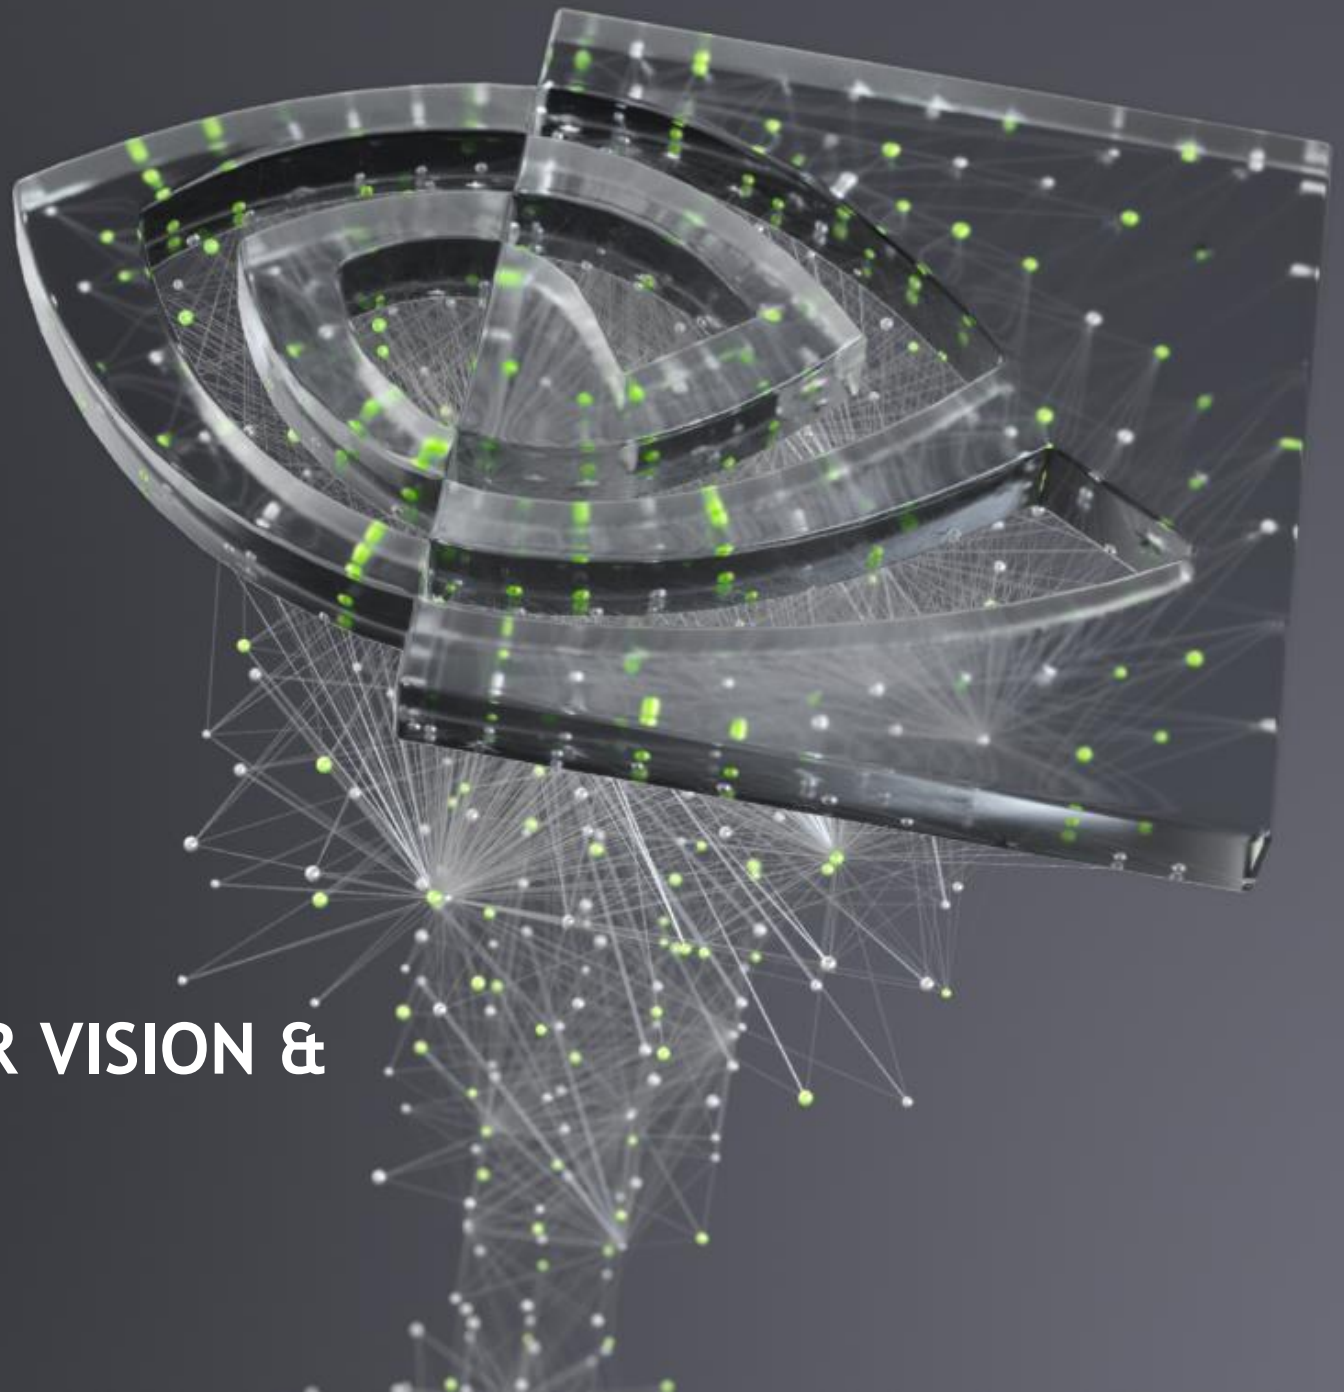

DEEP LEARNING FOR COMPUTER VISION & BIOMEDICAL IMAGE ANALYSIS

Jesse Tetreault, Solutions Architect, Nvidia

NIH.AI Workshop

October 19, 2021

AGENDA

1:05 PM **The Anatomy of a Deep Learning Problem**

Artificial Intelligence, Machine Learning, & Deep Learning

Deep Neural Networks

Deep Learning Techniques in Biomedical Image Analysis

1:40 PM **Instruments at Hand: Tools and Frameworks for AI Practitioners**

GPU Computing: A Computational Microscope

Deep Learning Frameworks & Software: MONAI

Nvidia Clara

2:30 PM **Lightning Talks from Intramural NIH Researchers**

LIGHTNING TALK SCHEDULE

"NCI Imaging Data Commons,"
Keyvan Farahani, Ph.D., NCI

"Recent AI Applications in the NCI CCR Artificial Intelligence Resource"
Stephanie Harmon, Ph.D., NCI

"Strategies to Mine Lymph Nodes from Incompletely Annotated PACS Data"
Tejas Mathai, Ph.D., CC

"AI for Imaging Metabolic Disorders"
Nader S. Metwalli, Ph.D., NIDDK

"Volume Electron Microscopy: Advances and Challenges in DL-based Segmentation"
Kedar Narayan, Ph.D., NCI

"Enhancing Fluorescence Microscopy with Computation"
Hari Shroff, Ph.D. - NIBIB

ARTIFICIAL INTELLIGENCE, MACHINE LEARNING, & DEEP LEARNING

ARTIFICIAL INTELLIGENCE, DEEP LEARNING & MACHINE LEARNING

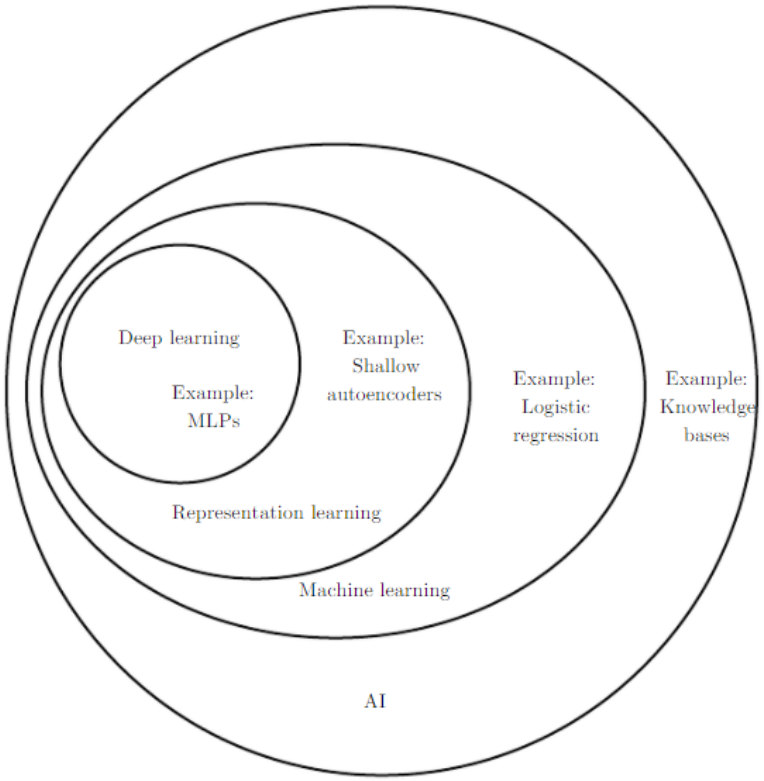

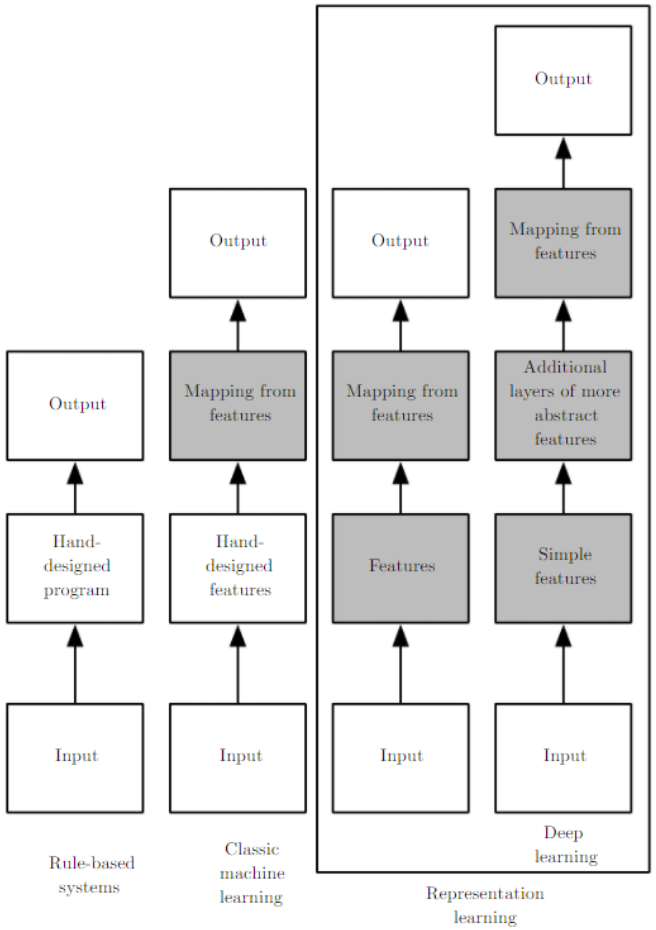

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

TYPES OF ML/DL

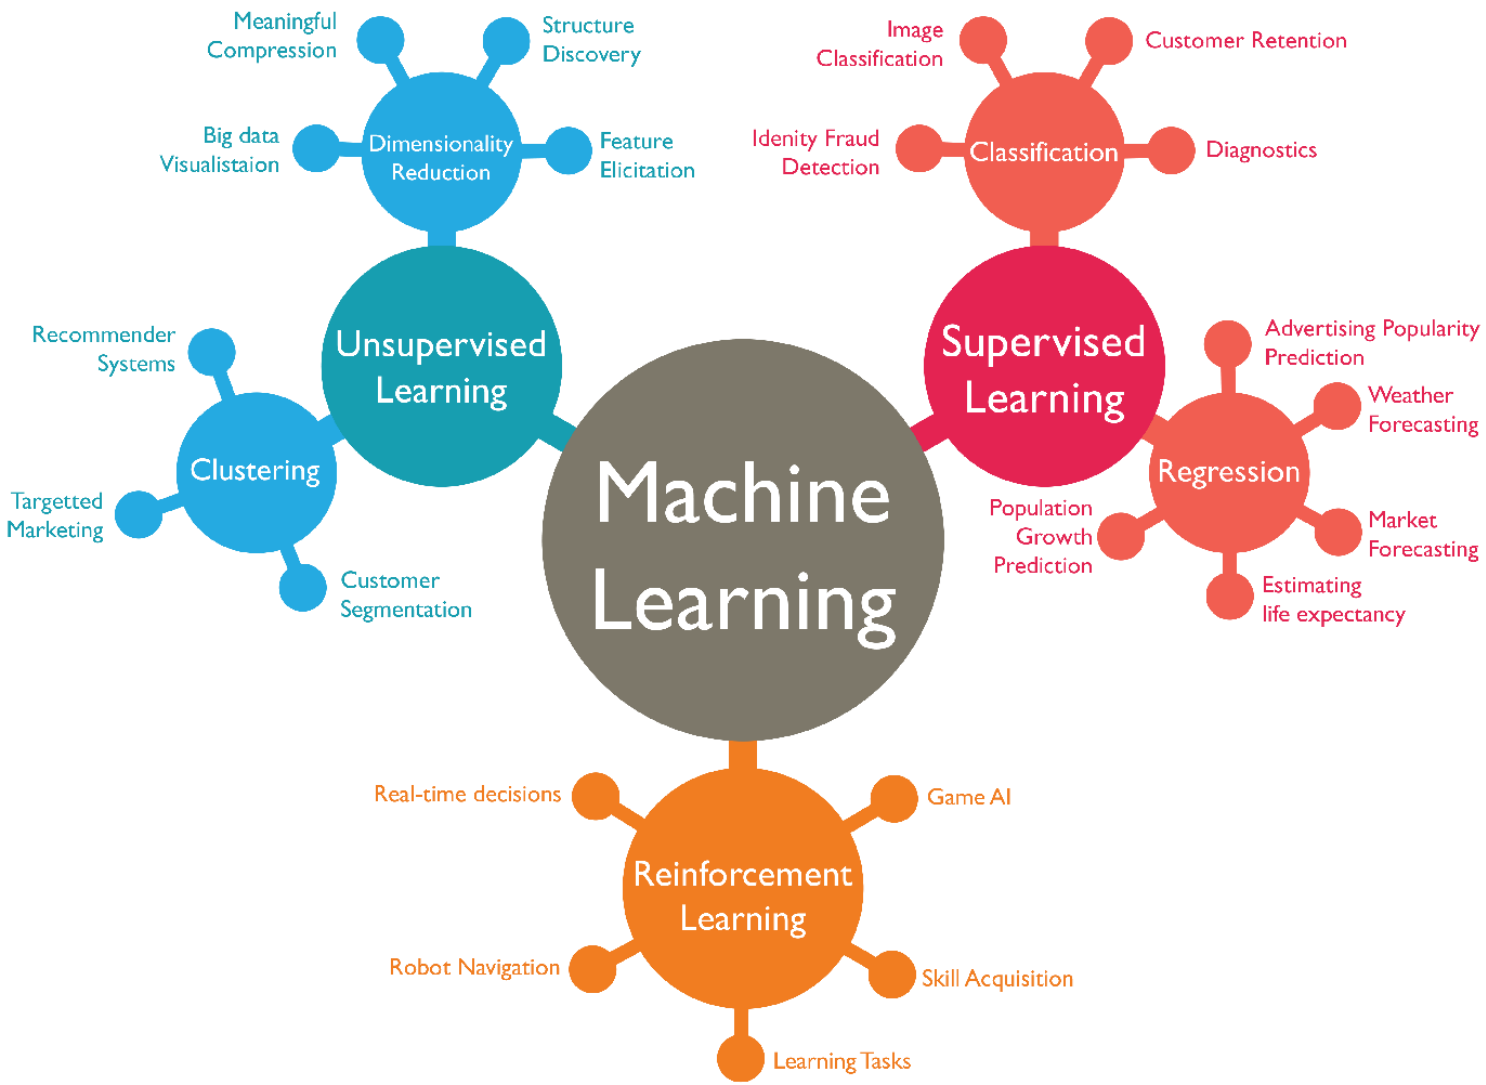

SOME KEY DECISIONS TO MAKE

FACTOR	DESCRIPTION
DL Challenge	Supervised or unsupervised, classification or regression, # of labels?
Architecture	What is the simplest architecture I can use?
Training Model	How am I going to tune my neural net? Kinds of non-linearity, loss function and weight initialization? Best training framework?
Data Quantity	How much data will be sufficient to train my model? How do I go about finding that data and is it evenly balanced?
Data Quality	Is my data directly relevant to the problem & real-world data?
Data Labels	Is training data is labeled same as raw data sets, how do I 'featurize'?
Data Similarity	Is data same length vectors or does it require pre-processing?
Data Storage & Access	Where is it stored, locally and or network? How do I plan to extract, transform and load the data (ETL)?
Infrastructure	Cloud, on-premise, Hybrid? GPUs, CPUs, or both? Single or distributed systems? Integration with languages, enterprise applications, databases?

DEEP NEURAL NETWORKS

CAMBRIAN EXPLOSION

Convolutional Networks

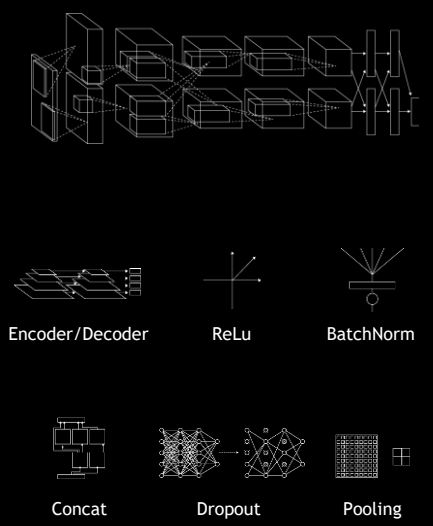

Recurrent Networks

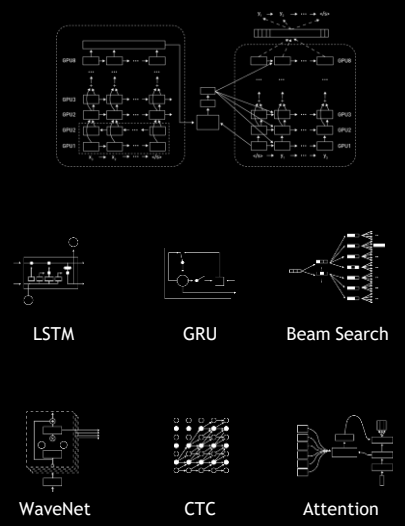

Generative Adversarial Networks

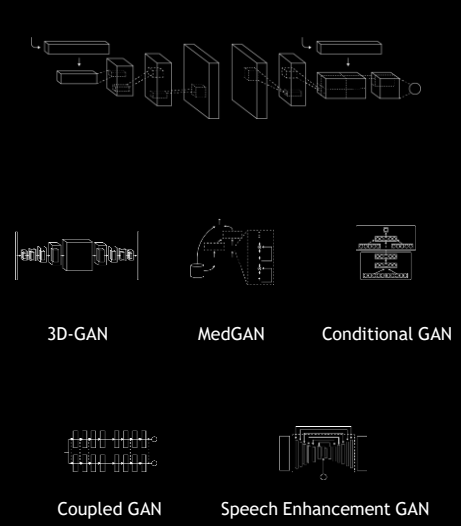

Reinforcement Learning

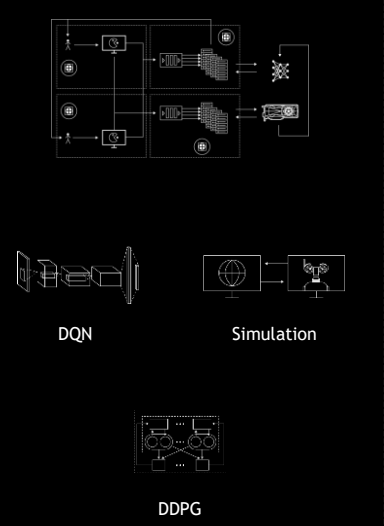

New Species

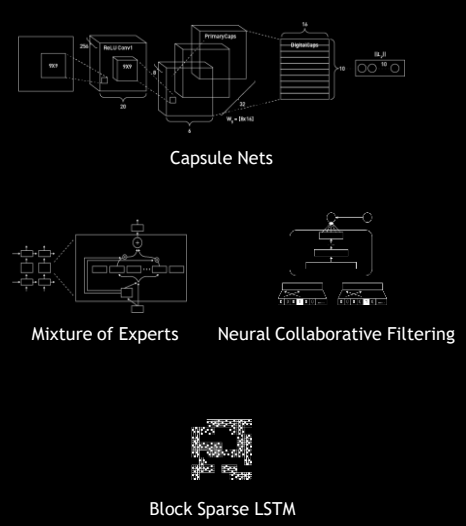

INGESTING IMAGE DATA

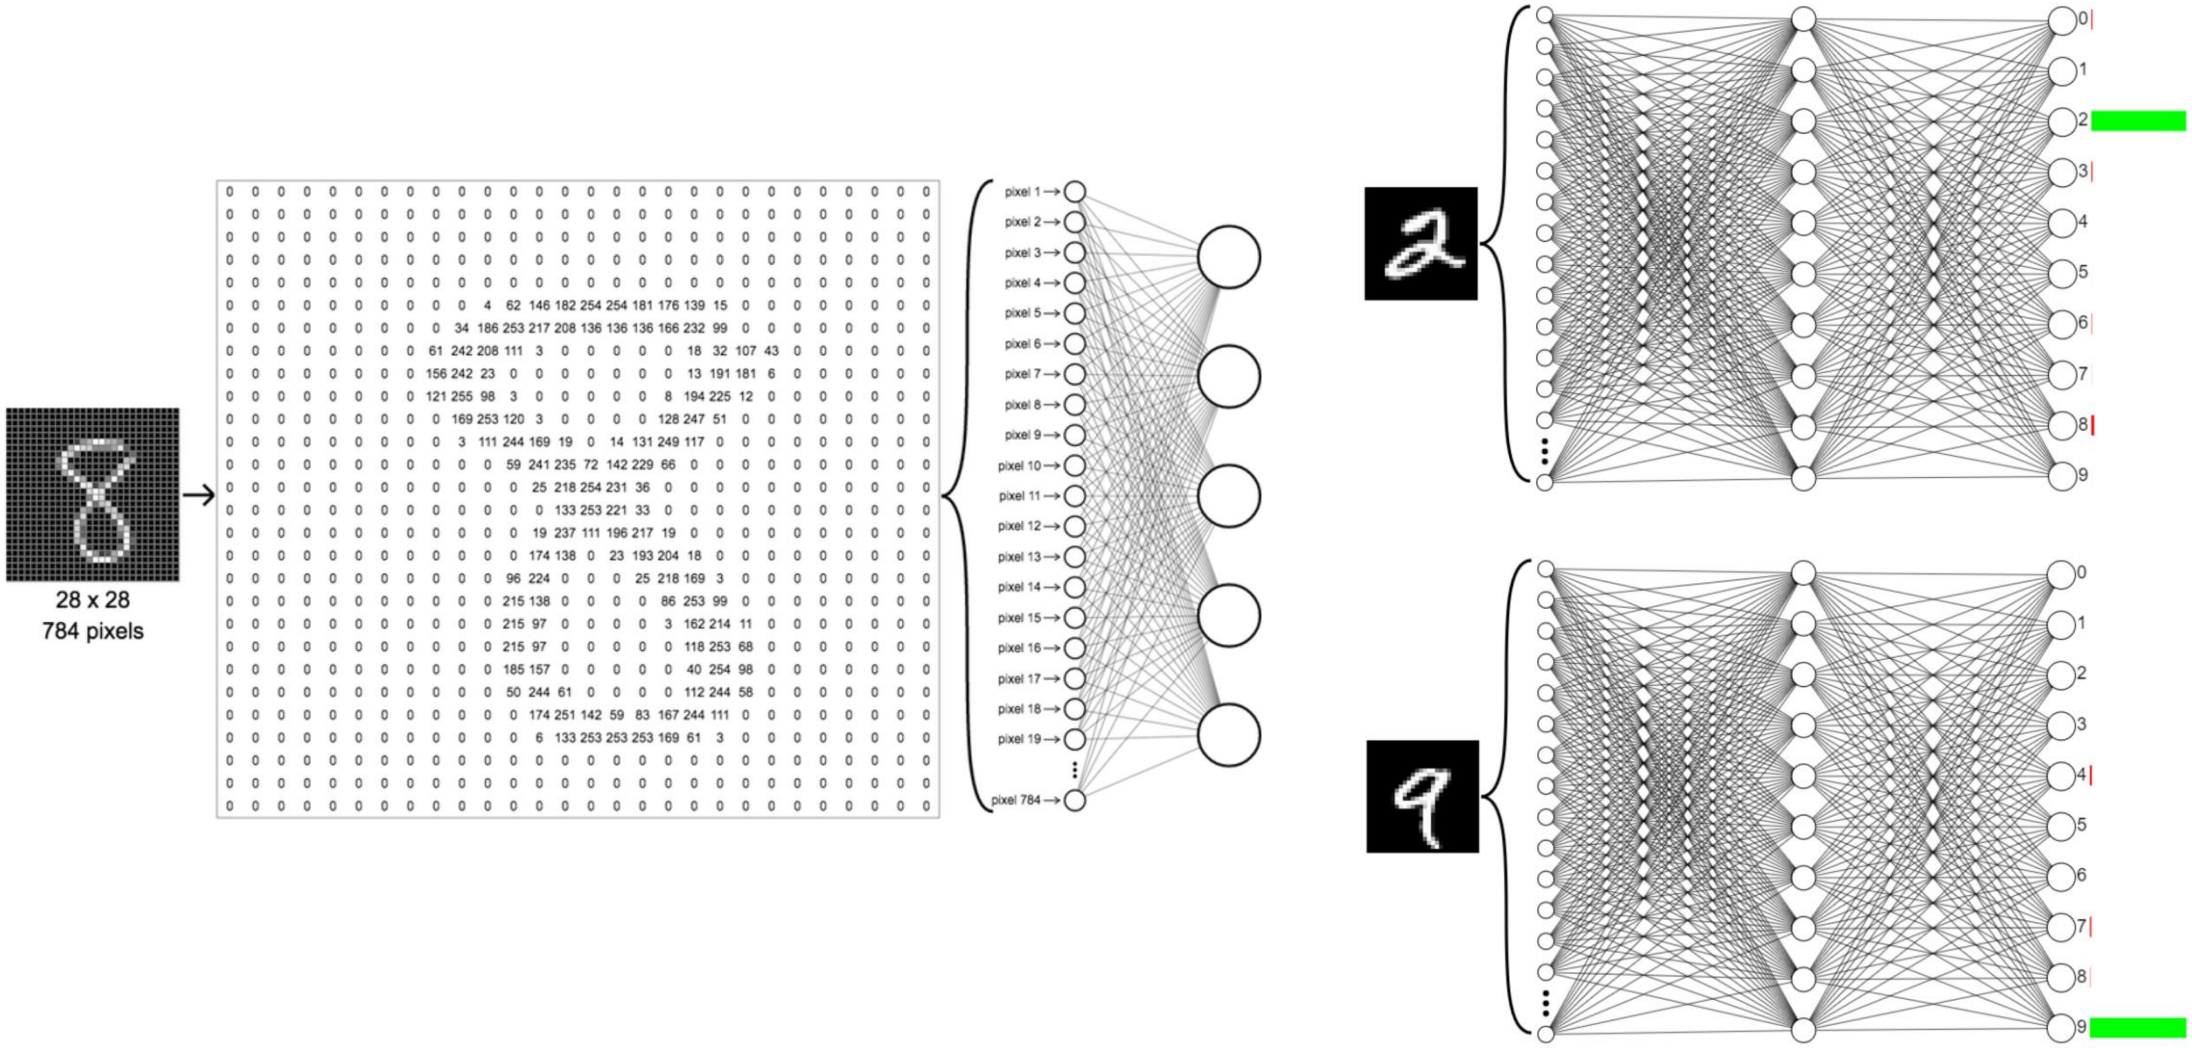

https://ml4a.github.io/ml4a/neural_networks/

CONVOLUTIONAL LAYERS

Title with Images

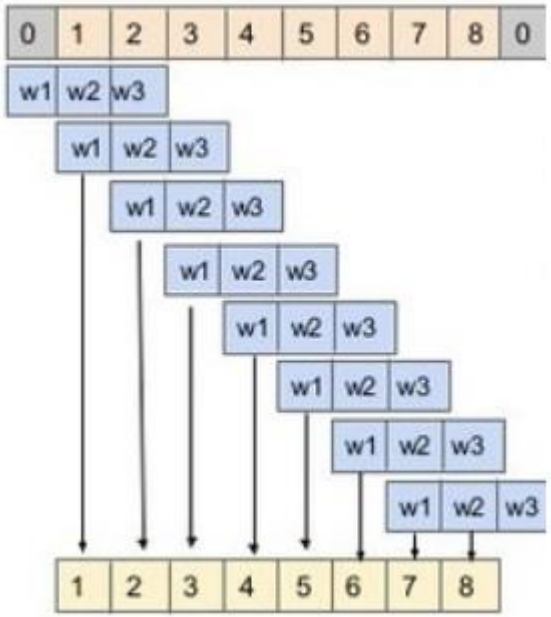

1D Convolutions
Supporting text here

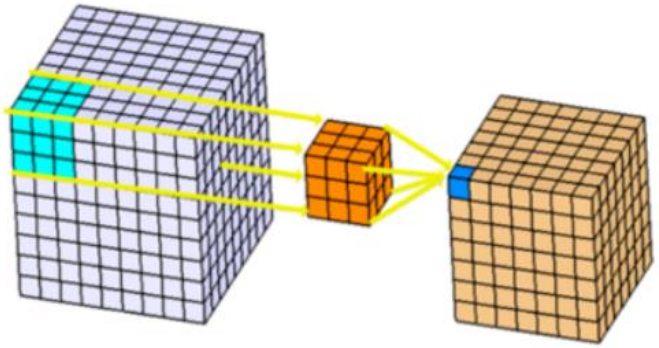

3D Convolutions
Supporting text here

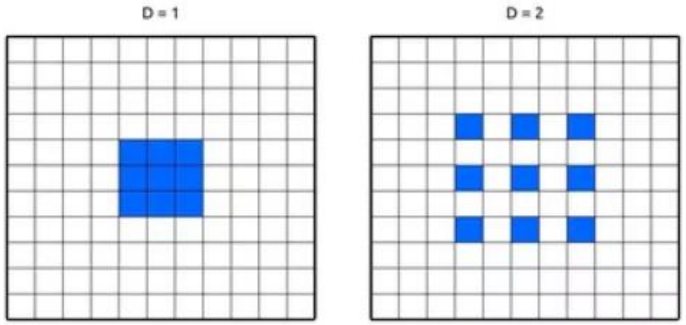

Dilated Convolutions
Supporting text here

https://www.kaggle.com/shivamb/3d-convolutions-understanding-use-case

IMAGE CLASSIFICATION EXAMPLE (ALEXNET)

Raw data

Low-level features

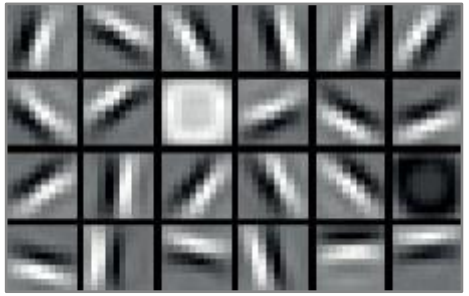

Mid-level features

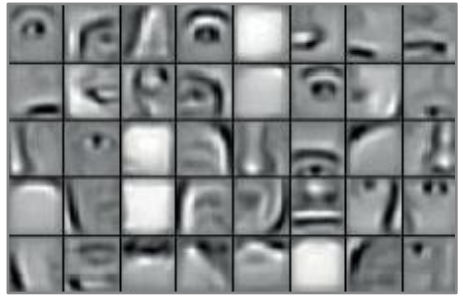

High-level features

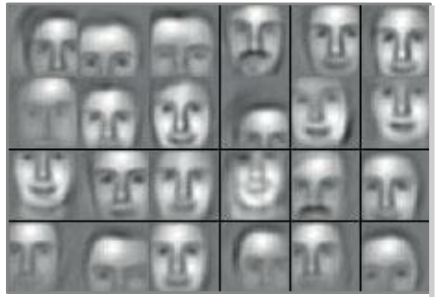

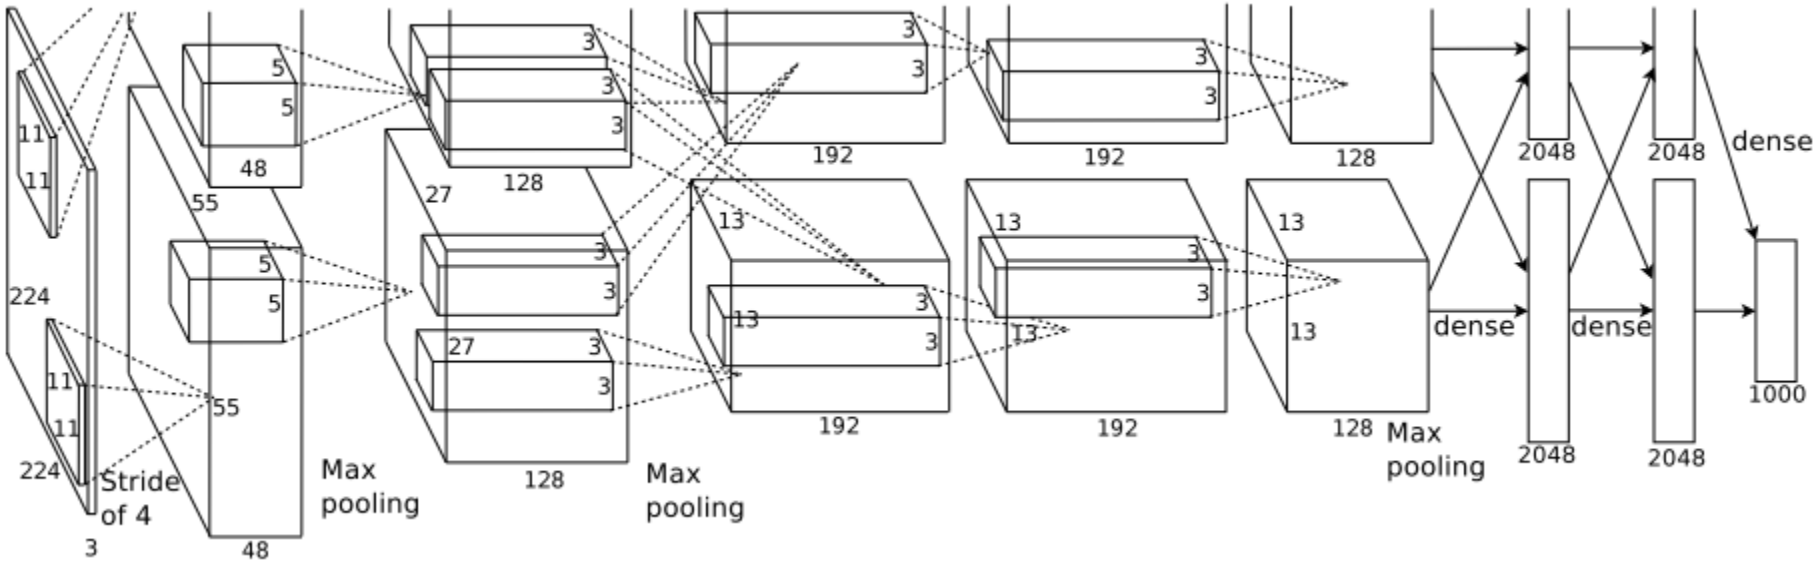

Task objective
Identify face

Training data
10-100M images

Network architecture
AlexNet
5 convolutional layers
3 fully-connected layers
61M parameters
650k ‘neurons’

Learning algorithm
Gradient descent

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems* 25 (2012): 1097-1105.

FEATURE VISUALIZATION

How neural networks build up their understanding of images

Edges (layer conv2d0) **Textures** (layer mixed3a) **Patterns** (layer mixed4a) **Parts** (layers mixed4b & mixed4c) **Objects** (layers mixed4d & mixed4e)

Feature visualization allows us to see how GoogLeNet[1], trained on the ImageNet[2] dataset, builds up its understanding of images over many layers. Visualizations of all channels are available in the appendix.

https://distill.pub/2017/feature-visualization/

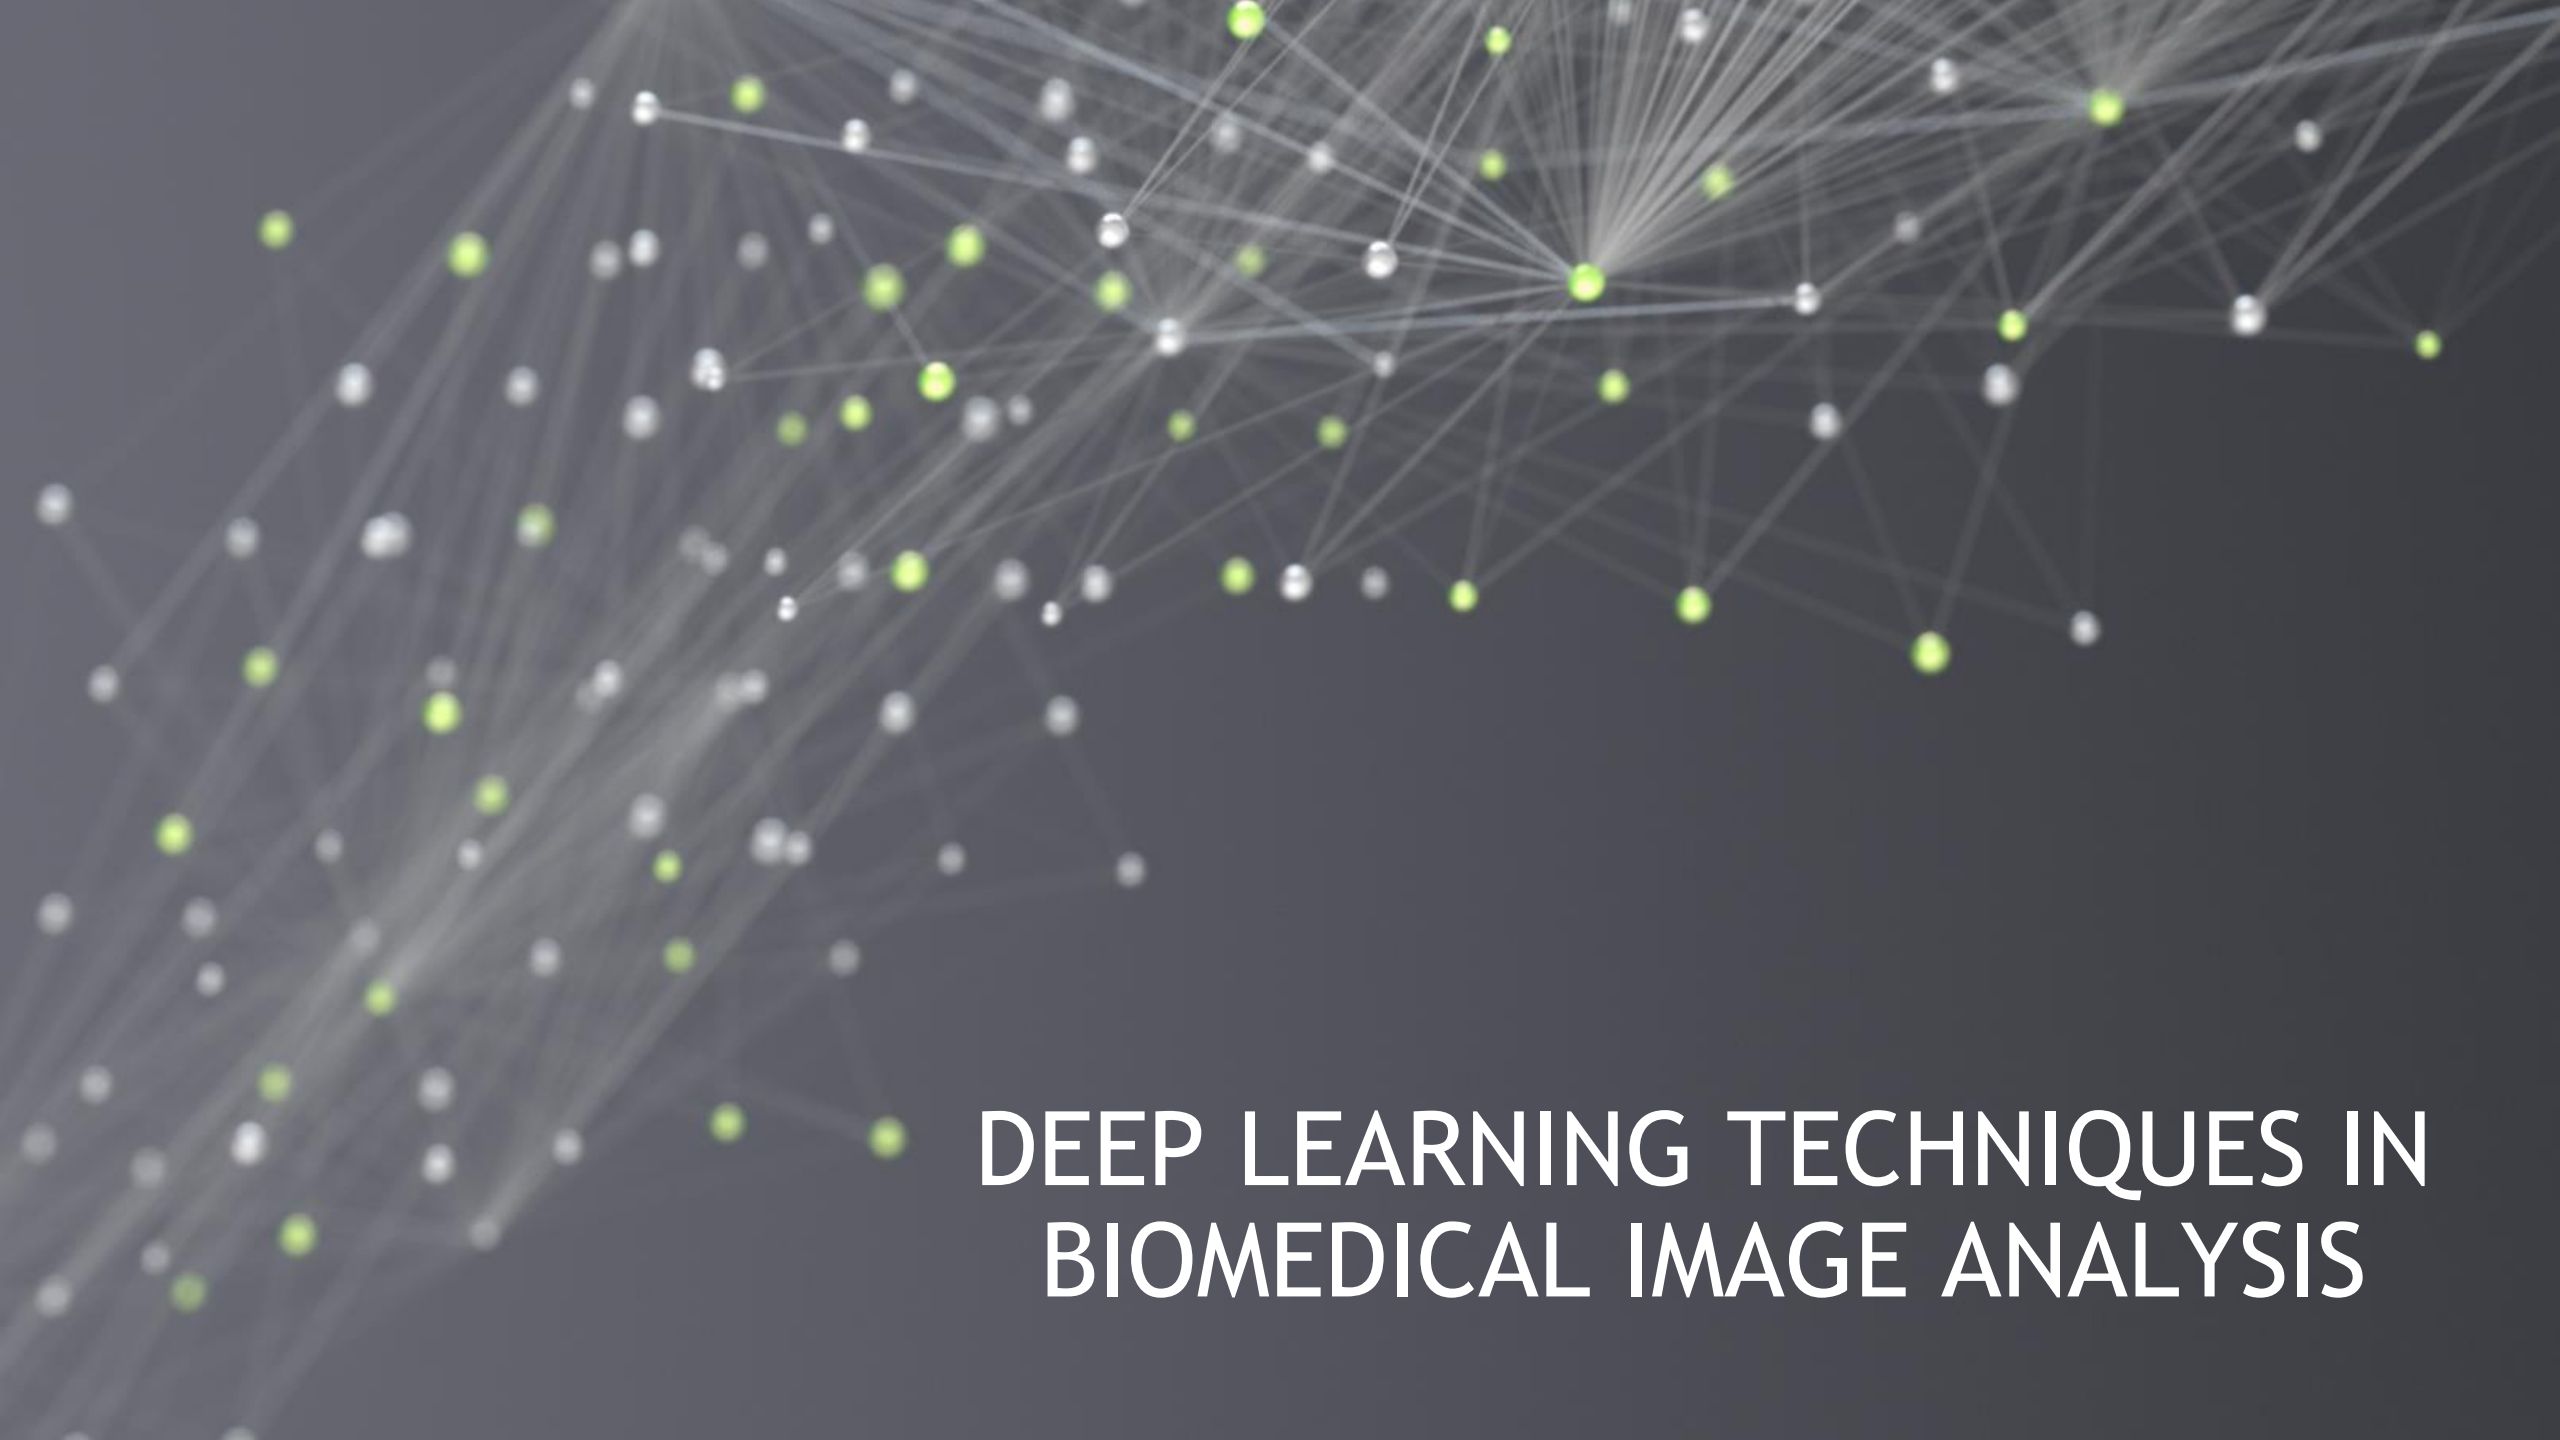

DEEP LEARNING TECHNIQUES IN BIOMEDICAL IMAGE ANALYSIS

BIOMEDICAL IMAGING MODALITIES

X-Ray

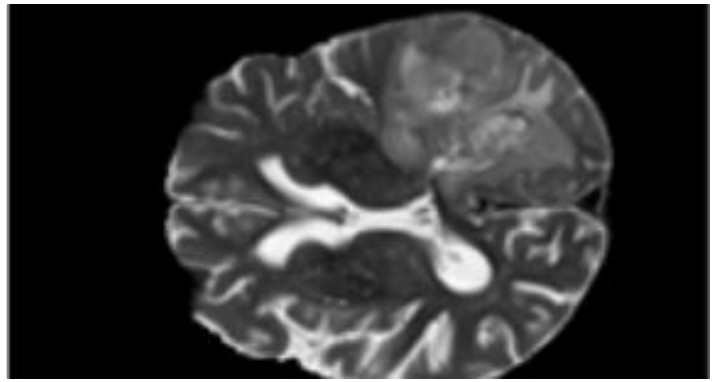

MRI

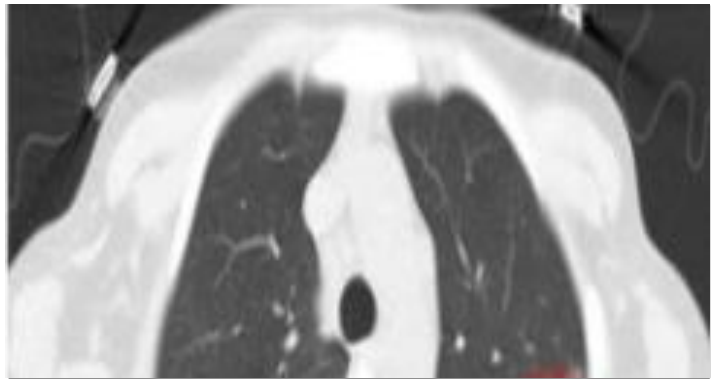

CT

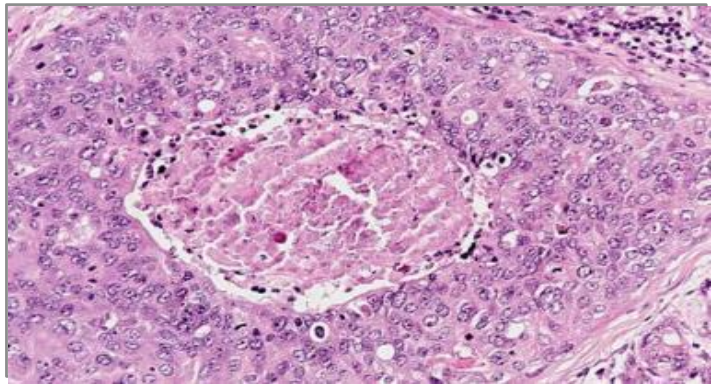

Light Microscopy

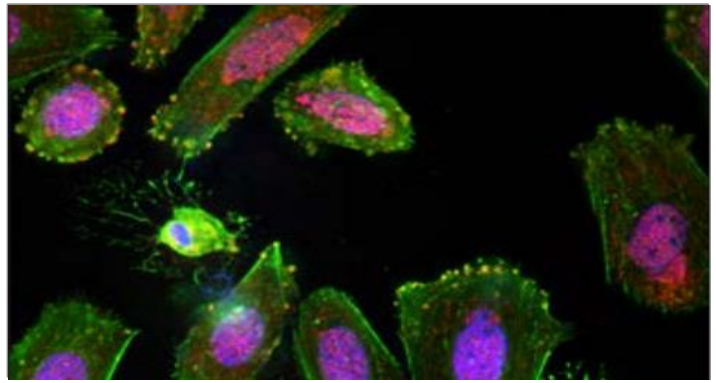

Electron Microscopy

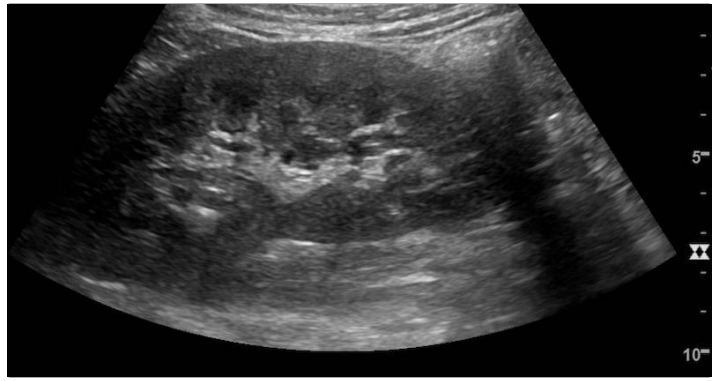

Ultrasound

MULTIPLE INSTANCE LEARNING FOR WSI

MICCAI 2021

- Multi-instance learning (MIL)-based classification of whole slide images (WSI)
  - Break the image into a set of patches
  - Generate the prediction based on these patches
- Account for dependencies between instances during training by embedding **self-attention blocks** to capture dependencies between instances
- Generate the mask labels based on the attention heat maps
- Achieved SOTA results on the Kaggle PANDA challenge, largest public WSI dataset (>11,000 images)

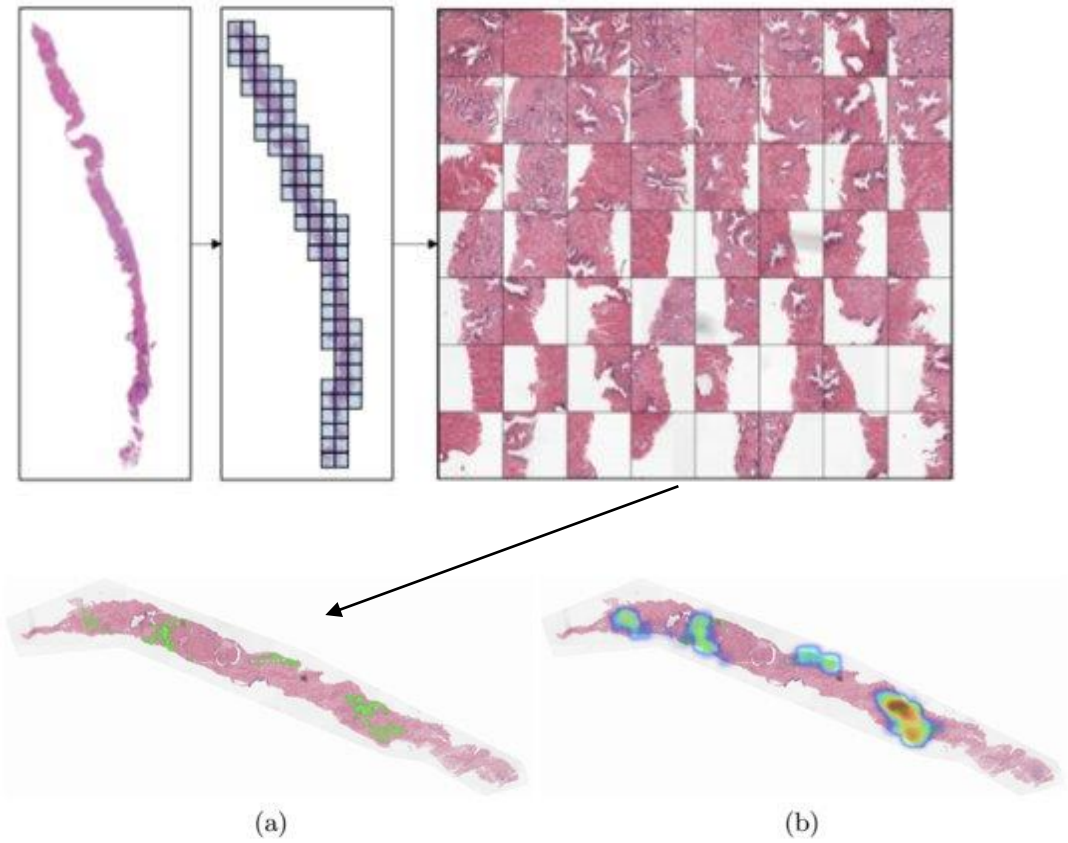

(a) (b)

Pos/Neg patches → Prediction

Myronenko, Andriy, et al. "Accounting for Dependencies in Deep Learning Based Multiple Instance Learning for Whole Slide Imaging." *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Cham, 2021.

TRANSFORMERS FOR 3D MEDICAL IMAGE SEGMENTATION

UNetR

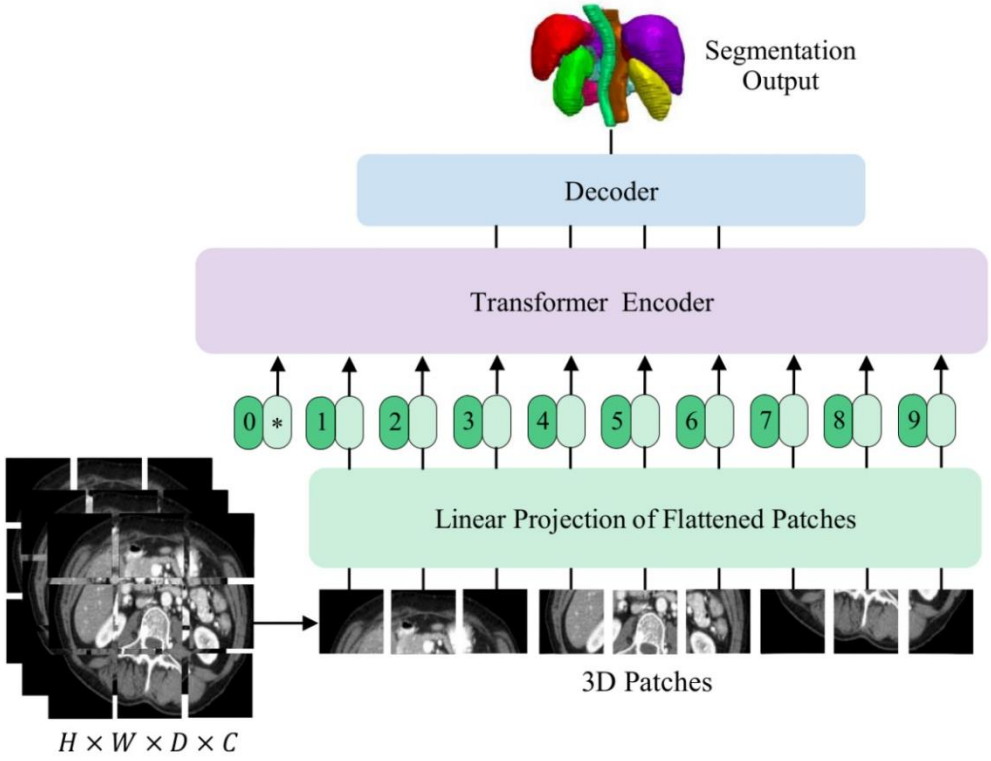

UNetR Architecture

Supporting text here

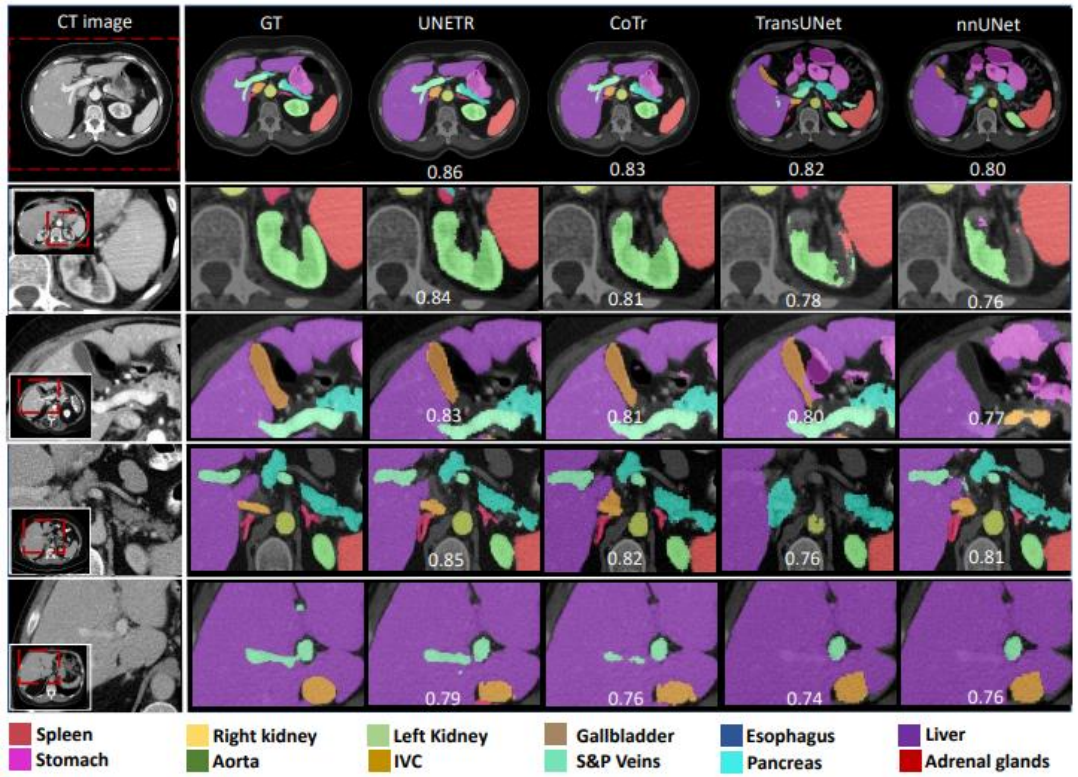

Segmentation Results

Supporting text here

Hatamizadeh, Ali, et al. "Unetr: Transformers for 3d medical image segmentation." *arXiv preprint arXiv:2103.10504* (2021).

INTELLIGENT RETINAL IMAGE ASSESSMENT

Observing the human retina is the gateway to diagnosing vision-threatening conditions and life-threatening vessel disorders. With a shortage of Ophthalmologists, retina images are often assessed by general physicians, which can lead to delayed or inaccurate diagnosis.

To help general physicians improve retinal image assessment VUNO offers a GPU-powered AI solution that detects and highlights twelve different findings — covering most clinically important retinal abnormalities — with over 90% accuracy.

NVIDIA VUNO

www.NVIDIA.com/Inception

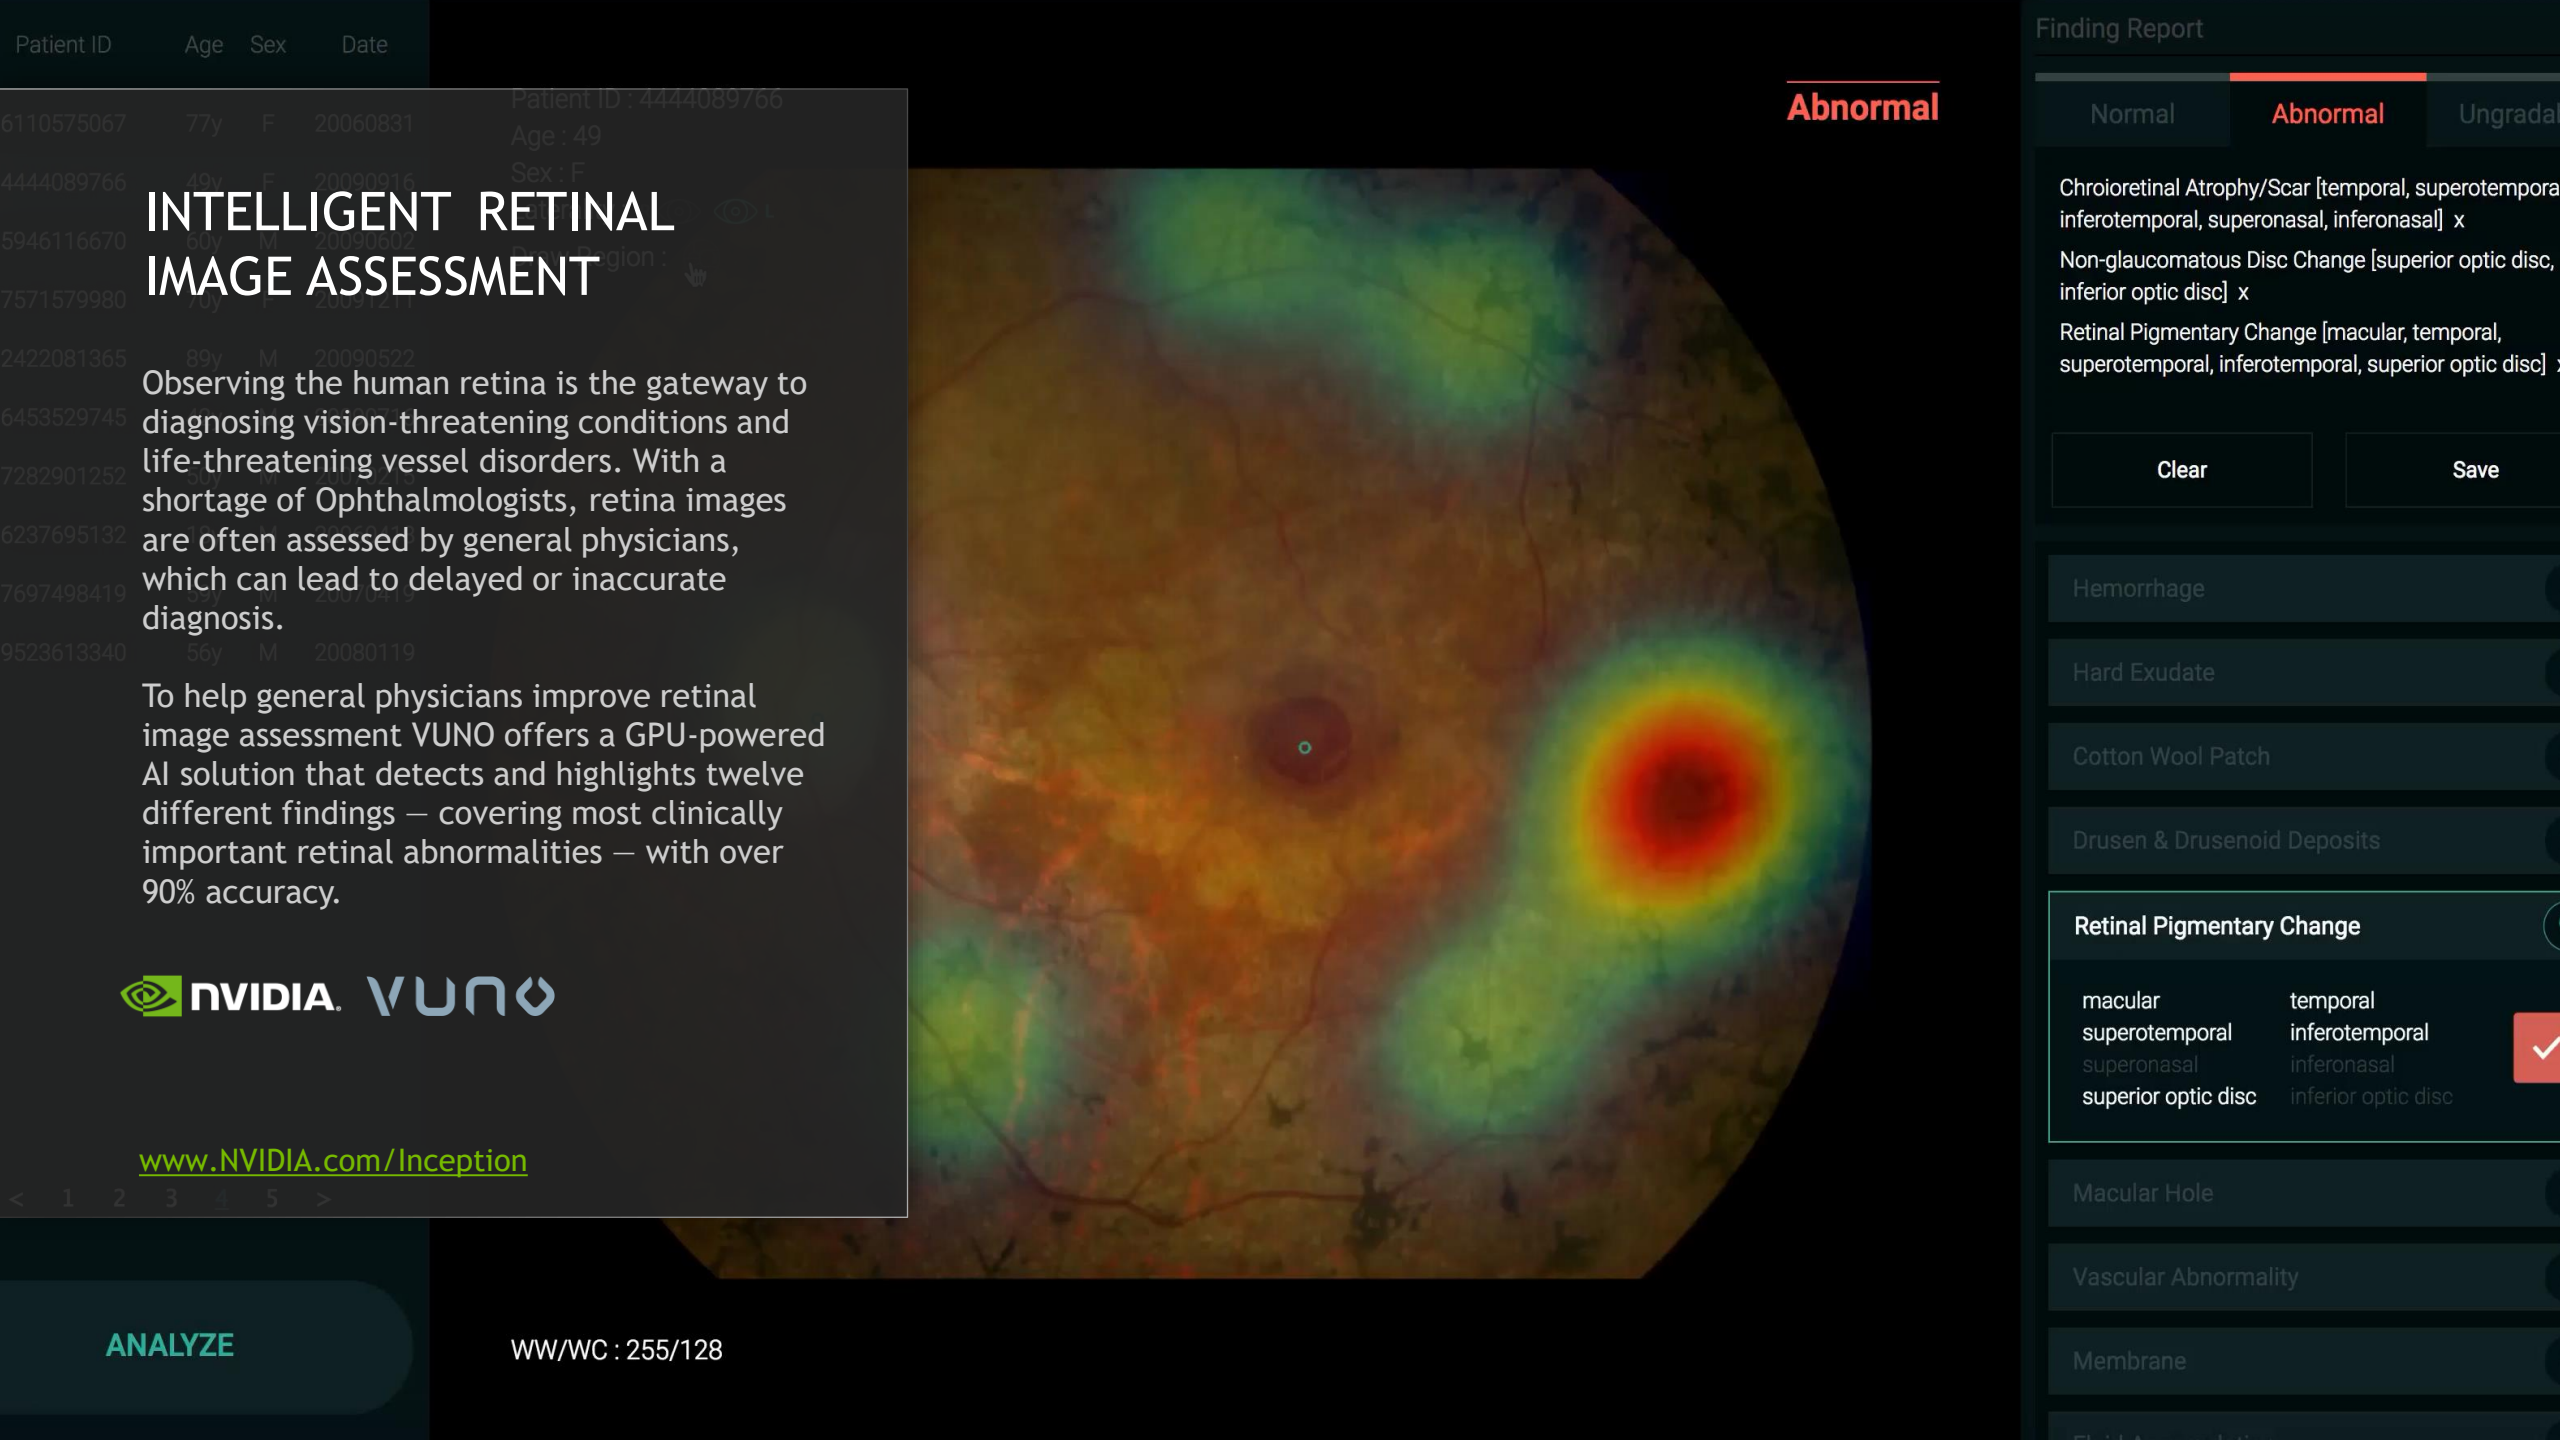

GPU COMPUTING:
A COMPUTATIONAL MICROSCOPE

GPU COMPUTING

Accelerating Scientific Applications with Parallelism

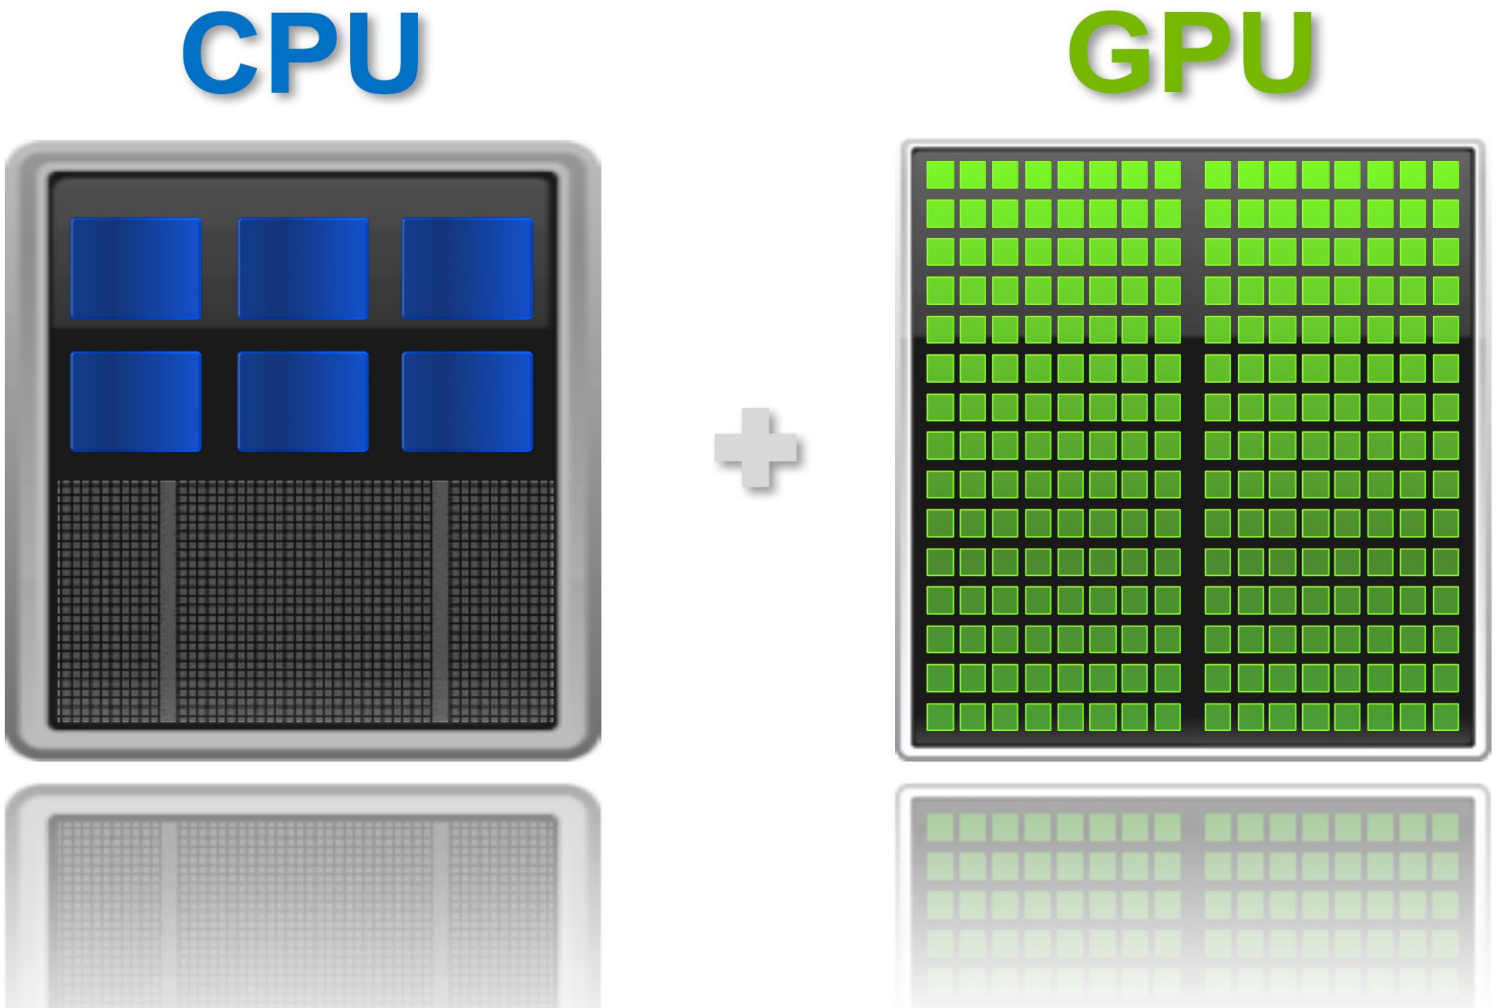

GIVING SCIENTISTS A TIME MACHINE

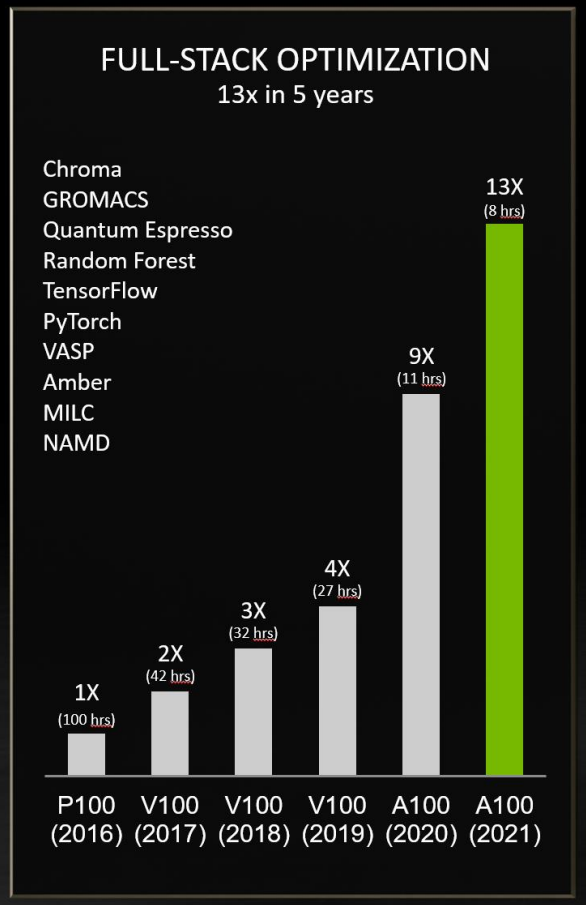

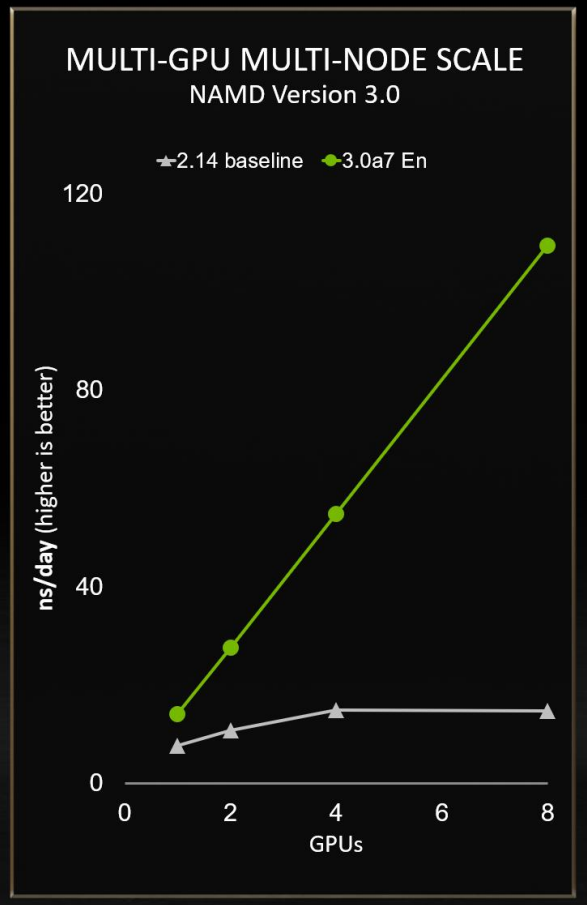

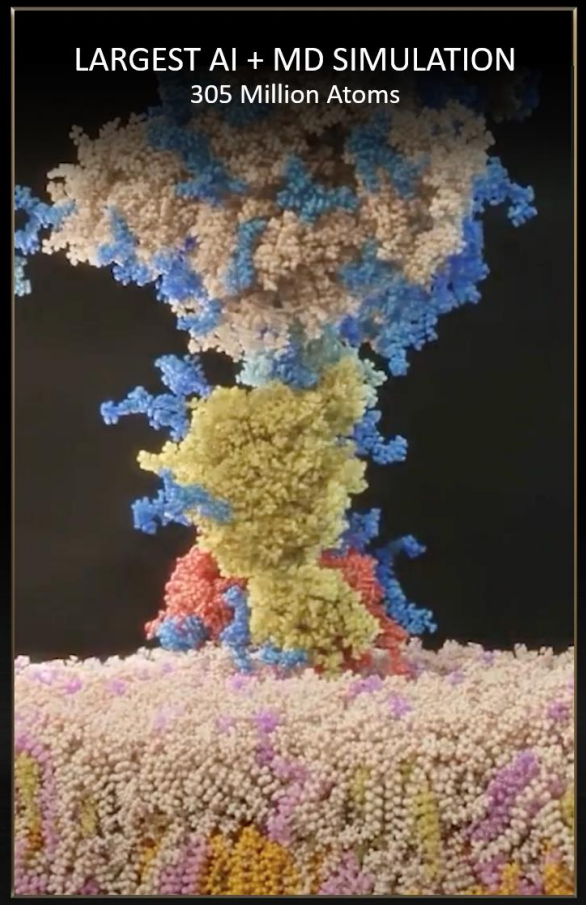

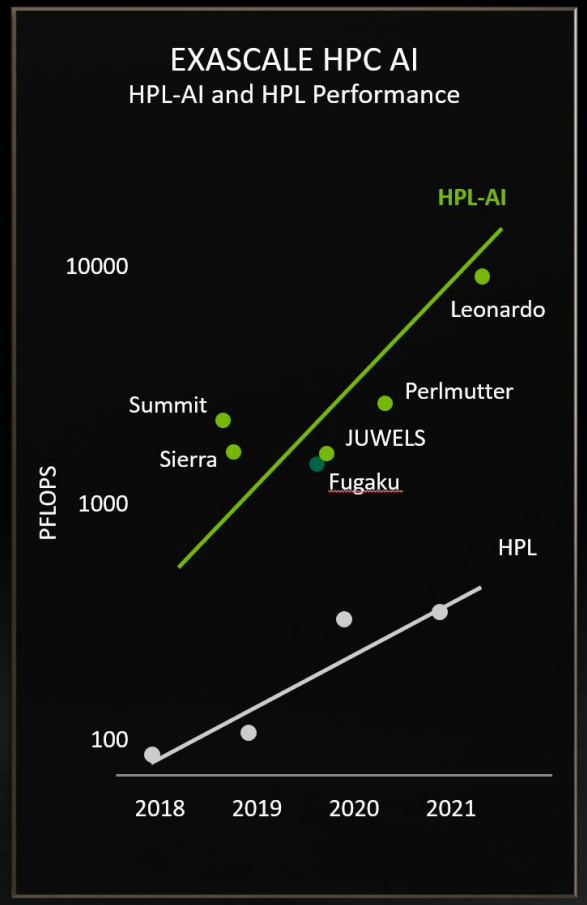

CHALLENGES: ACCELERATING BIG AND SMALL

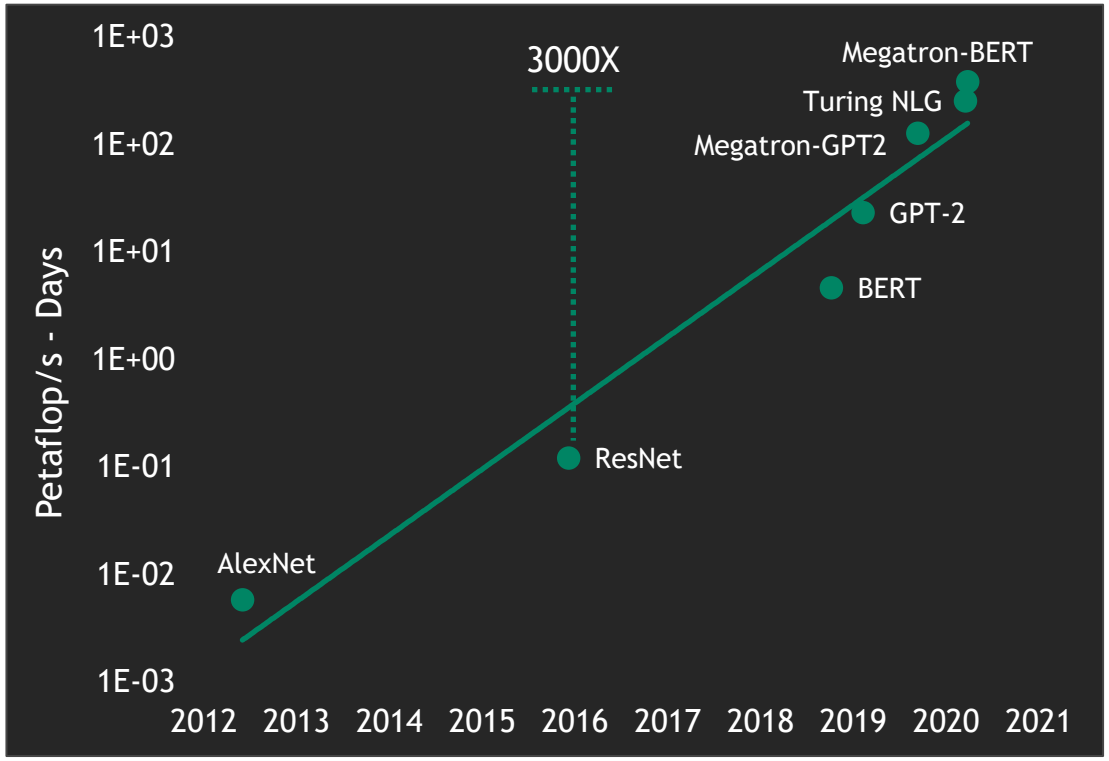

3000X Higher Compute Required to Train Largest Models Since Volta

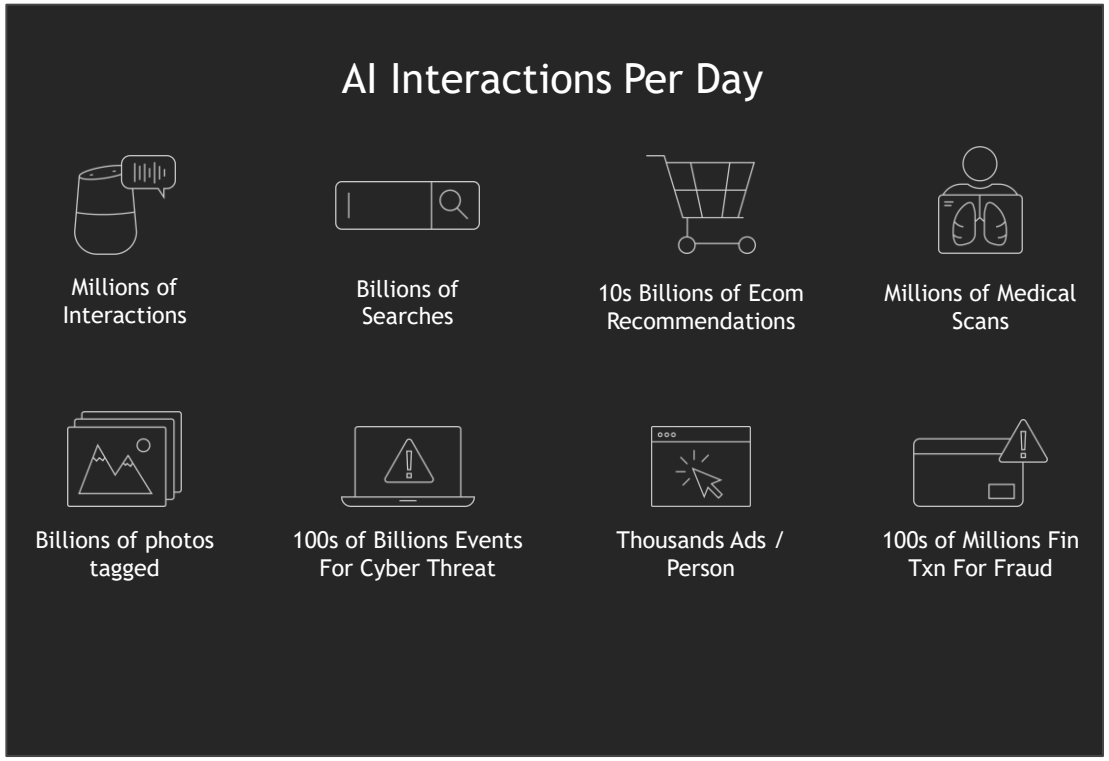

Every AI Powered Interaction Needs Varying Amount of Compute

Source: OpenAI, NVIDIA

DEEP LEARNING FRAMEWORKS: MONAI

NVIDIA CLARA TRAIN & MONAI

Domain Optimized AI Training Framework

NVIDIA Clara Train SDK

AI Application Framework

Using State-of-the-Art AI

Enterprise Grade AI Framework

High Level APIs Enable Workflows

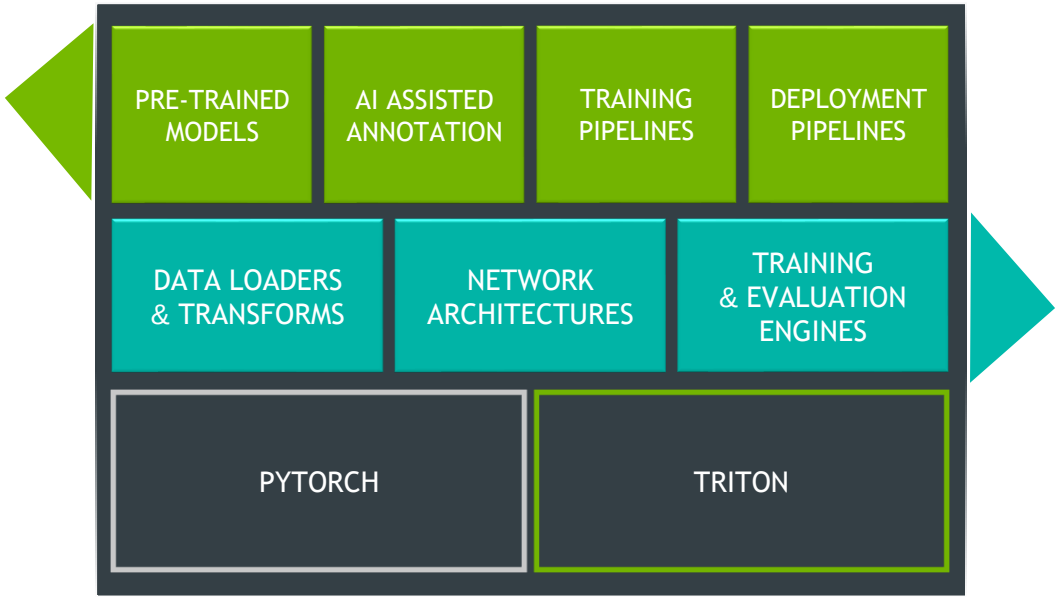

DOMAIN SPECIFIC TRAINING FRAMEWORK

MONAI+

Deep Learning Framework

Inventing State-of-the-Art AI

Opensource based on PyTorch

Optimized Building Blocks & Python APIs

2k Github stars, 50k downloads, 17 organizations

AI luminaries lead advisory board & 8 working groups

MONAI DESIGN GOALS

Ease of Integration & Enterprise Grade SW Development

CUSTOMIZABLE

Abstracted for customizable design for varying user expertise

COMPOSIBLE

Portable with ease of Integration into existing workflows

DOMAIN SPECIALIZED

3D Transformations, N/W Architecture & workflows for Medical Imaging

GPU OPTIMIZED

Multi GPU, CUDA acceleration data & Model Parallel Processing

REPRODUCIBLE

Built for reproducibility and comparison with state of the art

HIGH QUALITY

Tutorials for getting started and robust validation & documentation.

MONAI WORKFLOW MODULES

End-End Workflow for Medical Imaging Deep Learning

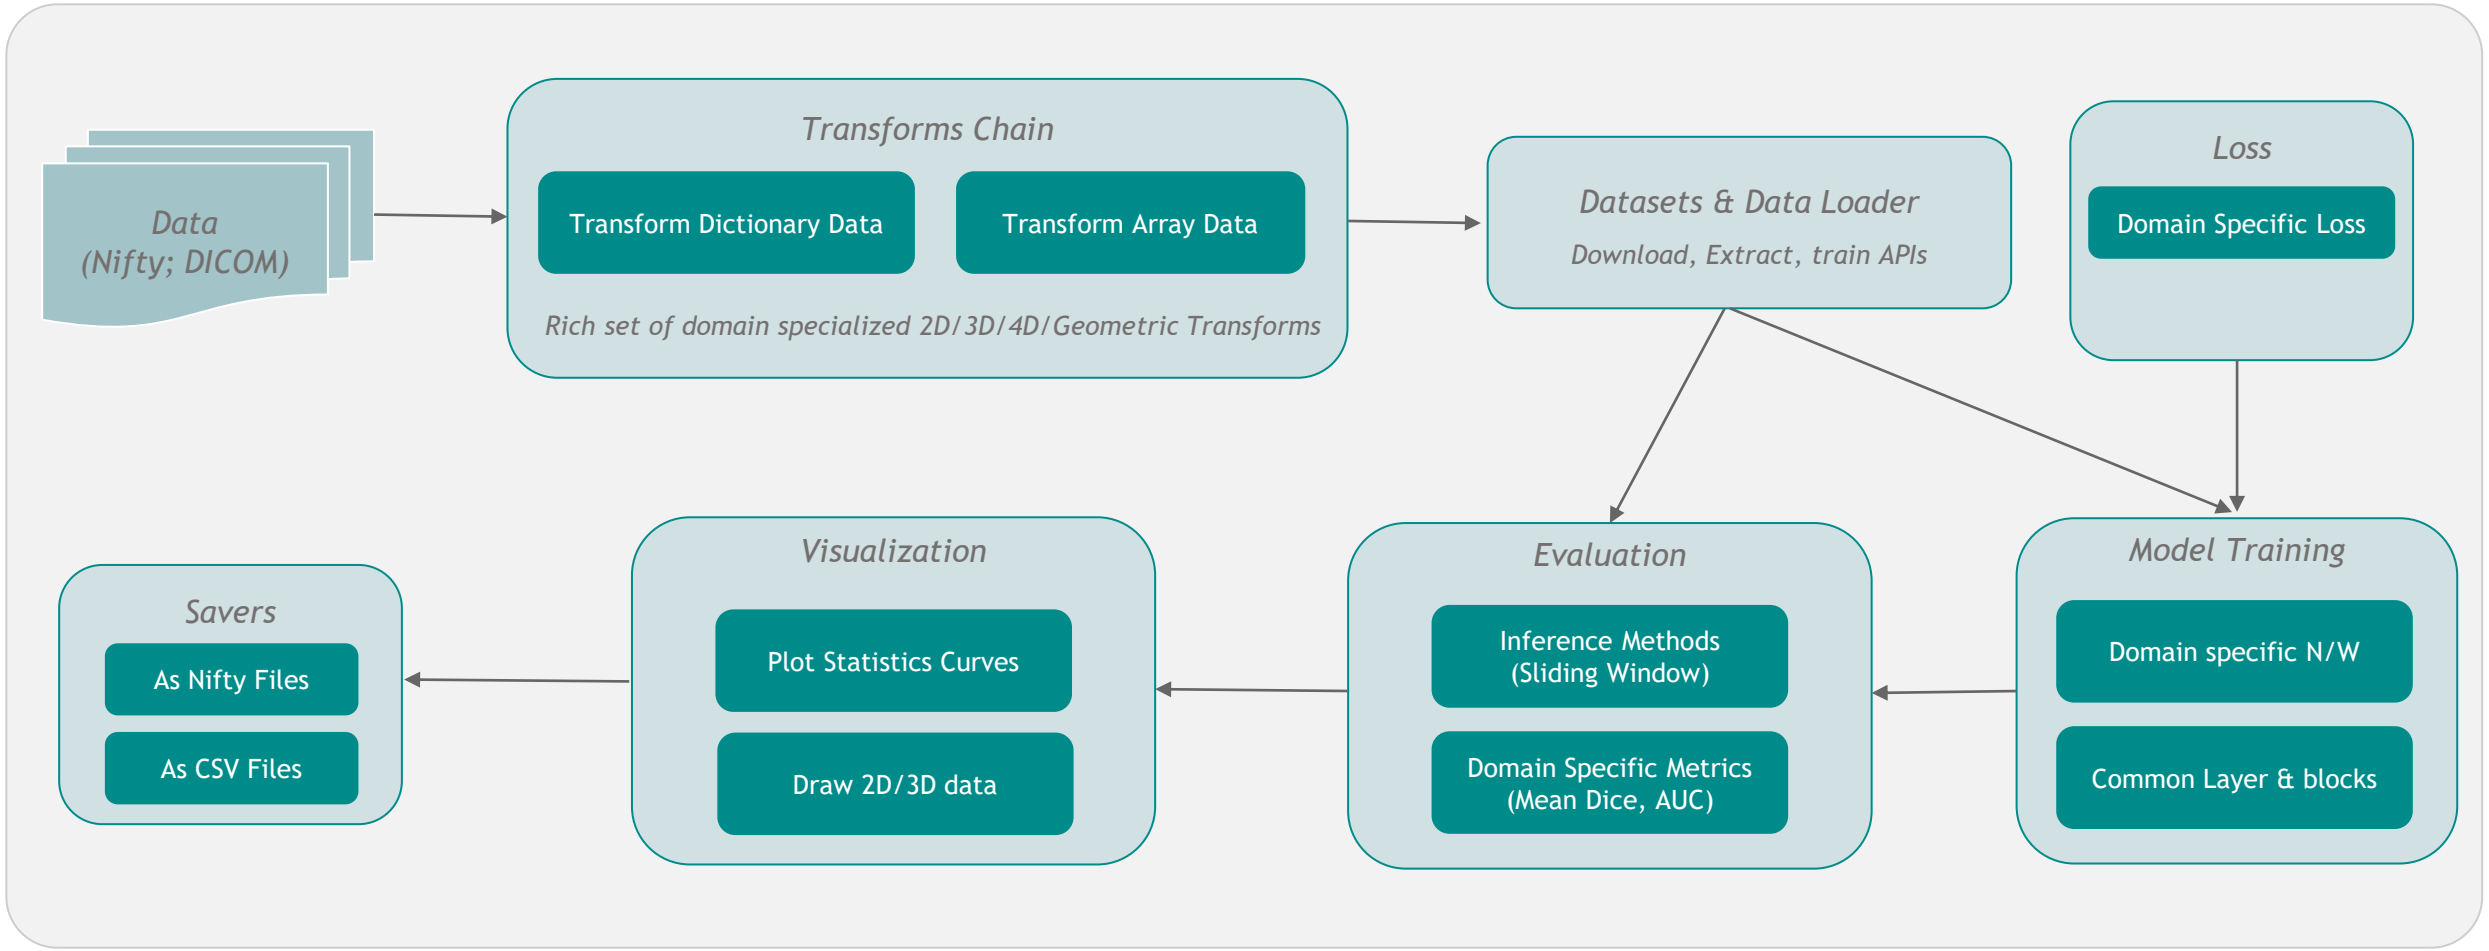

ACCESS MEDICAL DATA

Goal: Harmonize and simplify open data and biomedical challenges

- Participate in / use public challenges
- Define “challenges” (custom datasets) within your lab

Thin layer on top of PyTorch torch.data.utils.Dataset construct

- Automated (verified) download and unzip
- Caching of data as well as intermediate results of preprocessing
- Random splits of training, validation, and test

```python
from monai.apps import DecathlonDataset

dataset = DecathlonDataset(root_dir="./", task="Task05_Prostate", section="training", transform=None, download=True)
print(f"\nnumber of subjects: {len(dataset)}.\nThe first element in the dataset is {dataset[0]}.")
```

```
Task05_Prostate.tar: 100%|██████████| 229M/229M [03:15<00:00, 1.22MB/s]
Verified 'Task05_Prostate.tar.part', md5: 35138f08b1efaef89d7424d2bcc928db.
Verified 'Task05_Prostate.tar', md5: 35138f08b1efaef89d7424d2bcc928db.
Verified 'Task05_Prostate.tar', md5: 35138f08b1efaef89d7424d2bcc928db.
Load and cache transformed data: 100%|██████████| 26/26 [00:00<00:00, 196489.92it/s]
number of subjects: 26.
The first element in the dataset is {'image': 'Task05_Prostate/imagesTr/prostate_46.nii.gz', 'label': 'Task05_Prostate/label
```

MONAI: END TO END TRAINING WORKFLOW IN 10 LINES OF CODE

```
from monai.application import MedNISTDataset
from monai.data import DataLoader
from monai.transforms import LoadPNGd, AddChanneld, ScaleIntensityd, ToTensord, Compose
from monai.networks.nets import densenet121
from monai.inferers import SimpleInferer
from monai.engines import SupervisedTrainer

transform = Compose(
    [
        LoadPNGd(keys="image"),
        AddChanneld(keys="image"),
        ScaleIntensityd(keys="image"),
        ToTensord(keys=["image", "label"])
    ]
)

dataset = MedNISTDataset(root_dir="./", transform=transform, section="training", download=True)

trainer = SupervisedTrainer(
    max_epochs=5,
    train_data_loader=DataLoader(dataset, batch_size=2, shuffle=True, num_workers=4),
    network=densenet121(spatial_dims=2, in_channels=1, out_channels=6),
    optimizer=torch.optim.Adam(model.parameters(),lr=1e-5),
    loss_function=torch.nn.CrossEntropyLoss(),
    inferer=SimpleInferer()
)

trainer.run()
```

NVIDIA CLARA

NVIDIA CLARA PRE-TRAINED MODELS

Save Thousands of Hours | Millions of Dollars

3D SEGMENTATION AND INTERACTIVE ANNOTATION
Spleen | Liver | Brain | Prostate | DeepGrow 2D & 3D

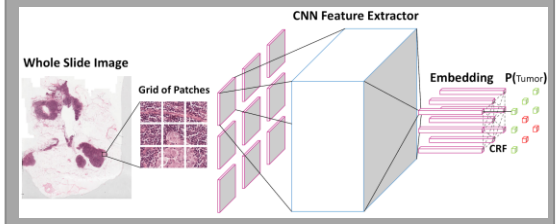

METASTSIS DETECTION – ResNet
Automated Detection of Metastasis WSI of lymph node sections

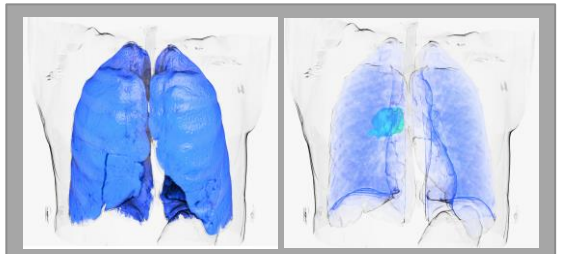

COVID-19 CT PROGNOSIS
Lung Seg | COVID Classification | Lesion Seg

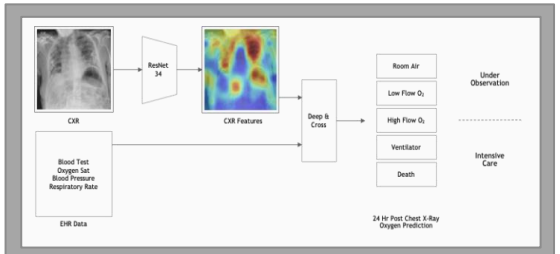

COVID-19 OXYGEN PREDICTION
Model Developed by Federated Learning | 20 Global Sites

Clara Imaging Pre-Trained Models in PyTorch
20+ Pre-Trained Models: CT, MRI, X-Ray, Digital Pathology

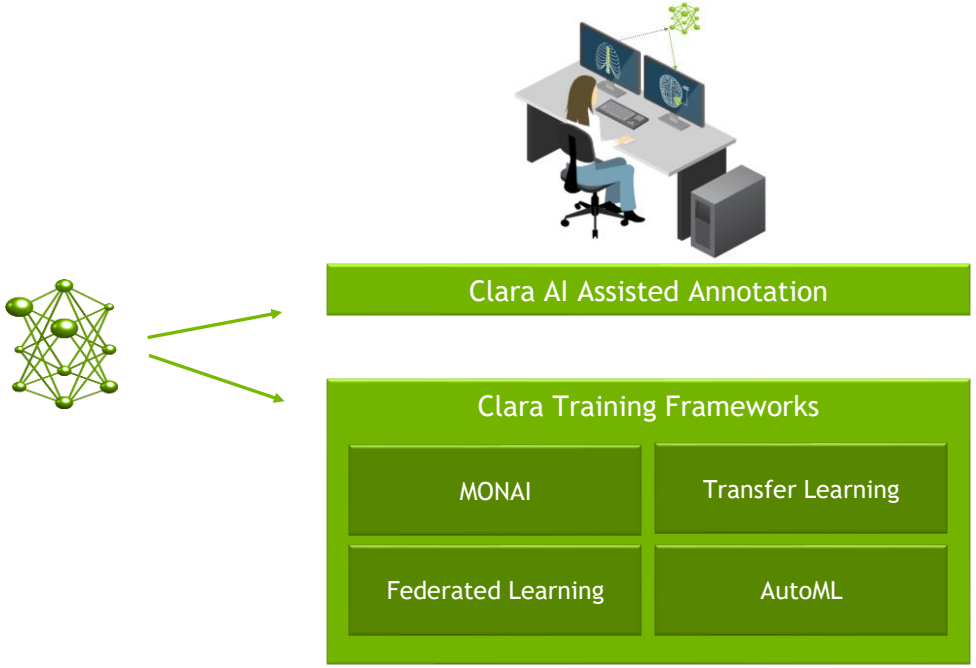

Data Labeling & Model Training
Jumpstart Complex 3D Data Labeling | Reduce Training Data Needed

INTERACTIVE ANNOTATION DEEPGROW 3D

General availability with ability to re-train the model

FOVIA AI

Next generation general purpose medical imaging annotation model in 3D

Foreground & background clicks to deliver high quality segmentation

Model prediction significantly improved for previously unseen images

Integrated into Fovia viewer SDK to enable commercial clinical integrations

Integrated into 3DSlicer, OHIF to enable open-source community

Video

CLARA DEPLOY APPLICATION CONCEPTS

Operators, Pipelines, Jobs & Services

Operator

An operator is an element of a pipeline. Each operator is typically designed to perform a specific function on incoming input data.

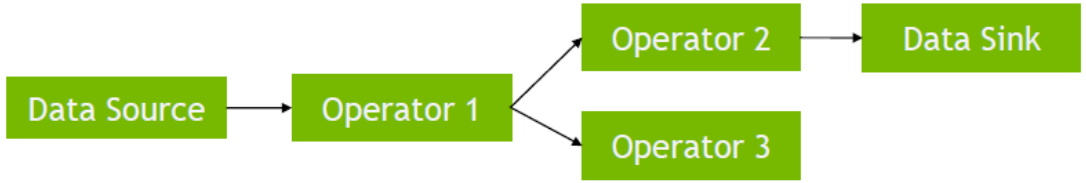

A pipeline is a collection of operators that are configured to work together to execute a Clara Deploy pipeline. A job is an instance of a Pipeline executed for a specific dataset.

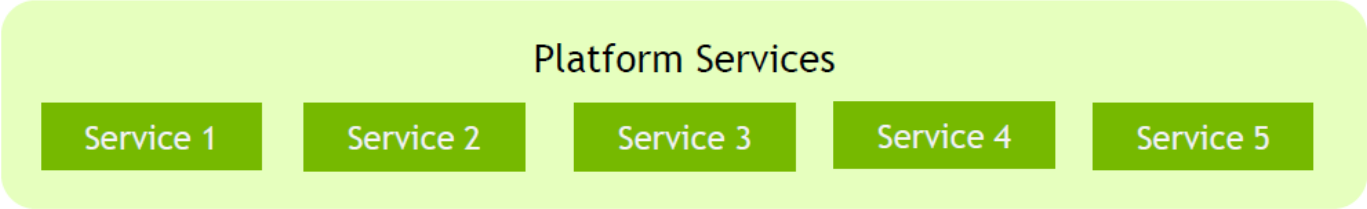

A service is an abstraction which defines a mechanism to access a computational facility available in the Clara Platform. Operators defined in a Clara pipeline could make use of one or more services at run-time.

MR PROSTATE PIPELINE EXAMPLE

17 Reference Pipelines

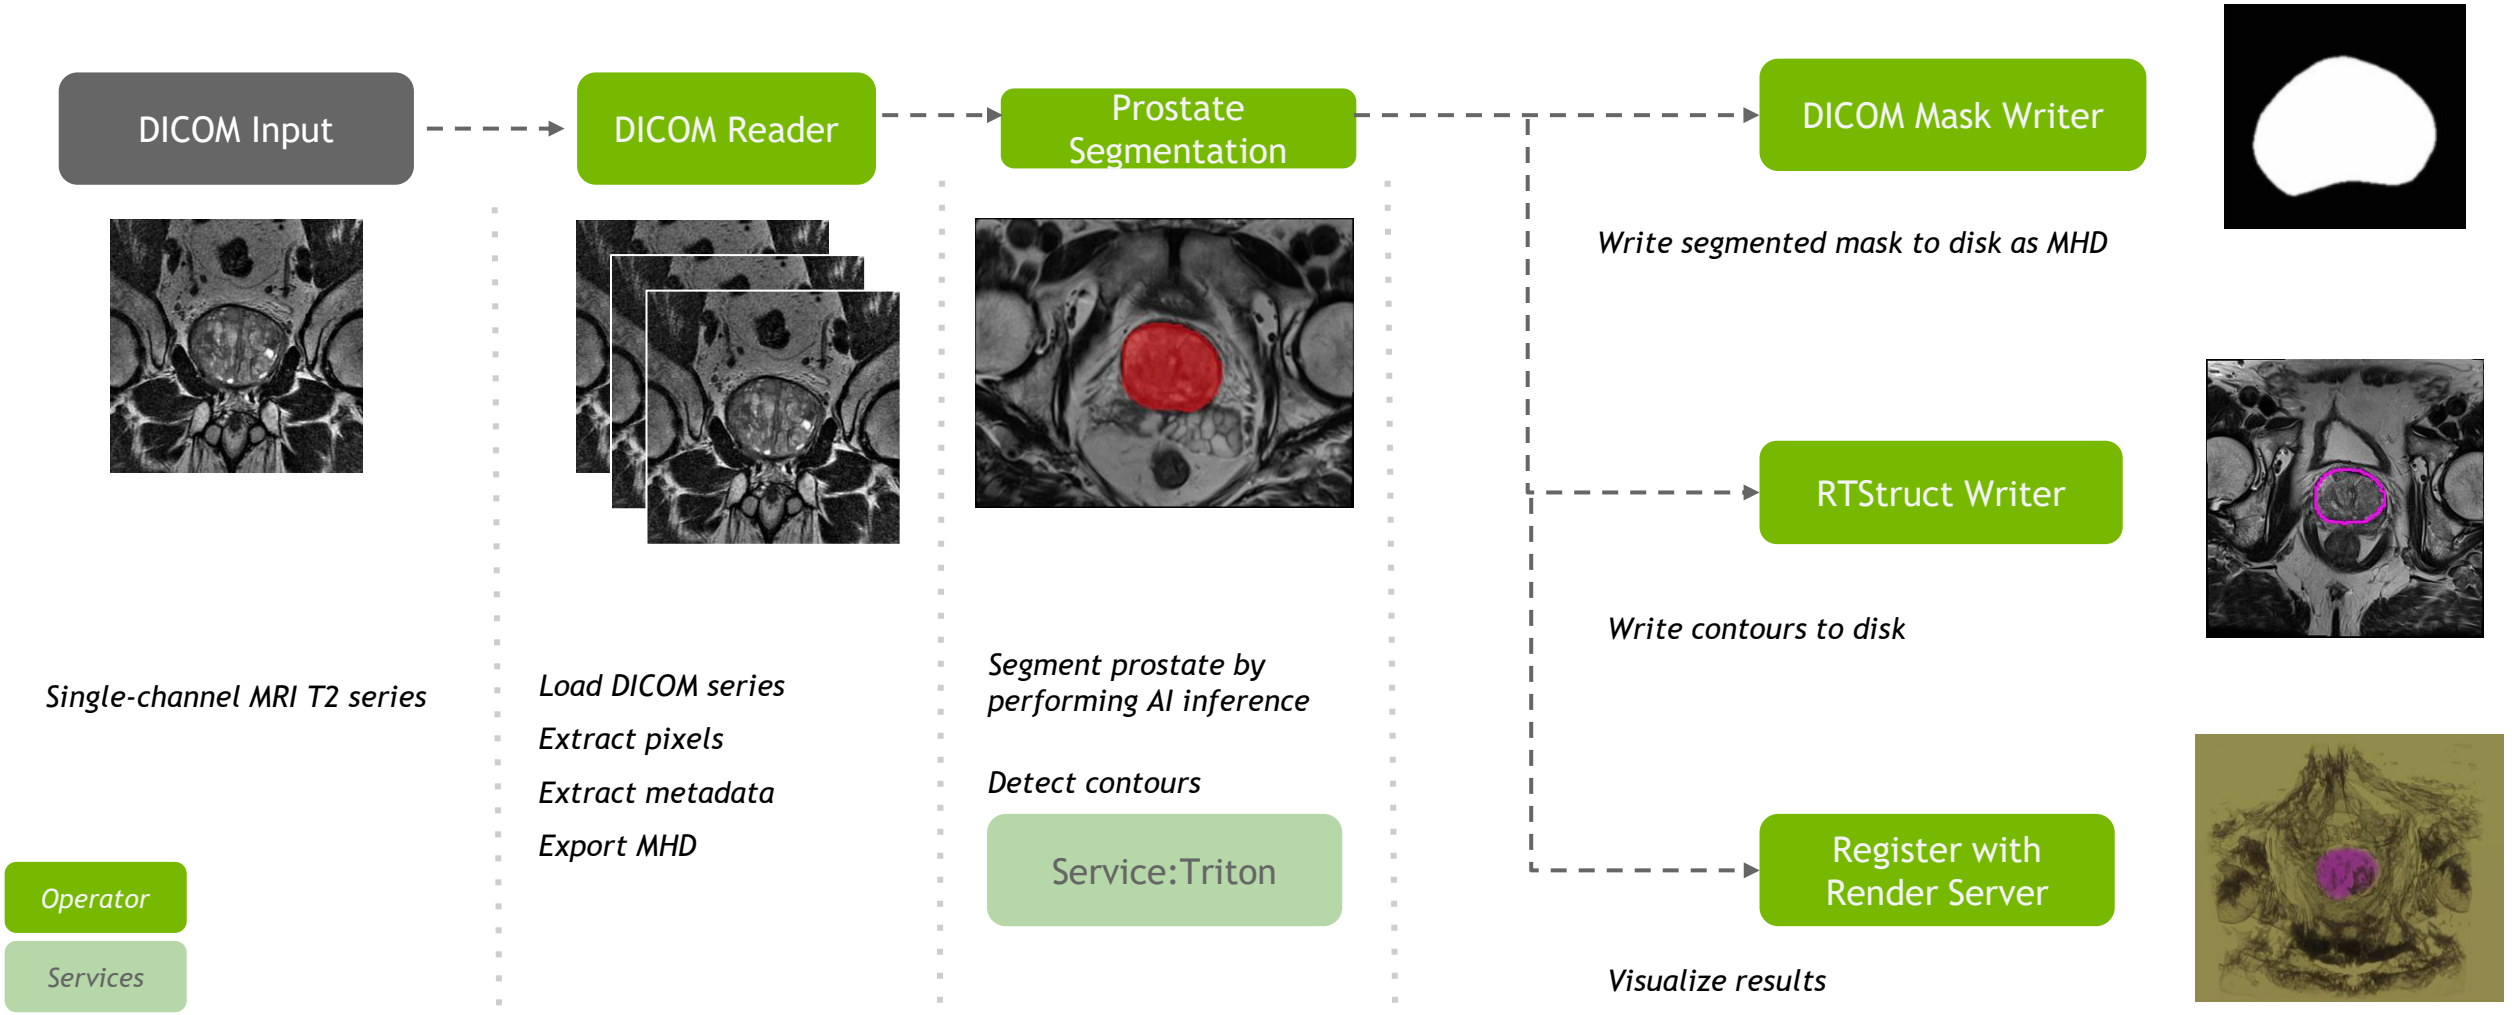

COVID-19 REFERENCE PIPELINE

Encapsulating multiple AI algorithms into an executable AI application

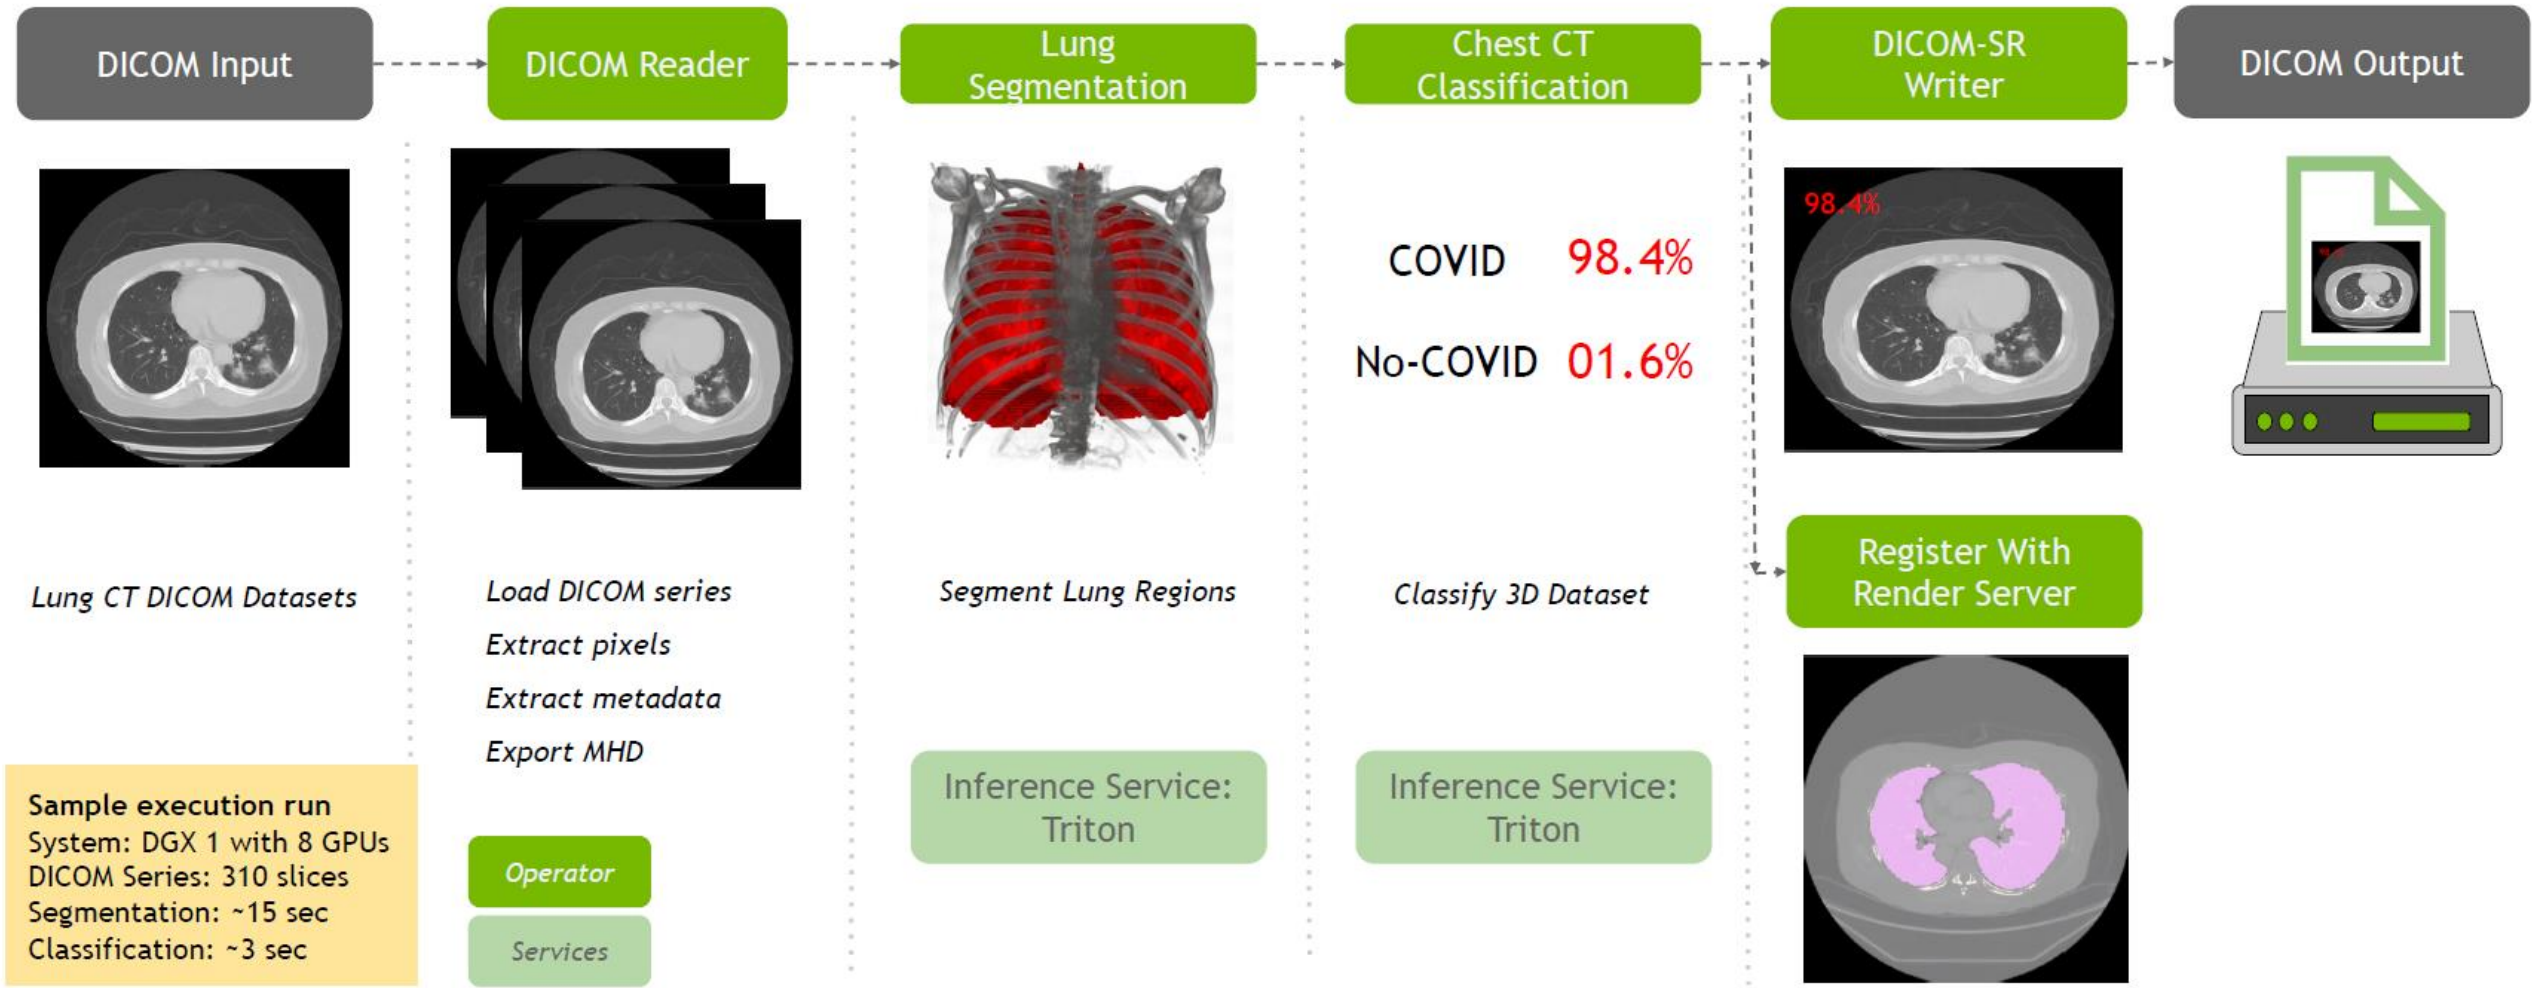

MULTI AI PIPELINE EXAMPLE

Multi-Organ Segmentation

Input

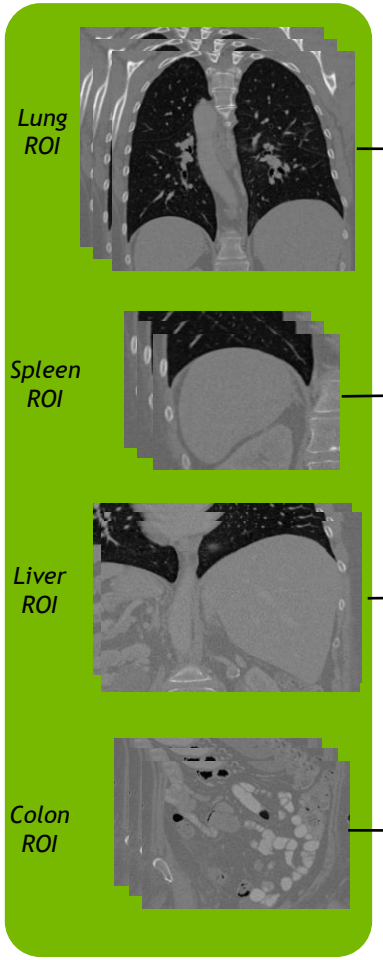

ROI-Generator (SPLIT)

Lung Tumor AI
Triton

Spleen AI
Triton

Liver Tumor AI
Triton

Colon Tumor AI
Triton

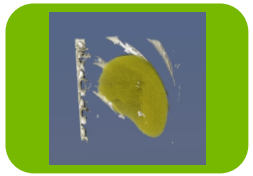

No Colon Tumors detected

ROI-Generator (MERGE)

Operator

Services

RESOURCES

NVIDIA GTC

BRILLIANT MINDS. BREAKTHROUGH DISCOVERIES.

The Conference for AI Innovators, Technologists, and Creators

Technology advancements move fast. Join us at NVIDIA GTC this November to keep up with the latest AI use cases and hear directly from the leaders driving innovation.

Register now at www.nvidia.com/GTC to get unique insights and training on the advanced technologies that are transforming today's industries.

NOVEMBER 8-11, 2021 | www.nvidia.com/GTC

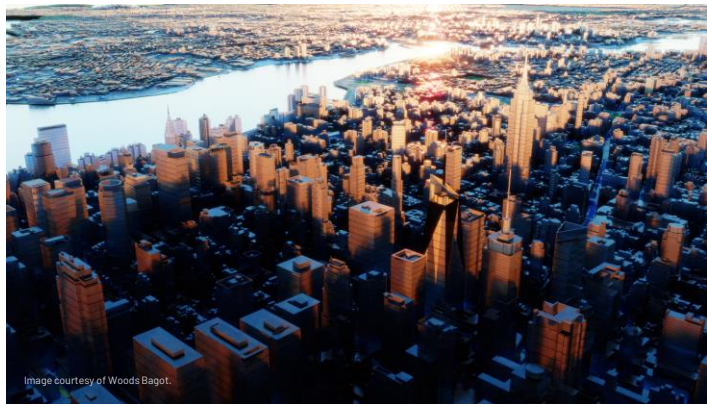

WHERE TO GO FROM HERE?

- https://www.kaggle.com
- https://ngc.nvidia.com
- https://www.deeplearningbook.org/
- https://monai.io
- https://distill.pub/
- https://www.deeplearningweekly.com/
- https://hpc.nih.gov/ (Biowulf)
- https://ncihub.org/groups/nihai
- https://btep.ccr.cancer.gov/ai/
- https://cbiit.github.io/p2p-datasci/

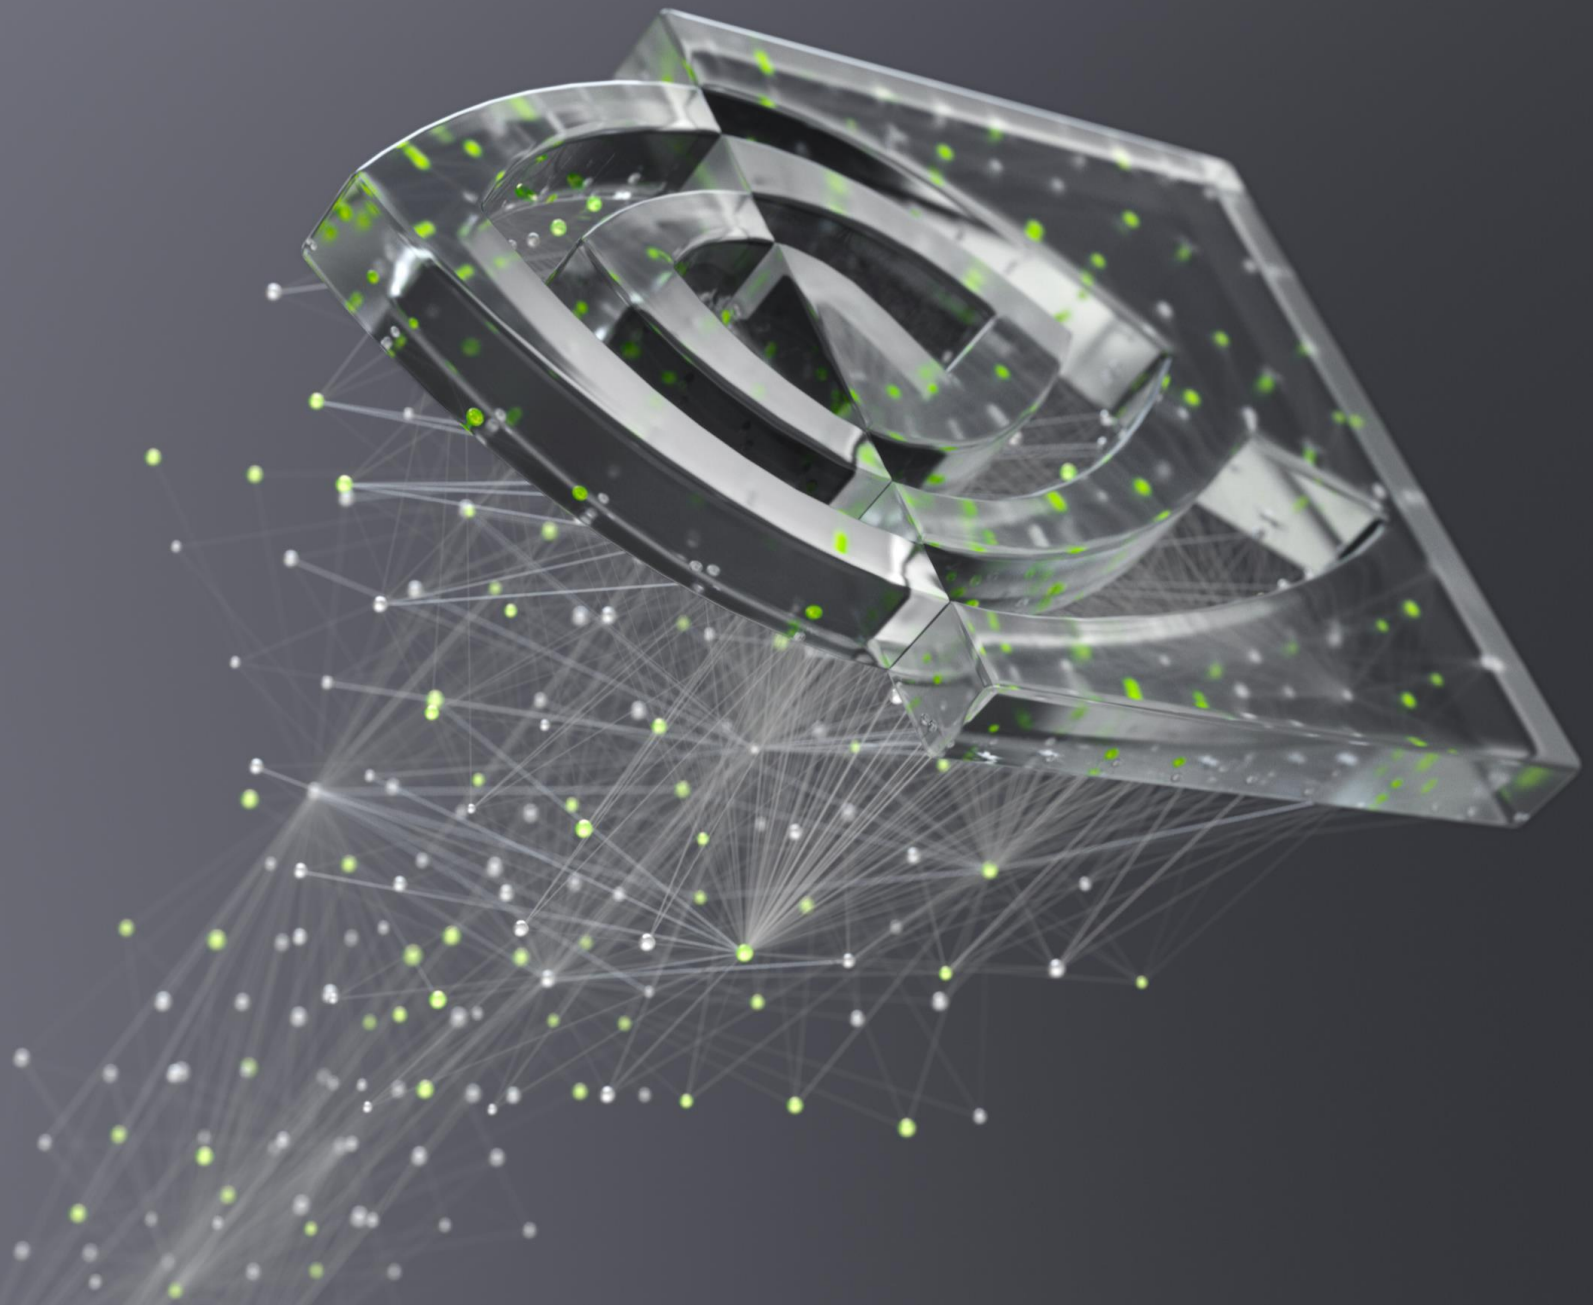

NVIDIA®