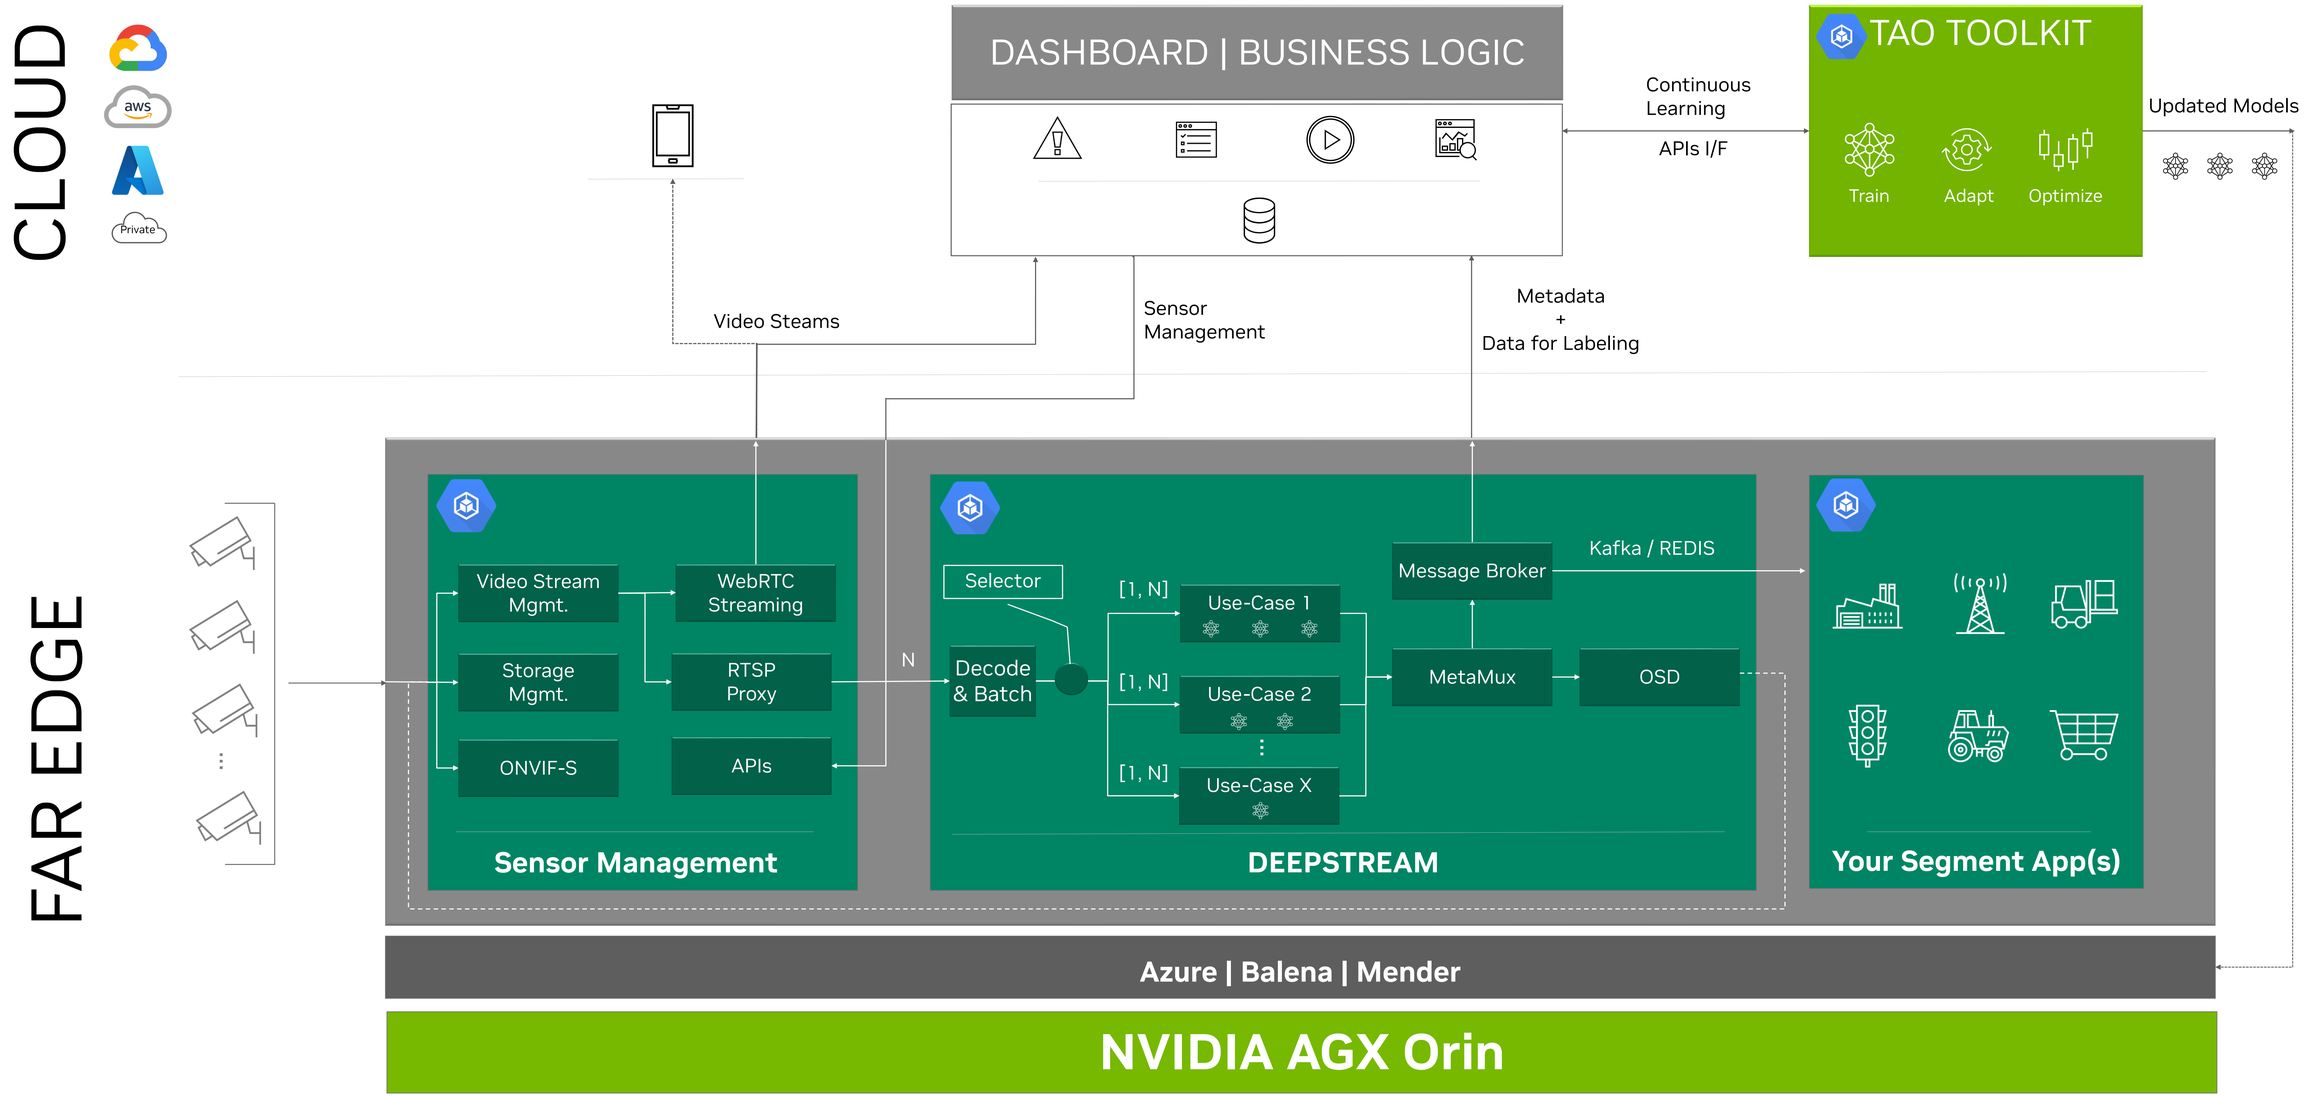


https://github.com/NVIDIA-AI-IOT/deepstream_parallel_inference_app

NVIDIA

# Connect with NVIDIA's Robotics Ecosystem


- To learn more about NVIDIA's robotics platform: https://developer.nvidia.com/isaac
- To learn more about Generative AI at NVIDIA: https://www.nvidia.com/en-us/ai-data-science/generative-ai/
- To get started with Jetson Generative AI Lab: https://www.jetson-ai-lab.com/
- To join NVIDIA's startups program Inception: https://www.nvidia.com/en-us/startups/