Deep Learning on Smartphones: A Detailed Overview

Andrey Ignatov

andrey@vision.ee.ethz.ch

ETH Zurich, Switzerland

Deep Learning: A Historical Overview

- **1940 — 1950s:** Basic deep learning concepts
  Hebbian learning, Rosenblatt's perceptron
- **1970s:** Backpropagation algorithm
- **1980 — 1990s:** Fundamental deep learning models
  Neocognitron, TDNN, LeNet, RNN, LSTM
- **2009 — 2010:** First conventional GPU-accelerated ANNs
- **2011 — 2012:** Max-Pooling CNNs, AlexNet
- **2013 — 2015:** Common deep learning problems and architectures

Why Deep Learning on Smartphones?

- Image Classification
- Image Enhancement and Super-Resolution
- Optical Character Recognition
- Face Detection and Recognition
- Augmented Reality
- Natural Language Translation
- Voice Assistants
- Human Activity Recognition
- Gesture Recognition
- Music Tracking

Deep Learning on Smartphones:

Software

Mobile Deep Learning Libraries

- **2015:** TensorFlow Mobile

  Runs standard TensorFlow .pb models

  Acceleration: CPU only, Arm NEON instructions

Mobile Deep Learning Libraries

- **2015:** TensorFlow Mobile

  Runs standard TensorFlow .pb models

  Acceleration: CPU only, Arm NEON instructions

- **2015:** Mobile DNN *(Lane, N. D., & Georgiev, P.)*

  Acceleration: Qualcomm SDM800 only / Hexagon DSP

- **2016:** CNNdroid *(Latifi Oskouei, et al.)*

  Acceleration: GPU (RenderScript-based)

- **2017:** RSTensorFlow *(Alzantot, Moustafa, et al.)*

  Acceleration: GPU (RenderScript-based)

Mobile Deep Learning SDKs

- **2016, Qualcomm:** Snapdragon Neural Processing Engine (SNPE) SDK
  - Frameworks: TensorFlow, Caffe, Caffe2, ONNX
  - Acceleration: Snapdragon Hexagon DSPs + Adreno GPUs
- **2017, Huawei:** HiAI SDK
  - Frameworks: TensorFlow, Caffe
  - Acceleration: Kirin NPUs
- **2018, MediaTek:** NeuroPilot SDK
  - Frameworks: TensorFlow, TFLite, Caffe, Caffe2, MXNet, NNabla
  - Acceleration: MediaTek APUs + GPUs
- **2019, Samsung:** Exynos Deep Neural Network (EDEN) SDK
  - Frameworks: TBA
  - Acceleration: Exynos NPUs + GPUs

Mobile Deep Learning SDKs

- **2016, Qualcomm:** Snapdragon Neural Processing Engine (SNPE) SDK
  - Frameworks: TensorFlow, Caffe, Caffe2, ONNX
  - Acceleration: Snapdragon Hexagon DSPs + Adreno GPUs
- **2017, Huawei:** HiAI SDK
  - Frameworks: TensorFlow, Caffe
  - Acceleration: Kirin NPUs
- **2018, MediaTek:** NeuroPilot SDK
  - Frameworks: TensorFlow, TFLite, Caffe, Caffe2, MXNet, NNabla
  - Acceleration: MediaTek APUs + GPUs
- **2019, Samsung:** Exynos Deep Neural Network (EDEN) SDK
  - Frameworks: TBA
  - Acceleration: Exynos NPUs + GPUs

Problem: Each framework supports only the corresponding vendor's hardware!

Android Neural Network API (NNAPI): Android 8.1+

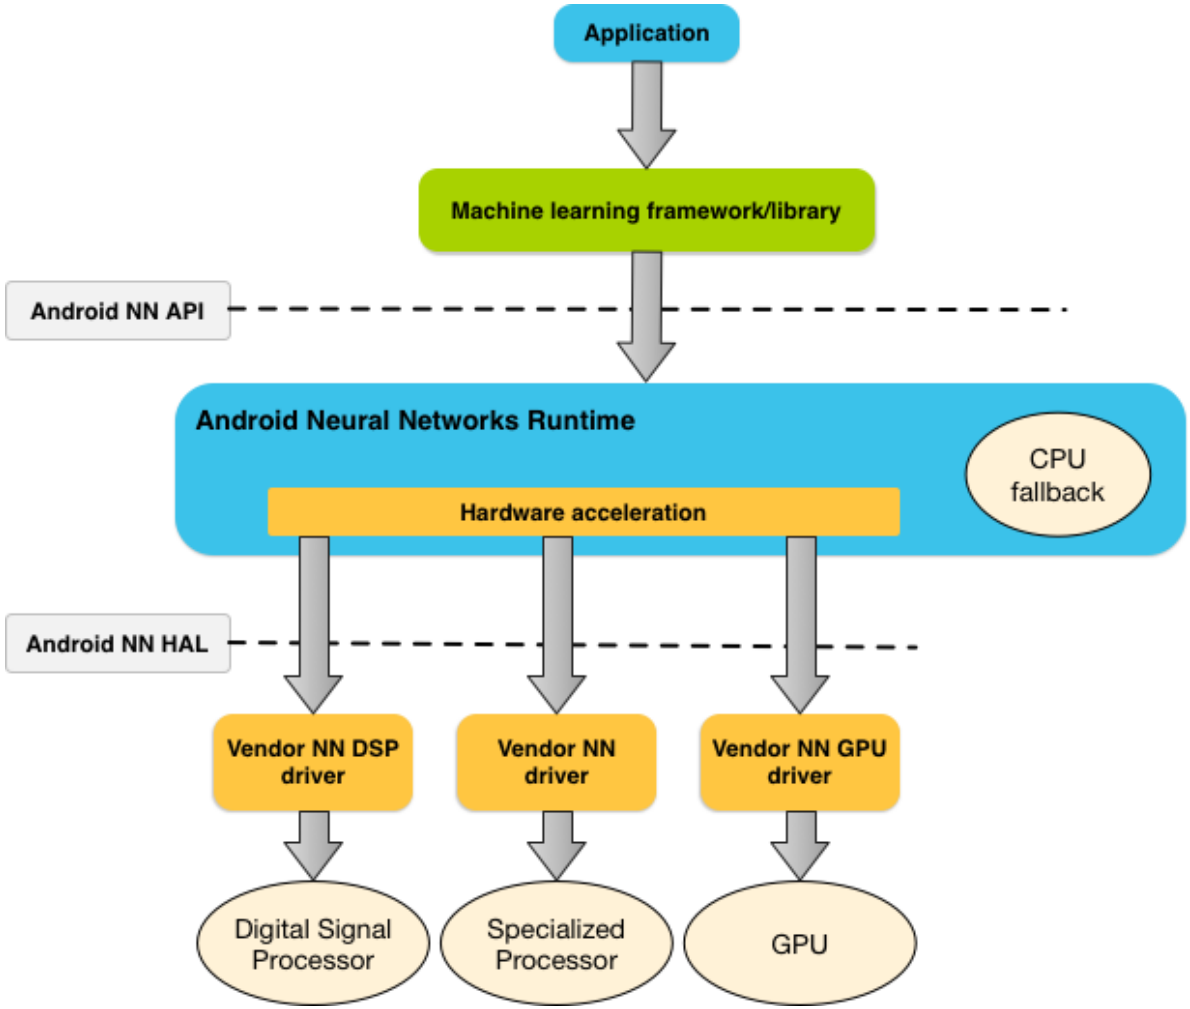

Intermediate layer between the higher-level DL frameworks and the device's hardware acceleration drivers.

TensorFlow Lite Framework

- Replaced *TensorFlow Mobile* library
- Supports *Android NNAPI*
- TensorFlow model has to be converted to *.tflite* format:

```
import tensorflow as tf
converter = tf.lite.TFLiteConverter.from_saved_model(saved_model_dir)
tflite_model = converter.convert()
open("mobile_model.tflite", "wb").write(tflite_model)
```

- Sample Android inference code:

```
// creating TFLite interpreter and enabling NNAPI
NnApiDelegate nnAPIDelegate = new NnApiDelegate();

private final Interpreter.Options tfLiteOptions = new Interpreter.Options();
tfLiteOptions.addDelegate(nnAPIDelegate);

private Interpreter tfLiteInterpreter = new Interpreter(loadModelFile(path_to_model), tfLiteOptions);

// run inference
tfLiteInterpreter.run(inputTensor, outputTensor);
```

TFlite + NNAPI: ISSUES

- Android NNAPI $\neq$ AI Hardware Acceleration
  - All computations are running by default on CPU!
  - Dedicated NNAPI hardware drivers are needed for NPUs, GPUs and DSPs

TFlite + NNAPI: ISSUES

- Android NNAPI $\neq$ AI Hardware Acceleration
  - All computations are running by default on CPU!
  - Dedicated NNAPI hardware drivers are needed for NPUs, GPUs and DSPs

---

- Android 8.1: **NNAPI 1.0** *(do not use!)*
  28 TensorFlow ops supported
  Numerous critical bugs
  Most devices were upgraded to Android 9+
- Android 9: **NNAPI 1.1**
  37 TensorFlow ops supported
- Android 10: **NNAPI 1.2**
  94 TensorFlow ops supported

TFlite + NNAPI: ISSUES

Version	Codename	API	Distribution
2.3.3 - 2.3.7	Gingerbread	10	0.3%
4.0.3 - 4.0.4	Ice Cream Sandwich	15	0.3%
4.1.x	Jelly Bean	16	1.2%
4.2.x		17	1.5%
4.3		18	0.5%
4.4	KitKat	19	6.9%
5.0	Lollipop	21	3.0%
5.1		22	11.5%
6.0	Marshmallow	23	16.9%
7.0	Nougat	24	11.4%
7.1		25	7.8%
8.0	Oreo	26	12.9%
8.1		27	15.4%
9	Pie	28	10.4%

TFlite + NNAPI: ISSUES

BUGS! BUGS! BUGS! BUGS!
BUGS! BUGS! BUGS! BUGS!
BUGS! BUGS! BUGS! BUGS!
BUGS! BUGS! BUGS! BUGS!
BUGS! BUGS! BUGS! BUGS!
BUGS! BUGS! BUGS! BUGS!

TFlite + NNAPI: ISSUES

Sample issues:

- TFLite 1.13: – Huge CPU performance degradation
- TFLite 1.14: – Random outputs during the first two inferences
- TFLite 1.15: – Increased RAM consumption (by up to 10 times)
- TFLite 1.15 / 2.0: – Transposed convolution op is broken
- TFLite 2.0 / 2.1: – Quantized models cannot run with acceleration

TFlite + NNAPI: ISSUES

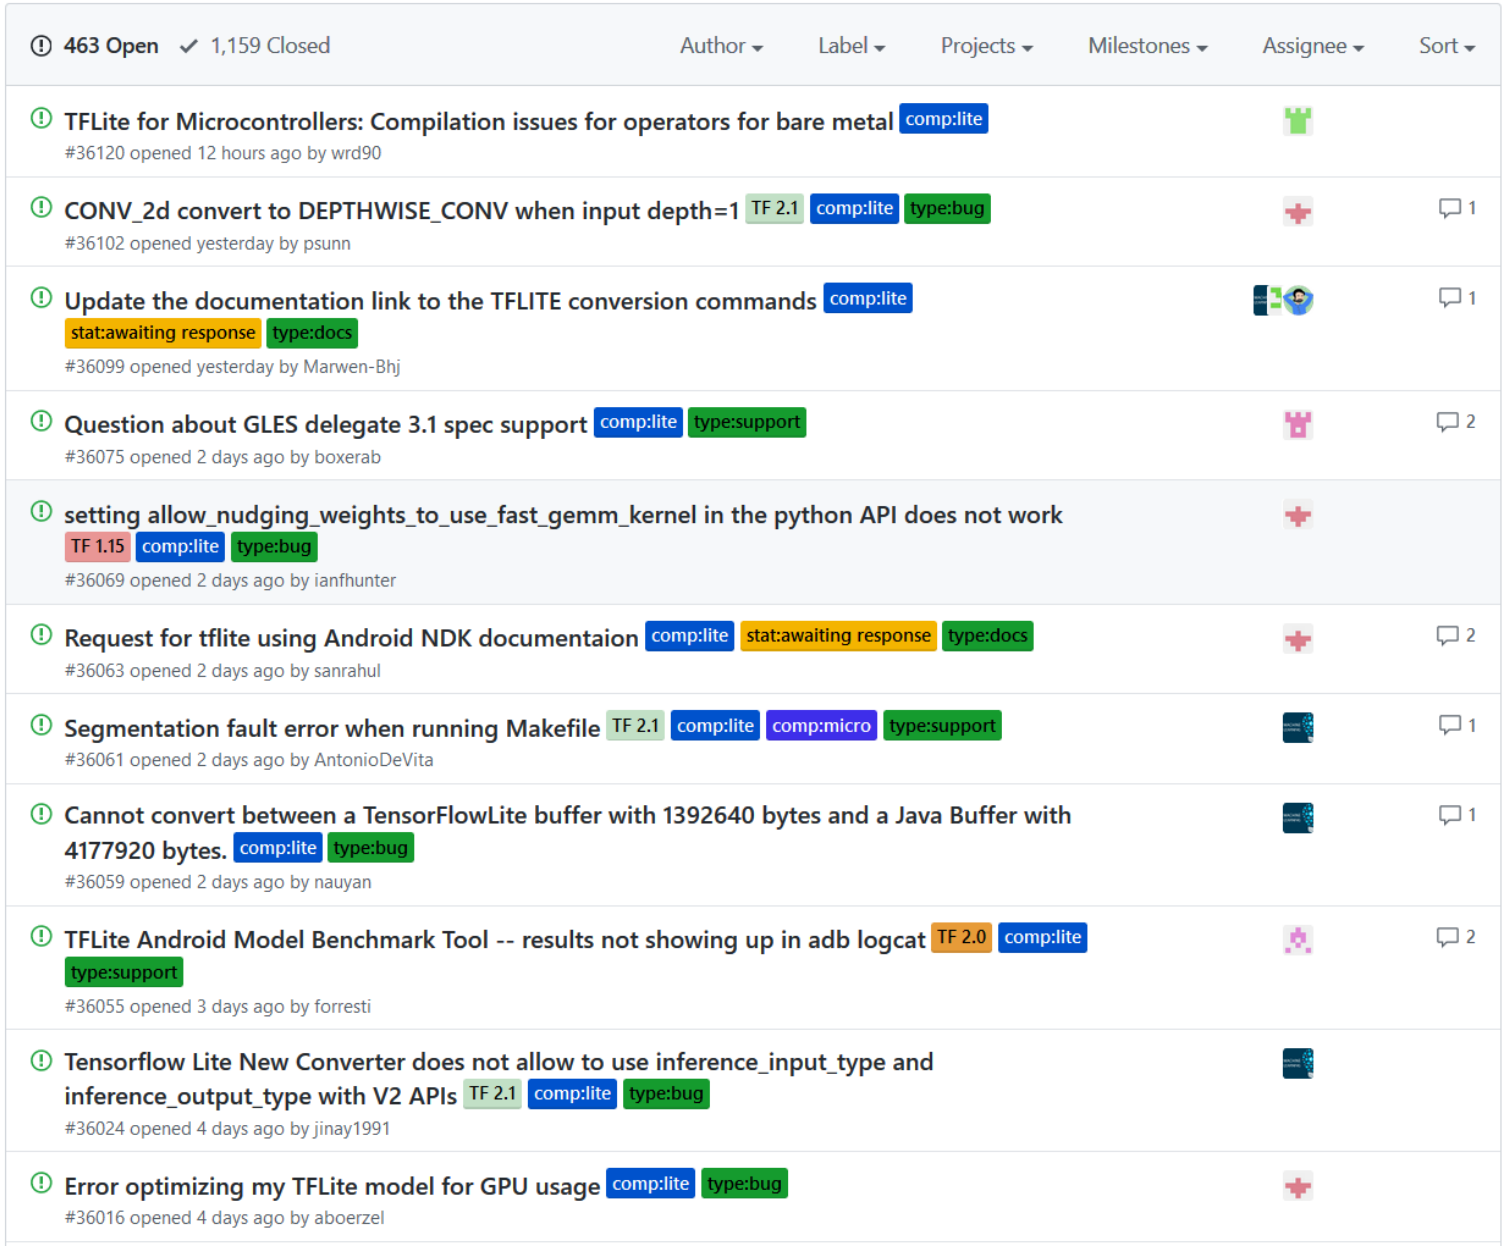

463 Open ✓ 1,159 Closed Author Label Projects Milestones Assignee Sort

- TFLite for Microcontrollers: Compilation issues for operators for bare metal comp:lite
  #36120 opened 12 hours ago by wrd90
- CONV_2d convert to DEPTHWISE_CONV when input depth=1 TF 2.1 comp:lite type:bug — 1
  #36102 opened yesterday by psunn
- Update the documentation link to the TFLITE conversion commands comp:lite stat:awaiting response type:docs — 1
  #36099 opened yesterday by Marwen-Bhj
- Question about GLES delegate 3.1 spec support comp:lite type:support — 2
  #36075 opened 2 days ago by boxerab
- setting allow_nudging_weights_to_use_fast_gemm_kernel in the python API does not work TF 1.15 comp:lite type:bug
  #36069 opened 2 days ago by ianfhunter
- Request for tflite using Android NDK documentaion comp:lite stat:awaiting response type:docs — 2
  #36063 opened 2 days ago by sanrahul
- Segmentation fault error when running Makefile TF 2.1 comp:lite comp:micro type:support — 1
  #36061 opened 2 days ago by AntonioDeVita
- Cannot convert between a TensorFlowLite buffer with 1392640 bytes and a Java Buffer with 4177920 bytes. comp:lite type:bug — 1
  #36059 opened 2 days ago by nauyan
- TFLite Android Model Benchmark Tool -- results not showing up in adb logcat TF 2.0 comp:lite type:support — 2
  #36055 opened 3 days ago by forresti
- Tensorflow Lite New Converter does not allow to use inference_input_type and inference_output_type with V2 APIs TF 2.1 comp:lite type:bug
  #36024 opened 4 days ago by jinay1991
- Error optimizing my TFLite model for GPU usage comp:lite type:bug
  #36016 opened 4 days ago by aboerzel

TensorFlow Lite Delegates

- Replace Android NNAPI:

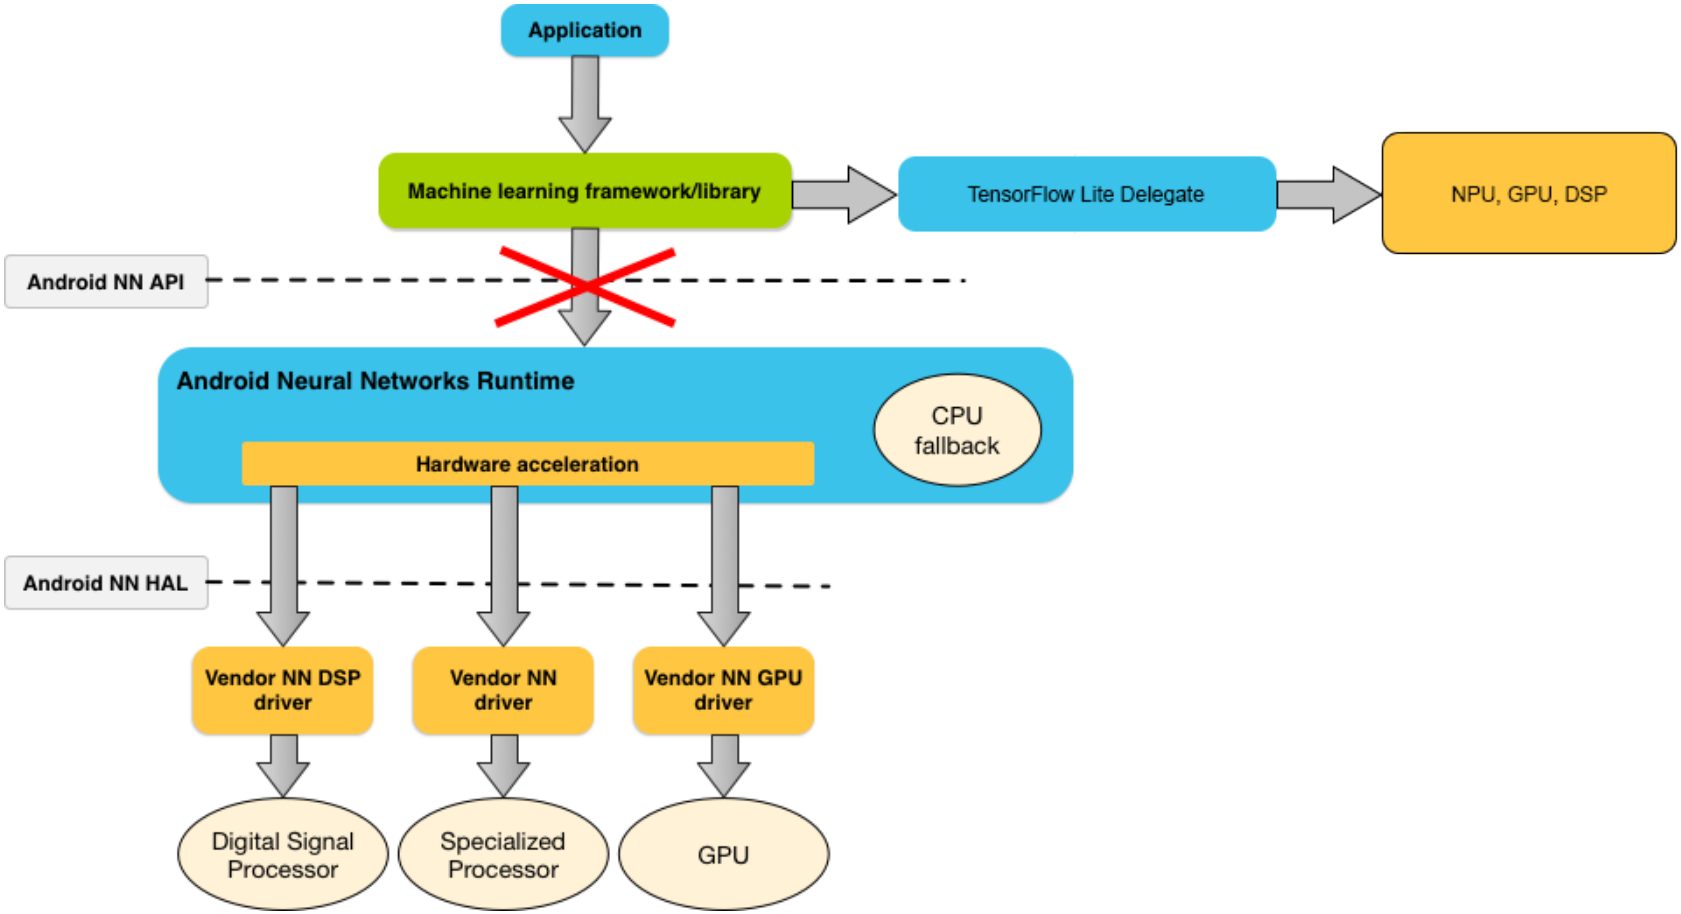

TensorFlow Lite Delegates

1) TFLite — changing one line of java code:

```
// creating TFLite interpreter and specifying the delegate
NnApiDelegate delegate = new NnApiDelegate();
GpuDelegate delegate = new GpuDelegate();

private final Interpreter.Options tfLiteOptions = new Interpreter.Options();
tfLiteOptions.addDelegate(delegate);

private Interpreter tfLiteInterpreter = new Interpreter(loadModelFile(path_to_model), tfLiteOptions);

// run inference
tfLiteInterpreter.run(inputTensor, outputTensor);
```

2) Advantages:

- Overcomes all Android NNAPI limitations
- Better performance and optimization
- Independent of Android OS version
- Support for legacy devices

Mobile Deep Learning Software

Summary:

- TensorFlow Mobile
  - Deprecated
  - CPU only inference, no NNAPI support
  - Can run the majority of TensorFlow models
- TensorFlow Lite
  - Much smaller number of ops
  - NNAPI support
    1) Be prepared for numerous bugs and issues
    2) Use only on Android 9+
    3) In general, better use custom TFLite delegates instead
- PyTorch Mobile
  - The first experimental release is available
  - CPU inference

Deep Learning on Smartphones:

Hardware

Mobile Deep Learning Hardware

- **2016, Qualcomm:** Snapdragon 820
  Accelerator: Hexagon V6 68x DSP
- **2017, Huawei:** Kirin 970
  Accelerator: NPU / Cambricon
- **2017, Samsung:** Exynos 8895
  Accelerator: Vision Processing Unit (not accessible)
- **2017, Google:** Pixel 2
  Accelerator: Pixel Visual Core (not accessible)
- **2018, MediaTek:** Helio P60
  Accelerator: AI Processing Unit

Can now distinguish 4 generations of mobile SoCs based on their AI performance

Generation 1

- Legacy chipsets
- No Android drivers for running AI models
- AI acceleration only with special SDKs or GPU-based libraries

  → Most neural networks = vector/matrix multiplications

  → Can run on mobiles GPUs supporting *OpenGL ES* or *OpenCL*

  → Usually much slower compared to CPU inference

Generation 2

– Released after 2017, support Android NNAPI

– Typical AI performance for the year 2018

- Qualcomm: Snapdragon 845 (Hexagon 685 + Adreno 630);
  Snapdragon 710 (Hexagon 685 + Adreno 616);
  Snapdragon 670 (Hexagon 685 + Adreno 615);
- HiSilicon: Kirin 970 (NPU, Cambricon, *floating-point models only*);
- Samsung: Exynos 9810 (Mali-G72 MP18);
  Exynos 9610 (Mali-G72 MP3);
  Exynos 9609 (Mali-G72 MP3);
- MediaTek: Helio P70 (APU 1.0 + Mali-G72 MP3);
  Helio P60 (APU 1.0 + Mali-G72 MP3);
  Helio P65 (Mali-G52 MP2).

Generation 3

– Released after 2018, support Android NNAPI

– Typical AI performance for the year 2019

- Qualcomm: Snapdragon 855 / 855+ (Hexagon 690 + Adreno 640);
  Snapdragon 730 (Hexagon 688 + Adreno 618);
  Snapdragon 675 (Hexagon 685 + Adreno 612);
  Snapdragon 665 (Hexagon 686 + Adreno 610);
- HiSilicon: Kirin 980 (dual-core NPU, Cambricon, *floating-point models only*);
- Samsung: Exynos 9825 / 9820 (Mali-G76 MP12 + NPU/*not yet accessible*);
- MediaTek: Helio P90 (APU 2.0);
  Helio G90 (APU 1.0 + Mali-G76 MP4).

Generation 4

– Released after 2019, support Android NNAPI

– Typical AI performance for the year 2020

- Qualcomm: Snapdragon 865 (Hexagon 698 + Adreno 650);
  Snapdragon 765 (Hexagon 696 + Adreno 620);
- HiSilicon: Kirin 990 5G (NPU, 3 cores, *Da Vinci series*);
  Kirin 990 (NPU, 2 cores, *Da Vinci series*);
  Kirin 810 (NPU, *Da Vinci series*);
- Samsung: Exynos 990 (*TBA*);
  Exynos 980 (Mali-G76 MP5 + NPU/*not yet accessible*);
- MediaTek: Dimensity 1000 (APU 3.0);
- Unisoc: Tiger T710 (NPU, 2 cores, Imagination Technologies);

Deep Learning on Smartphones:

Benchmarking

AI Benchmark (ETHZ)

- Android application for measuring smartphones' AI performance
- First public version: May 2018 *(12 public releases, last: 3.0.2)*
- TensorFlow Lite + Android NNAPI | TensorFlow Lite Delegates
- 11 Sections / 21 Tests:

- Section 1. Classification, MobileNet-V2: CPU (FP16) + NPU / GPU (FP16) + NPU / GPU / DSP (INT8)
- Section 2. Classification, Inception-V3: CPU (FP16) + NPU / GPU (FP16) + NPU / GPU / DSP (INT8)
- Section 3. Face Recognition, Inception-ResNet-V1: CPU (INT8) + NPU / GPU (FP16) + NPU / GPU / DSP (INT8)
- Section 4. Playing Atari Games, LSTM: CPU (FP16)
- Section 5. Deblurring, SRCNN: NPU / GPU (FP16) + NPU / GPU / DSP (INT8)
- Section 6. Super-Resolution, VGG19: NPU / GPU (FP16) + NPU / GPU / DSP (INT8)
- Section 7. Super-Resolution, SRGAN: CPU (FP32) + CPU (INT8)
- Section 8. Bokeh Simulation, U-Net: CPU (FP32)
- Section 9. Semantic Segmentation, ICNet: NPU / GPU (FP32) x 2 - two CNNs running in parallel
- Section 10. Image Enhancement, DPED ResNet: NPU / GPU (FP16) + NPU / GPU (FP32)
- Section 11. Memory limits, SRCNN: NPU / GPU (FP16)

AI Benchmark (ETHZ): Measurements

- FLOAT-16 performance (NPU, APU, GPU)
- FLOAT-32 performance (NPU, APU, GPU)
- INT-8 performance (NPU, APU, GPU, DSP)
- CPU single- and multi-thread performance (FP32, FP16, INT8)
- Single / throughput inference times
- Memory / RAM performance
- Initialization time
- Accuracy

AI Benchmark (ETHZ): Visualization

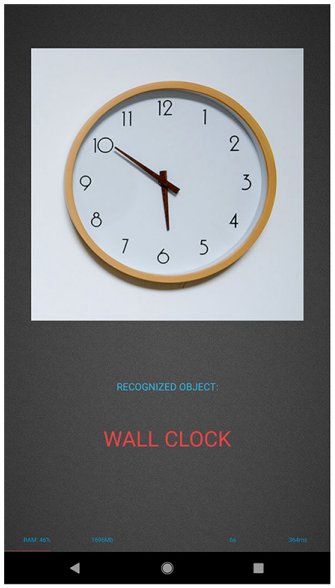

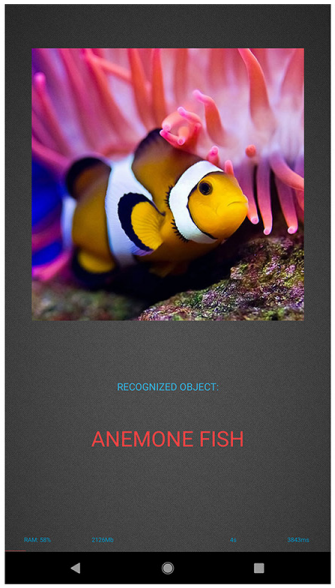

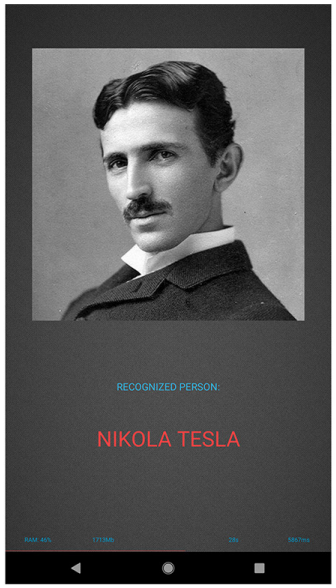

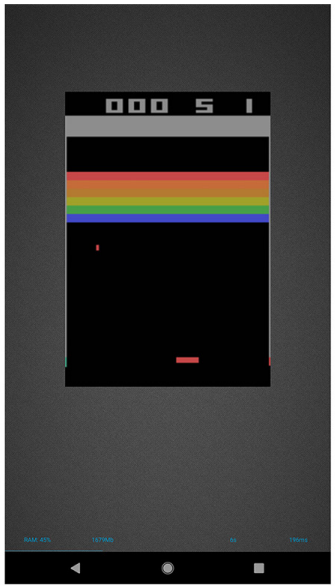

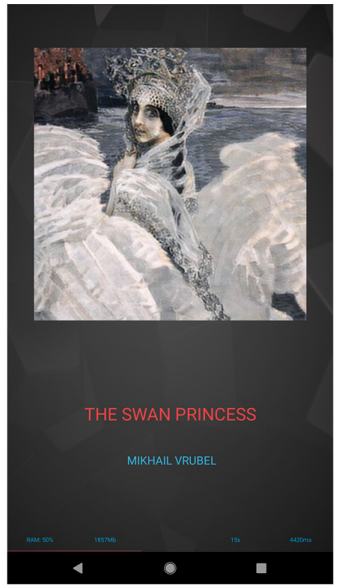

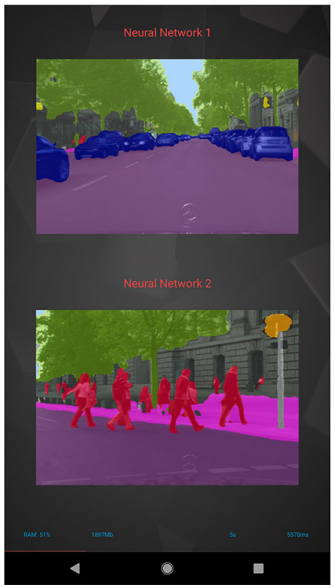

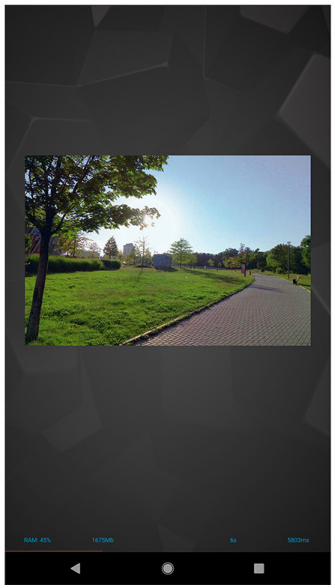

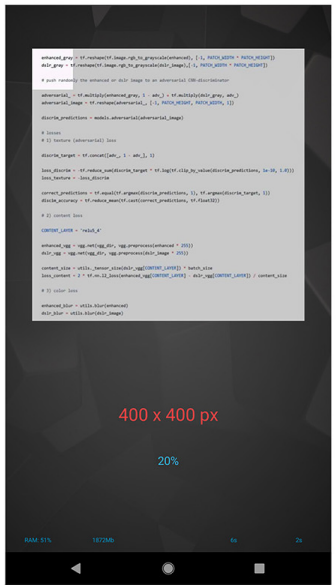

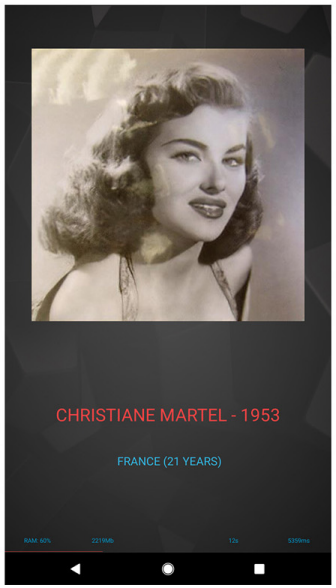

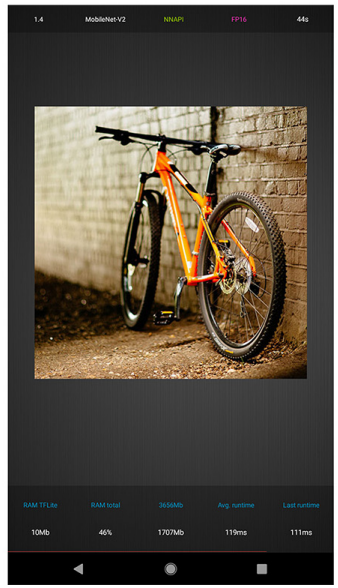

AI Benchmark (ETHZ): Visualization

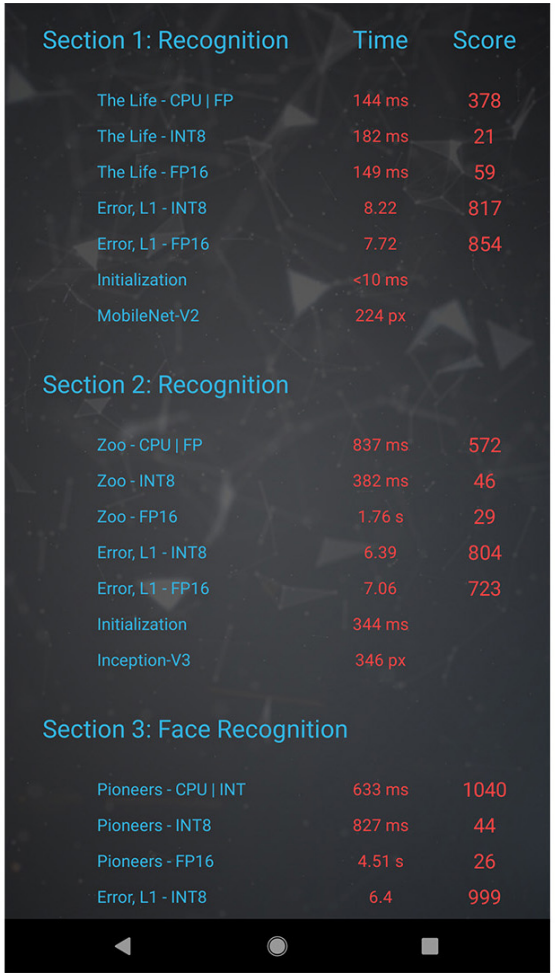

AI Benchmark: Results – FP16

SoC Model	MobileNet v2, ms	Inception v3, ms	Inc-ResNet v1, ms	SRCNN, ms	VGG-19, ms	DPED, ms	Relative Perf.
HiSilicon Kirin 990 5G	6	18	37	36	42	19	100%
MediaTek Dimensity 1000	4.2	31	73	28	73	25	78%
HiSilicon Kirin 990	7.2	23	47	44	65	26	76%
HiSilicon Kirin 810	10	34	82	72	122	42	47%
Unisoc Tiger T710	13	35	80	76	135	43	43%
HiSilicon Kirin 980	21	48	116	94	157	55	32%
Snapdragon 865	13	60	151	63	204	91	30%
Snapdragon 855 Plus	14	67	167	79	228	110	26%
Exynos 9825 Octa	19	72	156	127	166	73	26%
Mediatek Helio P90	8.3	101	263	75	309	66	26%
Exynos 9820 Octa	20	73	163	133	171	84	25%
Snapdragon 855	17	70	172	80	230	110	25%
Snapdragon 845	25	89	213	148	287	114	19%
HiSilicon Kirin 970	43	69	1514	141	235	83	14%
Exynos 980	32	160	328	273	400	183	12%
Snapdragon 730	25	158	427	202	609	193	11%
Snapdragon 670	31	234	636	300	934	318	8%
Snapdragon 710	30	207	565	267	816	287	8%
Mediatek Helio G90T	28	202	502	394	906	659	7%
Snapdragon 675	40	316	900	525	1253	351	5%
Exynos 9810 Octa	72	209	488	1574	843	787	4%
Exynos 8895 Octa	63	216	497	1785	969	909	4%
Mediatek Helio P70	66	374	932	1096	865	764	4%
Mediatek Helio P65	50	344	918	750	1672	940	4%
Snapdragon 665	50	483	1292	678	2174	553	4%
Mediatek Helio P60	68	353	948	1896	889	1439	3%
Exynos 9609	61	444	1230	1661	1448	731	3%
Exynos 9610	77	459	1244	1651	1461	773	3%
Exynos 8890 Octa	98	447	1012	2592	1062	855	3%
Snapdragon 835	181	786	1515	1722	3754	1317	1%
GeForce GTX 1080 Ti	1.5	4.5	9.5	4.7	10	4.6	449%
GeForce GTX 950	3.9	15	38	23	47	20	115%
Nvidia Tesla K40c	3.7	16	38	22	60	20	111%
GeForce GT 1030	9.3	31	81	44	97	47	53%
GeForce GT 740	12	89	254	238	673	269	14%
Intel Core i7-9700K	4.8	23	72	49	133	72	55%
Intel Core i7-7700K	7.4	42	121	75	229	100	34%
Intel Core i7-3770K	12	125	345	209	729	242	13%
Intel Core i7-950	36	287	728	448	1219	515	6%

AI Benchmark: Results – INT8

SoC Model	AI Accelerator	MobileNet v2, ms	Inception v3, ms	Inc-ResNet v1, ms	SRCNN, ms	VGG-19, ms	Relative Perf.
MediaTek Dimensity 1000	APU 3.0 (6 cores)	**2.9**	**12**	**23**	**13**	**31**	100%
Unisoc Tiger T710	NPU (2 cores, Imagination)	5	17	38	20	53	62%
Snapdragon 855 Plus	DSP (Hexagon 690)	6	16	35	25	46	60%
Snapdragon 865	DSP (Hexagon 698)	6.9	17	35	25	43	59%
HiSilicon Kirin 990 5G	NPU (3 cores, Da Vinci)	6.5	20	37	38	39	53%
Snapdragon 855	DSP (Hexagon 690)	8.6	20	39	27	47	51%
Mediatek Helio P90	APU 2.0	4	23	38	22	147	49%
HiSilicon Kirin 990	NPU (2 cores, Da Vinci)	7.9	24	48	46	60	41%
Snapdragon 675	DSP (Hexagon 685)	10	34	73	53	103	29%
Snapdragon 730	DSP (Hexagon 688)	12	46	81	66	106	25%
Snapdragon 845	DSP (Hexagon 685)	12	44	83	66	109	25%
Snapdragon 710	DSP (Hexagon 685)	12	46	86	67	111	24%
Snapdragon 670	DSP (Hexagon 685)	13	48	88	68	115	23%
Snapdragon 665	DSP (Hexagon 686)	13	52	118	94	192	18%
Exynos 9825 Octa	GPU (Mali-G76 MP12)	21	68	134	73	202	16%
Exynos 9820 Octa	GPU (Mali-G76 MP12)	22	69	142	78	216	15%
Mediatek Helio G90T	GPU (Mali-G76 MP4)	15	63	137	107	308	14%
HiSilicon Kirin 810	NPU (Da Vinci)	25	98	160	116	172	13%
Mediatek Helio P70	APU 1.0	26	89	181	163	474	9%
Mediatek Helio P60	APU 1.0	27	89	181	164	475	9%
Exynos 980	GPU (Mali-G76 MP5)	32	128	278	177	504	7%
Exynos 8895 Octa	GPU (Mali-G71 MP20)	44	118	228	416	596	6%
Exynos 9810 Octa	GPU (Mali-G72 MP18)	45	166	360	539	852	4%
Mediatek Helio P65	GPU (Mali-G52 MP2)	45	249	551	610	1269	3%
Snapdragon 835	none	136	384	801	563	1525	2%
Exynos 9609	GPU (Mali-G72 MP3)	50	383	937	1027	2325	2%
Exynos 9610	GPU (Mali-G72 MP3)	52	380	927	1024	2322	2%
Exynos 8890 Octa	GPU (Mali-T880 MP12)	43	388	932	1133	2095	2%

BENCHMARK NEWS RANKING AI-TESTS Google play Research ABOUT

From Kirin 970 to Snapdragon 855 Plus: Performance Review of All SoCs with AI capabilities

Phones | Mobile SoCs

Performance Ranking

Desktop GPUs and CPUs

View Detailed Results

Model	CPU	RAM	Year	Android	Updated	CPU Q AI Score	CPU F AI Score	QUANT Score	QUANT Accuracy	FP16 Score	FP16 Accuracy	FP32 Score	FP PAR Score	Memory	Accuracy	AI-Score
Huawei Mate 30 Pro 5G	HiSilicon Kirin 990 5G	8GB	2019	10	9.19	1575	3744	6899	79	38042	93	716	649	2000	88	76206 1,6
Honor V30 Pro 5G	HiSilicon Kirin 990 5G	8GB	2019	10	1.20	1619	3890	6658	79	38190	93	709	630	2000	88	76120 1,6
Huawei Mate 30 5G	HiSilicon Kirin 990 5G	8GB	2019	10	1.20	1604	3813	6792	79	37703	93	703	635	2000	88	75524 1,6
Huawei Mate 30	HiSilicon Kirin 990	8GB	2019	10	1.20	1552	3800	5458	79	29610	93	703	630	2000	88	61560 1,6
Huawei Mate 30 Pro	HiSilicon Kirin 990	8GB	2019	10	10.19	1552	3750	5385	79	29171	93	713	634	2000	88	60748 1,6
Honor 20S	HiSilicon Kirin 810	8GB	2019	9	9.19	1155	2808	1714	80	17977	95	451	396	1600	90	35130 1,6
Honor 9X Pro	HiSilicon Kirin 810	6GB	2019	9	9.19	1123	2568	1704	80	17896	95	459	366	1600	90	34495 1,6
Huawei Nova 5	HiSilicon Kirin 810	6GB	2019	9	9.19	1109	2388	1712	80	18005	95	461	360	1600	90	34432 1,6
Huawei Nova 5i Pro	HiSilicon Kirin 810	8GB	2019	9	7.19	1150	2657	1699	80	17571	95	448	368	1600	90	33970 1,6
Google Pixel 4	Snapdragon 855	6GB	2019	10	1.20	2073	3212	8792	94	8883	87	920	859	1000	90	33289 1
Asus ROG Phone II	Snapdragon 855 Plus	12GB	2019	9	10.19	2253	4147	6394	40	9211	30	1356	1073	1000	34	33185 1
Realme X2 Pro	Snapdragon 855 Plus	12GB	2019	9	1.20	2303	4243	6198	54	9034	37	1319	1088	1000	42	32834 1
Google Pixel 4 XL	Snapdragon 855	6GB	2019	10	1.20	1978	3149	8577	94	8945	87	913	804	1000	90	32793 1
Samsung Galaxy S10+	Snapdragon 855	8GB	2019	10	1.20	2055	4204	6945	59	9031	49	957	950	1000	52	32664 1
Samsung Galaxy Note10+	Snapdragon 855	12GB	2019	10	1.20	2096	4310	6870	59	8887	49	977	893	1000	52	32544 1
Xiaomi Redmi K20 Pro	Snapdragon 855	8GB	2019	10	1.20	1911	3534	6743	85	9404	87	954	855	1000	86	31740 1
Xiaomi Mi 9	Snapdragon 855	8GB	2019	10	1.20	1973	3603	6745	85	9286	87	969	839	1000	86	31725 1
Xiaomi Mi 9T Pro	Snapdragon 855	8GB	2019	10	1.20	1943	3558	6687	85	9281	87	964	846	1000	86	31552 1

Project website: ai-benchmark.com

AI Benchmark: Research

- AI Benchmark: All About Deep Learning on Smartphones in 2019

  Andrey Ignatov (ETH Zurich), Radu Timofte (ETH Zurich), Andrei Kulik (Google Research), Seungsoo Yang (Samsung, Inc), Ke Wang (Huawei, Inc.), Felix Baum (Qualcomm, Inc.), Max Wu (MediaTek, Inc.), Lirong Xu (Unisoc, Inc.), Luc Van Gool (ETH Zurich)

- AI Benchmark: Running Deep Neural Networks on Android Smartphones

  Andrey Ignatov (ETH Zurich), Radu Timofte (ETH Zurich), William Chou (Qualcomm, Inc.), Ke Wang (Huawei, Inc.), Max Wu (MediaTek, Inc.), Tim Hartley (Arm, Inc.), Luc Van Gool (ETH)

- PIRM Challenge on Perceptual Image Enhancement on Smartphones

  Andrey Ignatov (ETH Zurich), Radu Timofte (ETH Zurich), et al.

◀ Thank you for your attention! ▶

Research Paper

Project Website

Google Play

Desktop Version