LIVE WEBINAR

Take the Stress Out of Going from Training to Inferencing with the OpenVINO™ toolkit

Sit back and relax. The webinar will begin shortly!

intel®

Sit tight and get familiar with your live webinar view

Use the **chat** during the live event for technical assistance/help from the Event

1

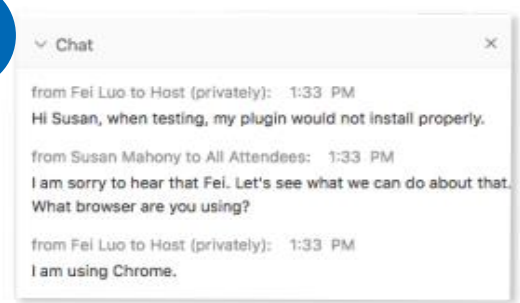

Use the **Q&A** panel to pose your questions to the presenters

- When submitting questions you must have **"All Panelists"** showing in the **"Ask"** text box

2

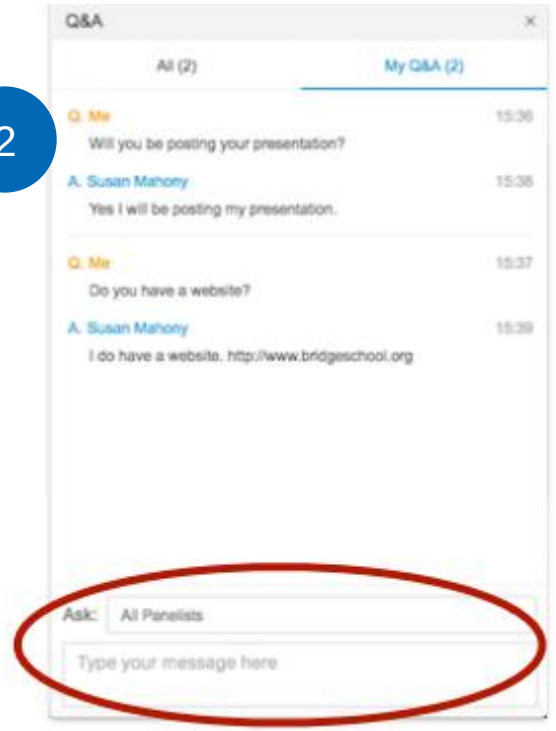

Use the **toolbar** as a shortcut to access your audio connections, chat, Q&A, and leave event

3

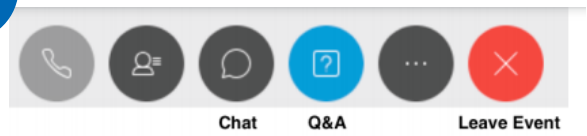

Taking AI to Production Can Be Treacherous

from idea..

..to enabling biopharmaceutical research using medical text mining

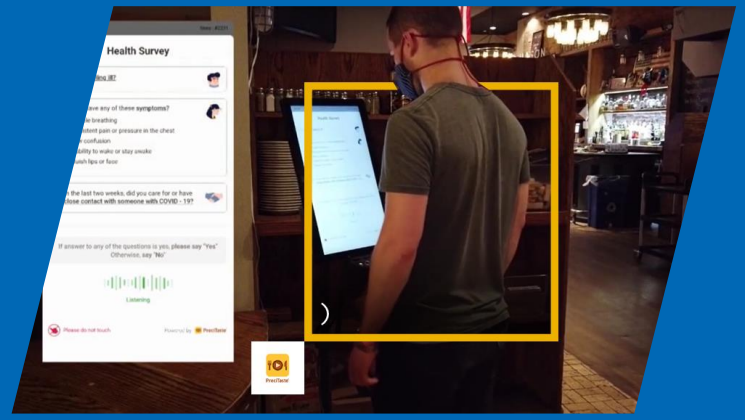

..to helping restaurants safely open during the pandemic

Real World Requirements Can Hamper Productivity for Many Dev Teams

PERFORMANCE

ACCURACY

TIME TO MARKET

NO ONE SIZE FITS ALL

Poll Time!

Answer in 30 seconds

Hi, nice to meet you!

Zoe Cayetano
Product Manager

Raymond Lo
Software Evangelist

Sergei Nosov
Engineering Manager

Poll Time!

Answer in 30 seconds

How do I turn data into a DL model maximum performance at minimum effort?

Practical meets Pragmatic Implementation

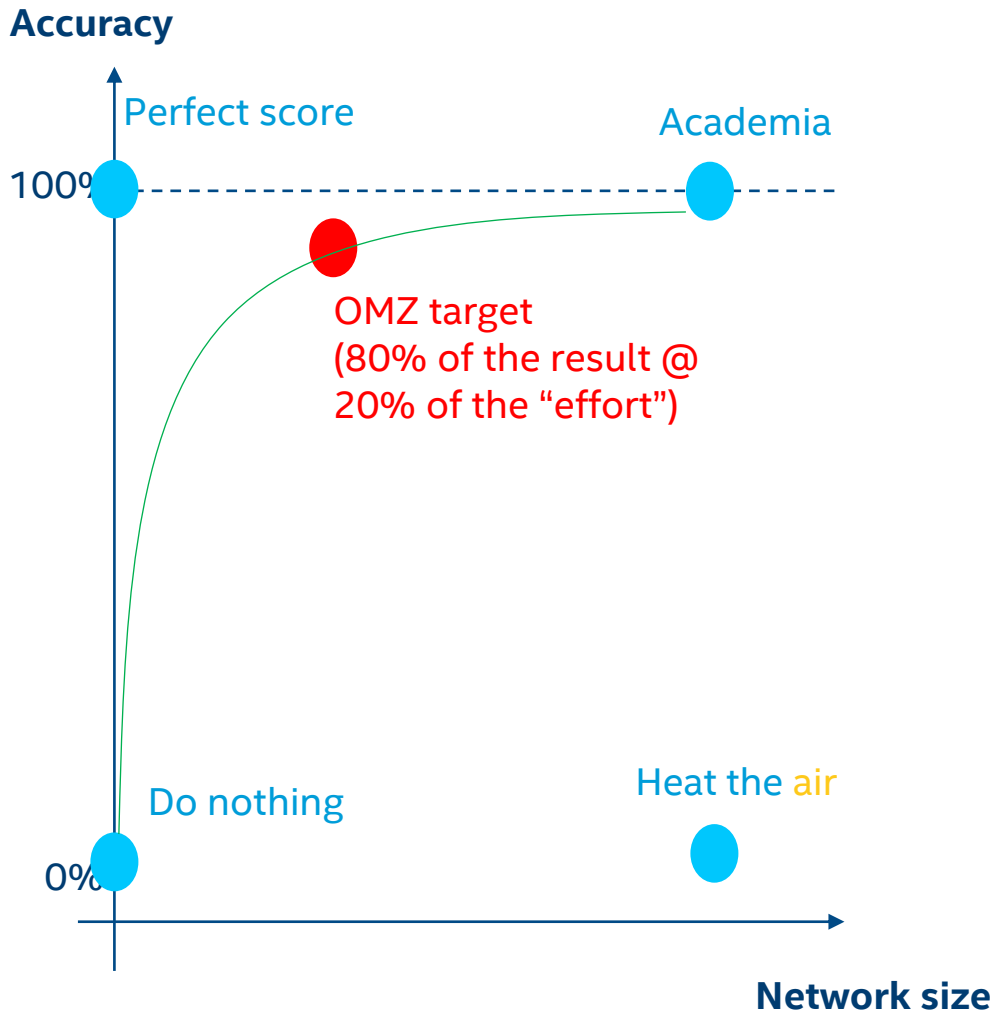

Deep Dive into Image Classification

#	Name	# of images	# of classes
1	Chest X-Ray Images (Pneumonia)	5232/624	2
2	Fashion MNIST	60000/10000	10
3	SVHN	73257/26032	10
4	CIFAR100	50000/10000	100
5	Food-101	75750/25250	101
6	Birdsnap	47386/2443	500
7	SUN397	92440/16314	397
8	Cars Dataset	8144/8041	196
9	Oxford-IIIT Pets	3680/3369	37
10	Oxford 102 Flowers	6614/1575	102
11	Caltech 101	6941/1736	101
12	Describable Textures (DTD)	4826/814	47

Experiment	AVG mAP	AVG TOP1	AVG TOP5	Latency, ms	FPS	Params, M	GFLOPs
baseline efficientnet-b0	88.39	86.42	96.99	9.7	447.5	4.14	0.76
final efficientnet-b0	90.54	88.46	97.89				
baseline MN-v3-large-1	81.04	82.08	94.82	5.58	792.2	4.33	0.44
final MN-v3-large-1	90.22	87.06	97.28				
baseline MN-v3-large-075	83.10	84.16	96.38	4.48	985.4	2.84	0.308
final MN-v3-large-075	89.20	86.15	97.16				
baseline MN-v3-small	77.28	80.85	95.07	2.79	1599	1.56	0.112
final MN-v3-small	85.60	83.75	95.96				

Academic

COCO Explorer

COCO 2017 train/val browser (123,287 images, 886,284 instances)

person ✖ car ✖

8878 results

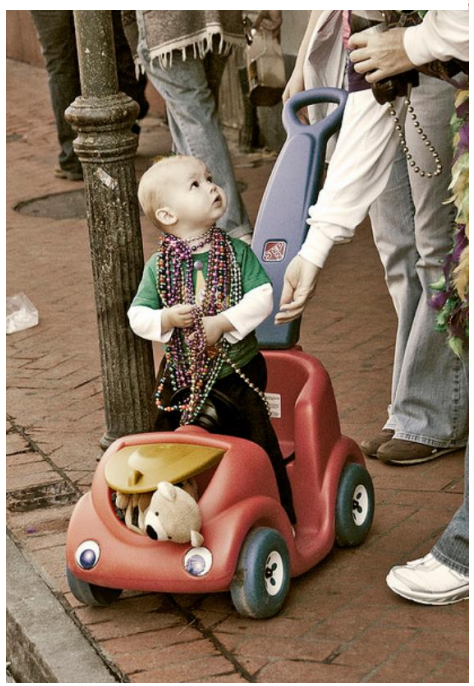

Production

- ~65x bigger than MS COCO in size

Type of object	Number of unique objects	Number of Frames
Pedestrian	1,020,114	7,997,092
Vehicle	700,843	7,282,679
Ignored	258,397	6,941,283
Non-vehicle	127,125	7,008,948

- ~$250K worth of data and annotation (excl. R&D costs)
- Contains data specifically for a safety & security scenario

Intel® Distribution of OpenVINO™ toolkit : The Stack

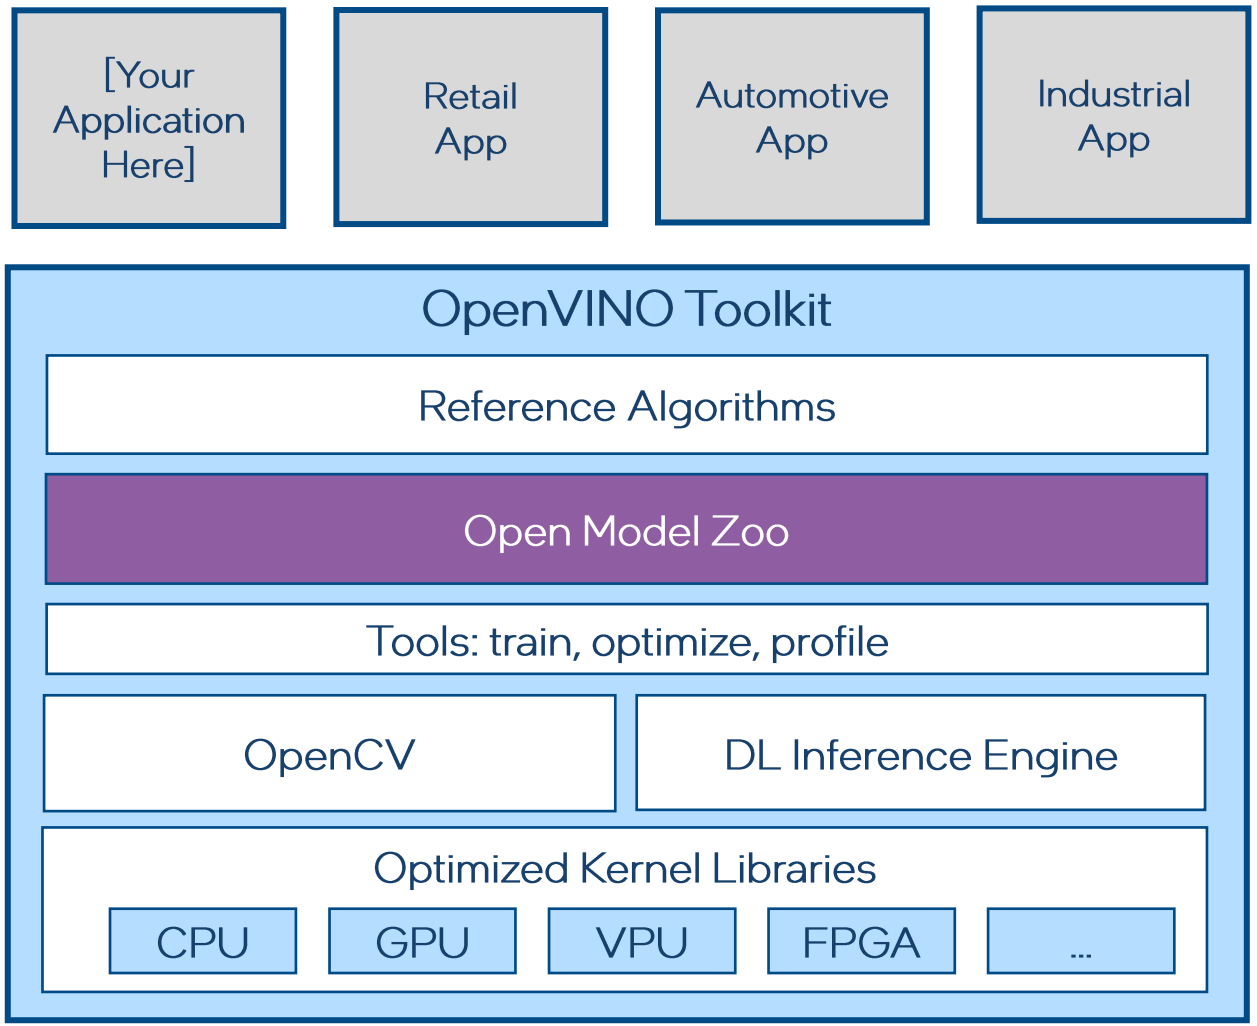

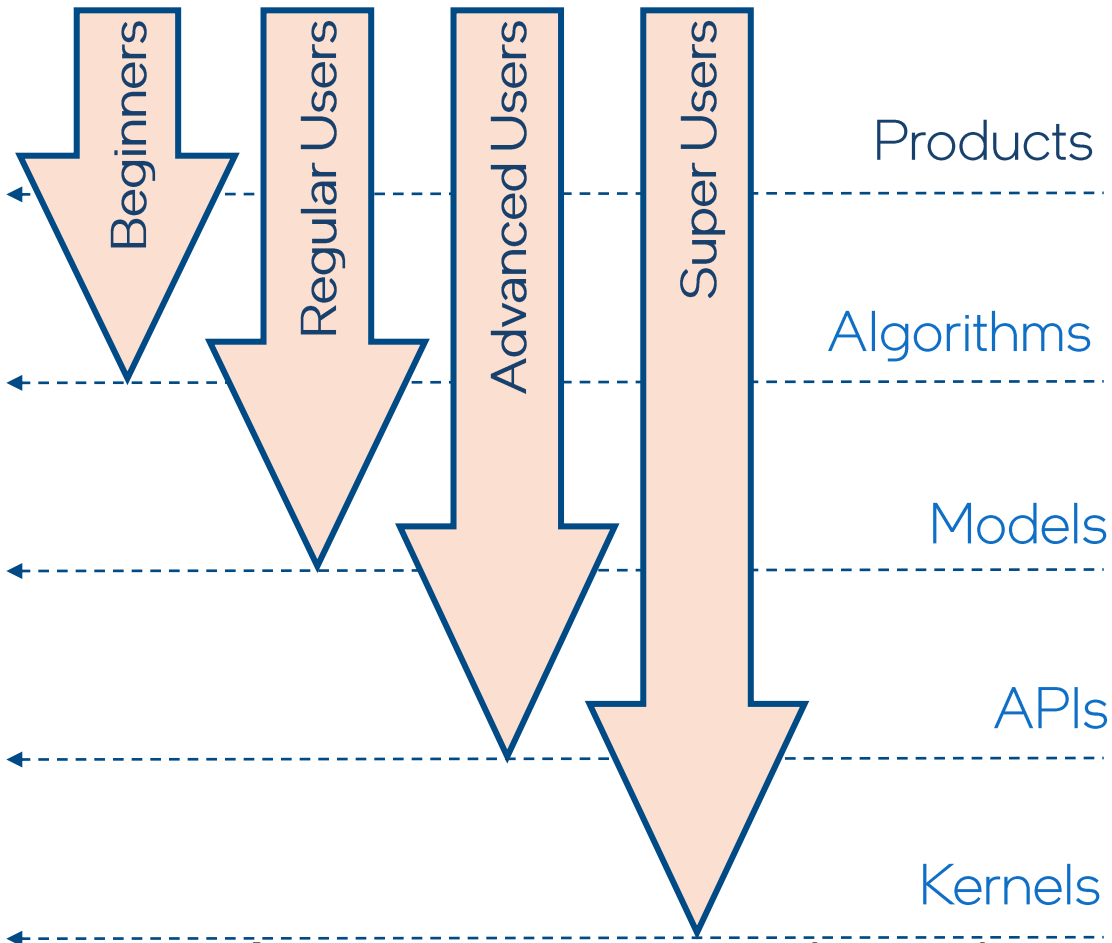

Layered access: every user can choose the proper API level for integration.


THE WESTIN GRAND BERLIN
truck
car
car
person
person
bicycle
handbag
person
bicycle
car

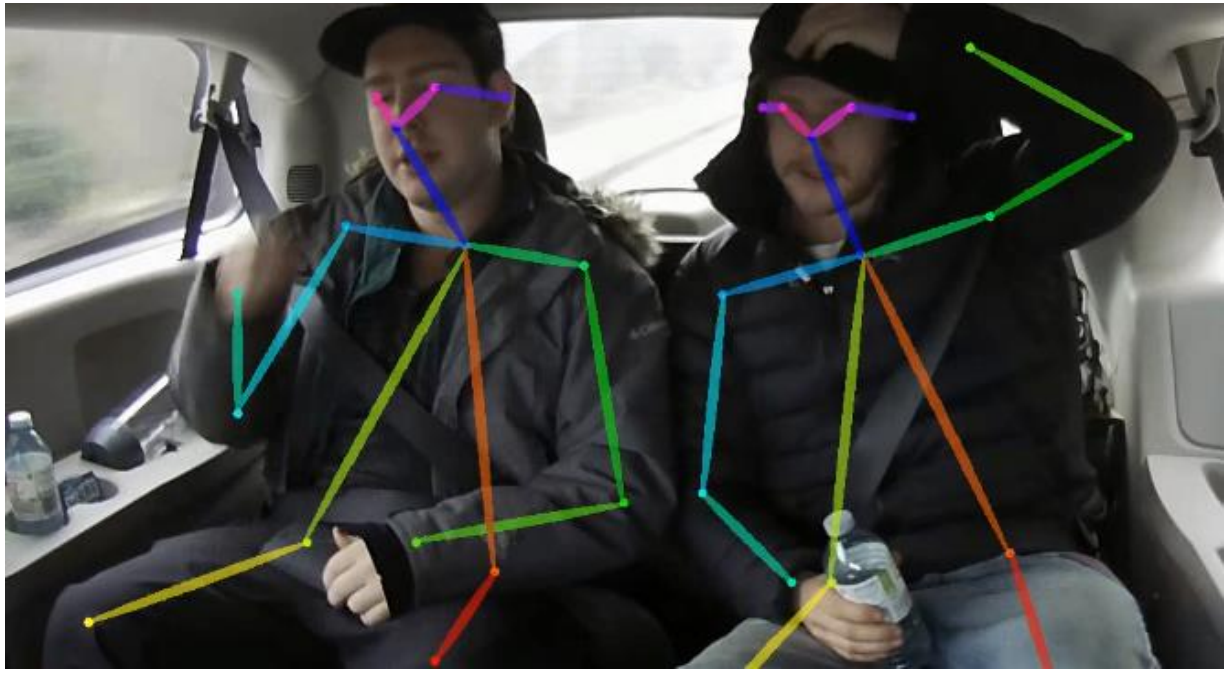


NUMBERS

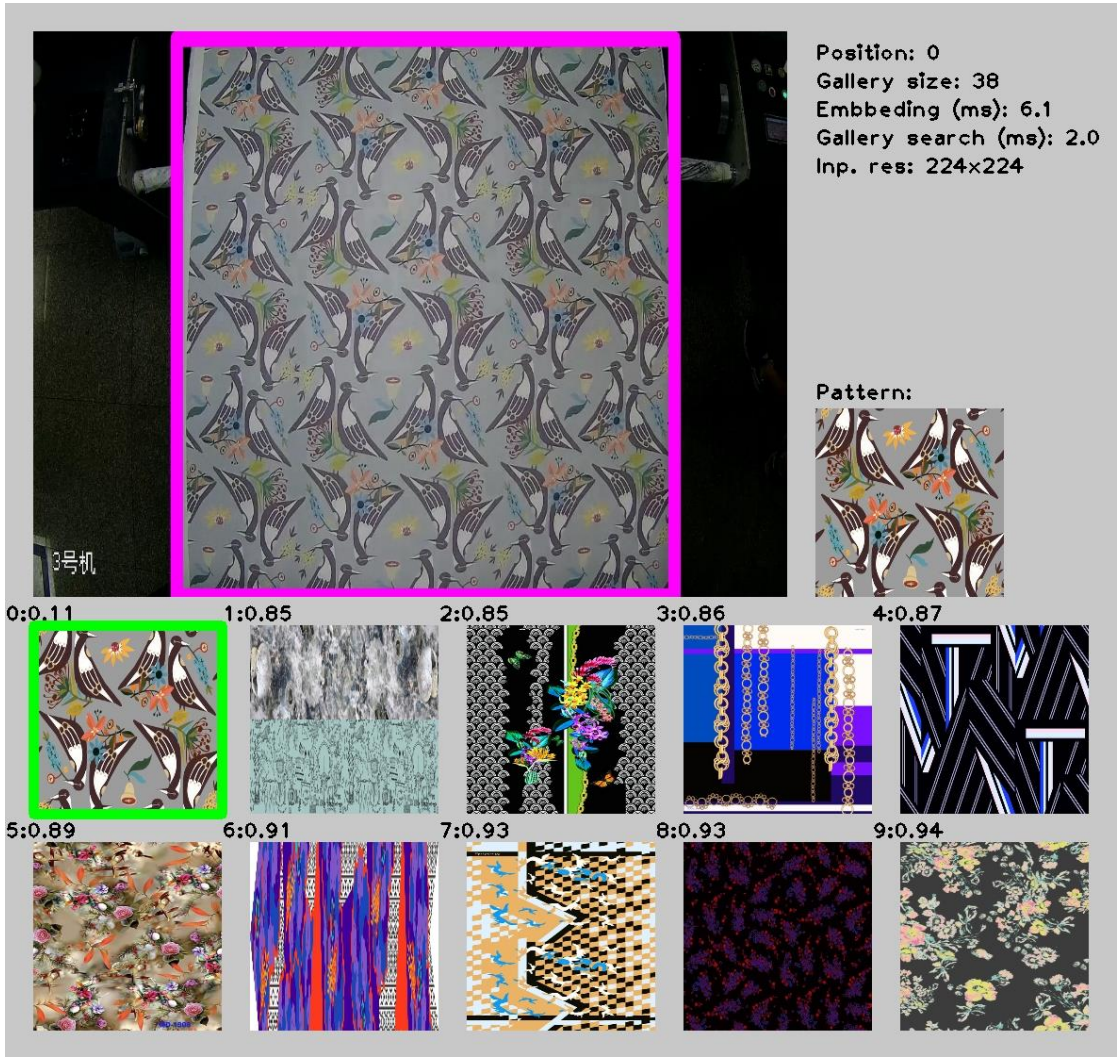

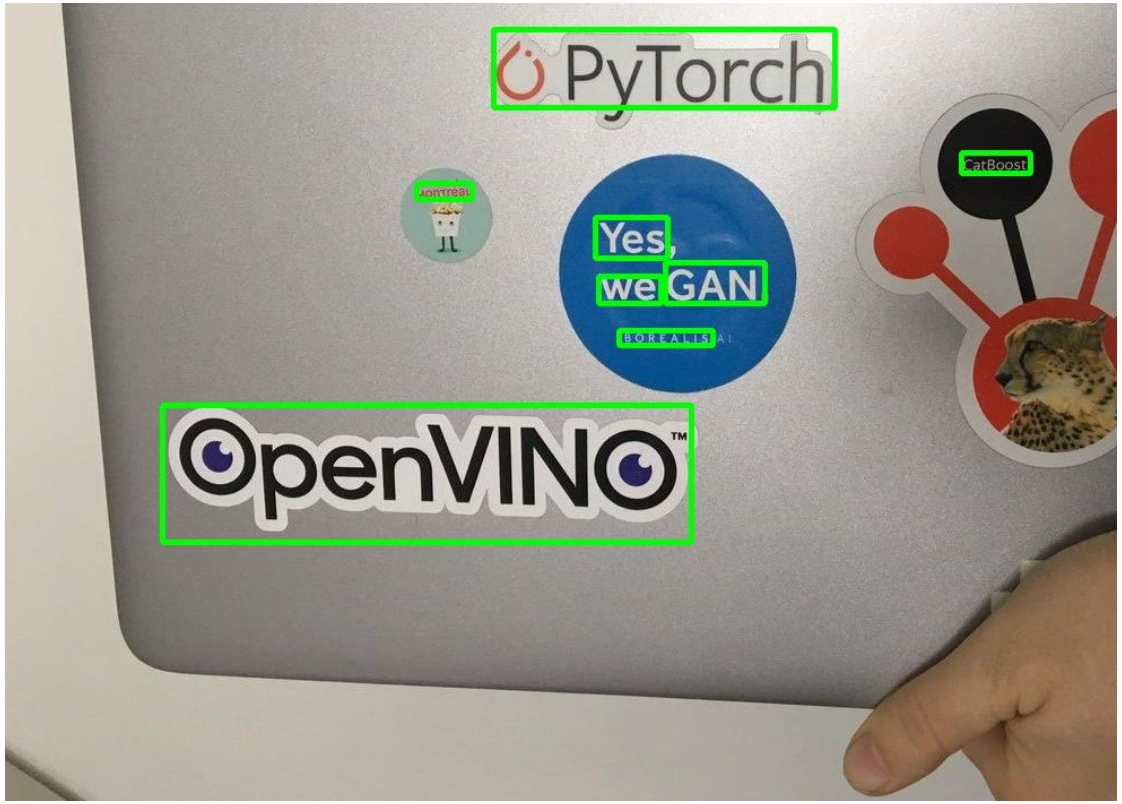

Enabling AI with the Open Model Zoo

220+ open-sourced models

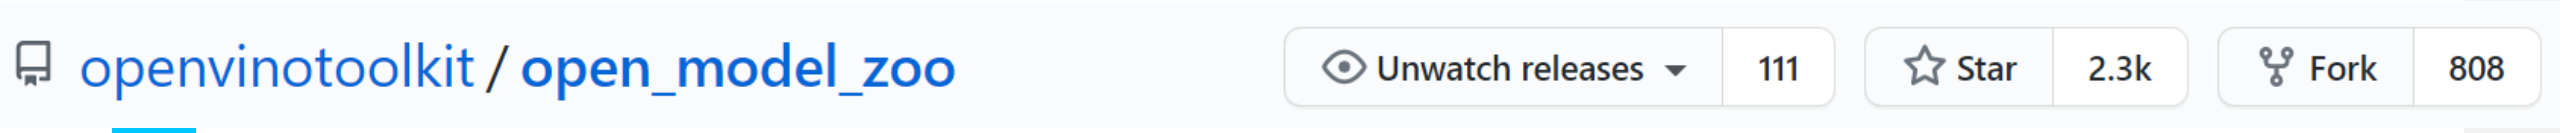

- **221 models** (123 public + 98 Intel)
  - **Areas:** Object Detection, Object Recognition, Reidentification, 2D/3D Semantic and Instance Segmentation, Pose estimation, Image Processing, Text Detection and Recognition, Action Recognition
  - **Precision:** FP32, FP16 + INT8, INT1 (for selected models)
- Models contain **a training source** and **compression tools**
- **37 demo applications** (16 native + 26 Python)

Seamless Optimization and Fine-tuning

Unified model templates for training/fine-tuning, evaluation, compression (NNCF) and OpenVINO export

- Object detection/Instance Segmentation and Text Spotting based on mmdetection
- Action recognition based on mmaction
- Image classification based on deep-object-reid

Why it's unique

- SOTA-level accuracy with edge networks
- Top performance on Intel HW
- Pre-trained models for common tasks (face/person/vehicle detection, etc.)

How I'm Using This

OpenVINO Training Extension

https://github.com/openvinotoolkit/training_extensions/tree/refactor

(Free Sample PyTorch-based Training Tutorials)

Today, you learned:

- Why the need for highly accurate and performant models **has encumbered developers** and what you can do about it
- What **tools and resources** are available for you to make your AI deployment journey more efficient
- How to **get started using** training extensions and frameworks with the OpenVINO™ toolkit
- **Live demos and Q&As** with our team of experts

Poll Time!

Answer in 30 seconds

Get additional help along the way

Online Trainings
Self-paced, Community and Certifications

Community Forum
Ask questions and engage with other developers

Documentation
Deep dive into concepts and how-to tutorials

Ready to get started?

- Download the Intel® Distribution of OpenVINO™ toolkit
- Efficiently train DL models ready for inferencing with the OpenVINO™ toolkit Training Extensions
- Optimize and fine-tune models with the Neural Network Compression Framework compatible with the OpenVINO™ toolkit

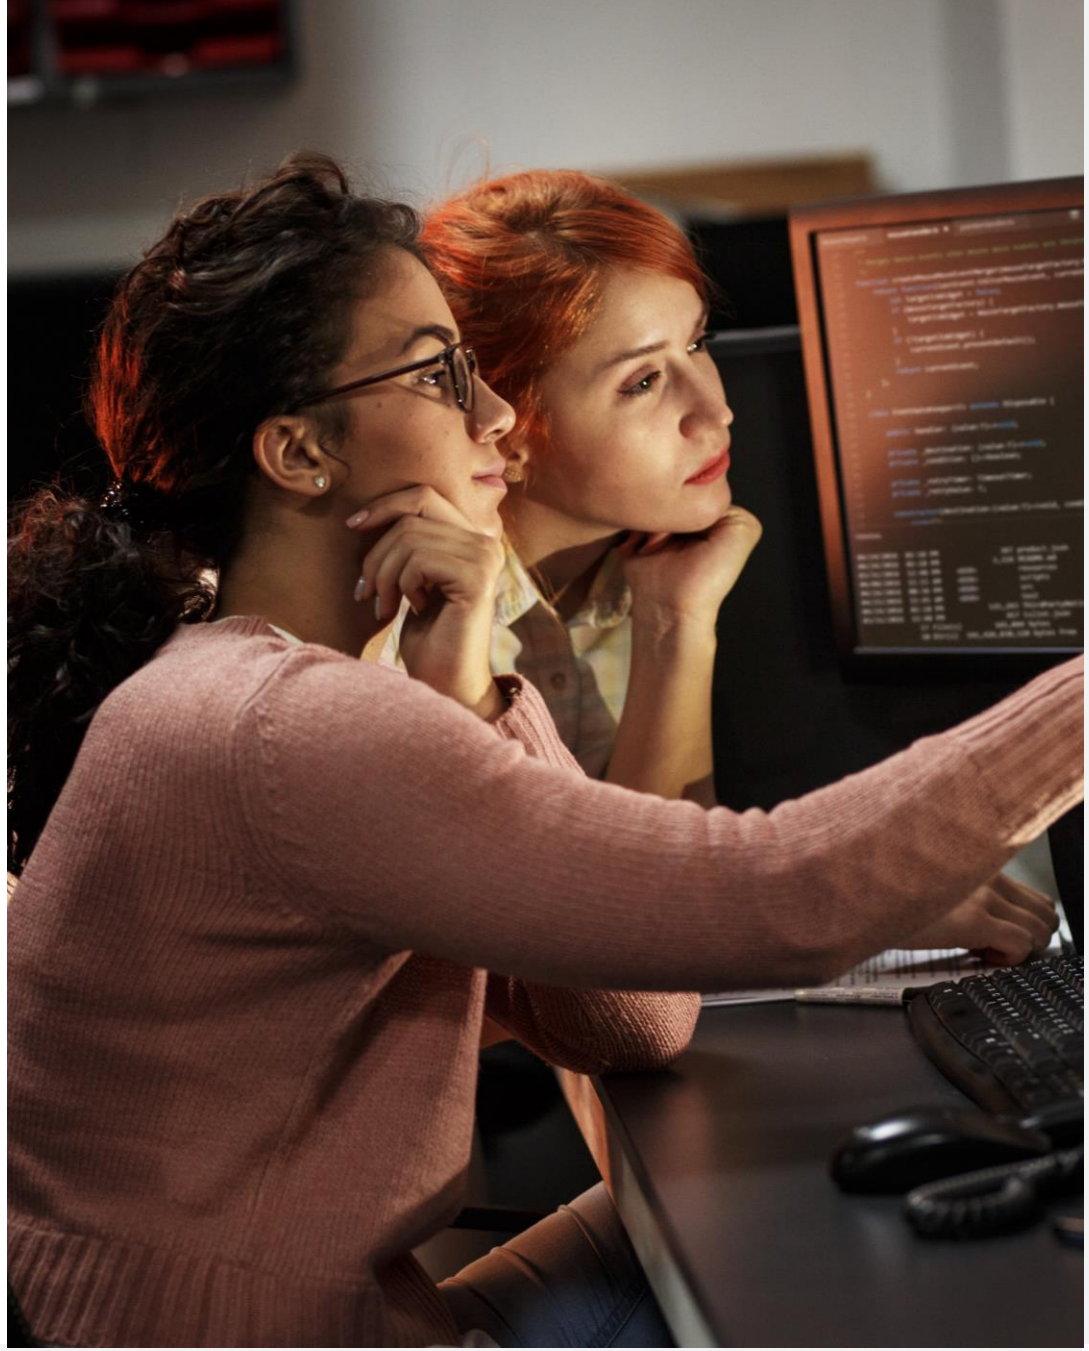

Notices & Disclaimers

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.


intel