

EdgeDRNN: Enabling Low-latency Recurrent Neural Network Edge Inference

Chang Gao*, Antonio Rios-Navarro[†], Xi Chen*, Tobi Delbruck*, Shih-Chii Liu*

*Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland

[†]Robotic and Technology of Computers Lab, Universidad de Sevilla, Seville, Spain

{chang, xi, tobi, shih}@ini.uzh.ch, {arios}@us.es

Abstract—This paper presents a Gated Recurrent Unit (GRU) based recurrent neural network (RNN) accelerator called EdgeDRNN designed for portable edge computing. EdgeDRNN adopts the spiking neural network inspired delta network algorithm to exploit temporal sparsity in RNNs. It reduces off-chip memory access by a factor of up to 10x with tolerable accuracy loss. Experimental results on a 10 million parameter 2-layer GRU-RNN, with weights stored in DRAM, show that EdgeDRNN computes them in under 0.5 ms. With 2.42 W wall plug power on an entry level USB powered FPGA board, it achieves latency comparable with a 92 W Nvidia 1080 GPU. It outperforms NVIDIA Jetson Nano, Jetson TX2 and Intel Neural Compute Stick 2 in latency by 6X. For a batch size of 1, EdgeDRNN achieves a mean effective throughput of 20.2 GOP/s and a wall plug power efficiency that is over 4X higher than all other platforms.

Index Terms—edge computing, FPGA, embedded system, deep learning, RNN, GRU, delta network

I. INTRODUCTION

Recurrent Neural Networks (RNN) are a subset of deep neural networks that are particularly useful for regression and classification tasks involving time series inputs. Gated RNNs which use Long Short-Term Memory units (LSTM) [1] and Gated-Recurrent Unit (GRU) [2] are used to overcome the vanishing gradient problem frequently encountered during RNN training with backpropagation through time. RNN models are frequently used in state-of-the-art models for automatic speech recognition tasks [3], [4].

In edge computing, computations are done locally on end-user devices to reduce latency and protect privacy [5]. RNNs achieve high accuracy at the cost of large memory footprint and expensive computation. RNNs are usually computed on the cloud with results sent to edge devices, which introduces high and variable latency, making it hard to guarantee real time performance for human computer interaction, robotics, and control applications. Previous work exploits weight pruning [6] [7], structured weight matrix [8], and temporal sparsity [9] to accelerate RNN computation by reducing the memory bottleneck of RNNs. However, these works used expensive FPGA boards with greater than 15 W power consumption and did not target portable edge devices with low latency demands and a limited power budget.

This paper describes an RNN accelerator for edge applications. The accelerator exploits temporal sparsity using the delta network (DeltaGRU) [10] algorithm. It achieves sub-

millisecond inference of big multi-layer RNNs comparable with a desktop-level GPU, but with 38 times less power.

II. GATED-RECURRENT UNIT & DELTA NETWORK

The equations for a GRU layer of M neurons and N -dimensional input are given as:

$$\begin{aligned} \mathbf{r}_t &= \sigma(\mathbf{W}_{ir}\mathbf{x}_t + \mathbf{W}_{hr}\mathbf{h}_{t-1} + \mathbf{b}_r) \\ \mathbf{u}_t &= \sigma(\mathbf{W}_{iu}\mathbf{x}_t + \mathbf{W}_{hu}\mathbf{h}_{t-1} + \mathbf{b}_u) \\ \mathbf{c}_t &= \tanh(\mathbf{W}_{ic}\mathbf{x}_t + \mathbf{r}_t \odot (\mathbf{W}_{hc}\mathbf{h}_{t-1}) + \mathbf{b}_c) \\ \mathbf{h}_t &= (1 - \mathbf{u}_t) \odot \mathbf{c}_t + \mathbf{u}_t \odot \mathbf{h}_{t-1} \end{aligned} \quad (1)$$

where $\mathbf{r}, \mathbf{u}, \mathbf{c} \in \mathcal{R}^M$ are respectively the reset gate, the update gate and the cell state. $W_i \in \mathcal{R}^{M \times N}$, $W_h \in \mathcal{R}^{M \times M}$ are weight matrices and $b \in \mathcal{R}^M$ are bias vectors. σ denotes the logistic sigmoid.

Inspired by spiking neural networks, the DeltaGRU [10] reduces operations in GRU-RNNs while maintaining high prediction accuracy. In DeltaGRU, weights are multiplied with the delta vectors $\Delta\mathbf{x}_t = \mathbf{x}_t - \mathbf{x}_{t-1}$, $\Delta\mathbf{h}_{t-1} = \mathbf{h}_{t-1} - \mathbf{h}_{t-2}$ between the current and the previous time steps and then added to a memory term $\mathbf{M}_t = \sum_{i=0}^{t-1} (\mathbf{W}\Delta\mathbf{x}_i + \mathbf{W}\Delta\mathbf{h}_{i-1})$ that is the accumulation of all previous products. The initial states are $\mathbf{M}_0 = \mathbf{b}$, $\mathbf{x}_{-1} = 0$ and $\mathbf{h}_{-1} = \mathbf{h}_{-2} = 0$.

By setting the elements of a delta vector to zero when their individual values are less than a defined **Delta Threshold** Θ , the number of matrix-vector multiply-and-accumulate (MAC) operations is reduced by 5X to 100X, depending on the dynamics of the input and hidden units [10]. It allows skipping entire columns of the weight matrix. That way, DRAM weight memory reads are still in efficient burst mode.

III. EDGEDRNN ACCELERATOR

A. Accelerator Design

The design of EdgeDRNN aims to achieve low-latency RNN inference with batch size of 1, which are needed for real-time operation with minimum latency. 2D arithmetic unit arrays are not suitable here due to limited weight reuse, scarce on-chip memory resources and narrow external memory interface on embedded systems like MiniZed. The vector processing element (PE) array in EdgeDRNN is able to fully utilize the external memory bandwidth.

Fig. 1a shows the design of the EdgeDRNN accelerator. The number of PEs, K , in EdgeDRNN is $K =$

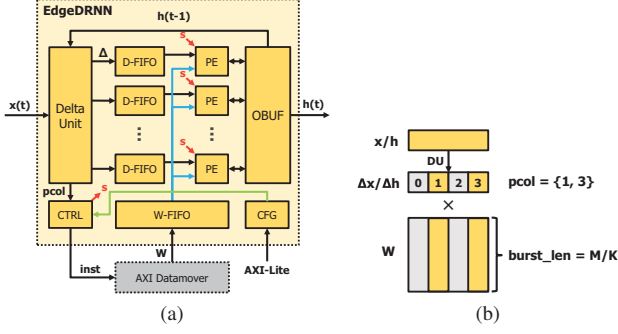


Fig. 1: (a) EdgeDRNN accelerator architecture; (b) Flow chart of the sparse matrix-vector multiplication.




Fig. 2: Architecture of the EdgeDRNN processing element (PE).




Fig. 3: Top-level diagram of the EdgeDRNN implementation on the MiniZed development board.

$BW_{\text{DRAM}}/BW_{\text{W}} = 64/8 = 8$, where $BW_{\text{W}} = 8$ is the weight precision and $BW_{\text{DRAM}} = 64$ the external memory interface bit-width. EdgeDRNN can be configured to support 1, 2, 4, 8, 16-bit fixed-point weights and 16-bit fixed-point activations; in this paper we used only 8-bit weights. The delta unit (DU) includes BRAM memory that records previous states x_{t-1} and h_{t-2} to be used for calculating delta vectors Δx and Δh . The DU checks one element in a delta vector per cycle. Elements that exceed Θ result in non-zero elements and are broadcast to all D-FIFOs that drive PEs. As shown in Figs. 1a and 1b, DU computes column pointers (pcol) to non-zero delta vector elements that are sent to the global controller (CTRL). Using pcol , CTRL generates instructions, containing the physical start address of a weight column and the burst length given in Fig. 1b, to control the AXI Datamover IP to fetch weights (biases are appended to weights). On MiniZed, DRAM data moves through the PL's DMA and Datamover.

TABLE I: Resource utilization of MiniZed.

	LUT	LUTRAM	FF	BRAM (36Kb)	DSP
Available	14400	6000	28800	50	66
Used	10464	552	11665	33	9
Percentage	72.67%	9.20%	40.50%	66%	13.64%




Fig. 4: Mean effective throughput and word error rate evaluated on the TIDIGITS test set versus various delta thresholds (shown as hex values corresponding to 0~0.5 floating point threshold) used in both training and inference of a 2L-768H-DeltaGRU network.

TABLE II: Word error rate (WER) of GRU and DeltaGRU networks trained with $\Theta = 0x40$, $\beta = 1e-5$ on TIDIGITS.

Network Size	#Param.	WER (GRU)	WER (DeltaGRU)	Degradation
1L-256H	0.23 M	1.83%	3.19%	+1.36%
2L-256H	0.62 M	1.13%	1.83%	+0.69%
1L-512H	0.85 M	1.04%	1.49%	+0.44%
2L-512H	1.86 M	0.89%	1.64%	+0.75%
1L-768H	2.42 M	1.27%	1.38%	+0.11%
2L-768H	5.40 M	0.77%	1.30%	+0.53%

Fig. 2 shows the design of the PE. The PE has a 16-bit multiplier **MUL** and two adders, 32-bit **ADD0** and 16-bit **ADD1**. Multiplexers are placed before operands of **MUL** to reuse it in both matrix-vector multiplications between delta vectors Δ and weights W , and any element-wise multiplication. The nonlinear unit (**NLU**) uses look-up tables (**LUT**) to compute quantized **sigmoid** and **tanh** functions. The multiplexer below **ADD0** selects between **BRAM** data and '0' for accumulation and necessary **BRAM** initialization respectively. Signal s from **CTRL** is used to control multiplexers and select target nonlinear function of **NLU**. **ADD1** is responsible for element-wise additions and sends the output activation h to output buffer **OBUF**.

B. Implementation on MiniZed

Fig. 3 shows the implementation of EdgeDRNN on the Zynq-7007S system-on-chip (SoC) on the \$89 MiniZed development board [11]. EdgeDRNN is implemented in the programmable logic (**PL**). I/O is managed by an AXI Direct Memory Access (**DMA**) IP. The AXI Datamover fetches weights from DDR3 memory on the Processing System (**PS**) side through a 64-bit (BW_{DRAM}) AXI-Full High Performance (**HP**) slave port. The AXI-Lite General Purpose (**GP**) master port is used for the single-core ARM Cortex-A9 CPU to

TABLE III: Latency and throughput of EdgeDRNN on DeltaGRU networks trained with $\Theta = 0x40$, $\beta = 1e-5$.

Network Sizes	Op (Timestep)	Latency (μ s)		Effective Throughput (GOp/s)		MAC Efficiency	Sparsity $\Gamma_{\Delta x}$	Sparsity $\Gamma_{\Delta h}$
		Mean (min, max)	Est.	Mean (min, max)	Est.			
1L-256H	0.45 M	46.4 (16.5, 142.4)	43.3	9.8 (3.2, 27.5)	10.5	490%	25.6%	90.0%
2L-256H	1.24 M	91.0 (29.3, 259.1)	91.6	13.6 (4.8, 42.4)	13.6	682%	78.9%	89.1%
1L-512H	1.70 M	130.7 (40.8, 331.2)	129.8	13.0 (5.1, 41.6)	13.1	649%	25.6%	89.5%
2L-512H	3.72 M	252.8 (57.2, 657.0)	262.9	19.2 (7.4, 84.6)	18.4	958%	85.5%	91.2%
1L-768H	4.84 M	224.3 (64.3, 616.8)	224.8	16.6 (6.0, 57.9)	16.6	830%	25.6%	91.3%
2L-768H	10.80 M	535.7 (96.6, 1344.7)	541.6	20.2 (8.0, 111.8)	19.9	1008%	87.0%	91.6%




Fig. 5: EdgeDRNN power breakdown on MiniZed.

control the DMA and write configurations, including network size, delta threshold and offset address of weights, to the EdgeDRNN. The PL is globally driven by a 125 MHz clock from the PS.

Table I shows the resource utilization of the PL. BRAMs are used to synthesize previous state memory in DU, accumulation memory in PE and FIFOs. 8 DSPs are used for the MAC units in 8 PEs while the remaining DSP in CTRL produces weight column addresses. The most consumed resources are LUTs (72%).

IV. EXPERIMENTAL RESULTS

We trained 6 different sizes of GRU and corresponding DeltaGRU networks to compare their word error rate (WER) on the TIDIGITS audio digit dataset, evaluated using the greedy decoder. Inputs of all networks are 40-dimensional log filter bank features extracted from audio sampled at 20kHz and framed with 25ms frame size and 10ms frame stride. Networks are trained for 50 epochs using the Connectionist Temporal Classification (CTC) loss function [12] and L1 regularizer with factor $\beta=1e-5$ [10]. The Adam optimizer was used to update network parameters with learning rate of $3e-4$ and batch size of 32. EdgeDRNN was configured to use INT16 activations and INT8 weights and these networks were trained in PyTorch 1.2.0 with a quantization method similar to [13]. We used DeltaGRU Θ from 0 to 0.5 (0x80). Training was coded in Python with PyTorch 1.2.0 and ran on an NVIDIA GTX 1080 GPU with CUDA 10 and cuDNN 7.6. Latency and throughput of EdgeDRNN were evaluated on DeltaGRU networks of different sizes using the first 10,000 timesteps of the test set. The latency is the elapsed time from when input data is fetched for RNN computation to when RNN output data is available in DRAM.

A. Accuracy and Throughput

Figure 4 shows the EdgeDRNN throughput and WER versus the Θ used in training and testing of a 2L-768H-DeltaGRU network. With 8 PEs at 125 MHz, EdgeDRNN has a theoretical peak throughput of 2 GOp/s. At $\Theta = 0$, there is still a speedup of about 2X from natural sparsity of the delta vectors. Higher Θ leads to better effective throughput, but with gradual slight WER degradation. The optimal point is at $\Theta = 0x40$ (0.25), just before a dramatic increase of WER, where EdgeDRNN achieves an effective throughput around 20.2 GOp/s with 1.3% WER. We use the same $\Theta = 0x40$ to train all other DeltaGRU networks and their accuracy is compared with GRU networks of the same size in Table II. The smallest network 1L-256H-DeltaGRU has a 1.36% WER increase. The largest network 2L-768H-DeltaGRU achieves a 0.53% higher WER but 4X more effective throughput. Setting $\Theta \leq 0x08$ shows that INT16/INT8 arithmetic achieves the same accuracy as FP32 (Table IV), but here the effective throughput is reduced to 6.5 versus 20.2 GOp/s/W.

B. Theoretical & Measured Performance

The theoretical estimated mean effective throughput ν of EdgeDRNN running a DeltaGRU layer is given as:

$$\nu = \frac{Op}{\tau_M + \tau_A} \quad (2)$$

$$\approx \frac{2(3MN + 3M^2(L-1) + 3M^2L)}{\frac{(3MN + 3M^2(L-1))(1-\Gamma_{\Delta x}) + 3M^2L(1-\Gamma_{\Delta h})}{Kf} + \frac{3M}{Kf}} \quad (3)$$

where Op is the number of operations in a DeltaGRU layer per timestep, τ_M the latency of $M \times V$ and τ_A the latency of remaining operations to produce the activation. $\Gamma_{\Delta x}$ and $\Gamma_{\Delta h}$ are the mean sparsity of input and hidden units respectively, L the number of hidden layers and f the clock frequency.

Table III compares benchmark results of different sizes of DeltaGRU networks on EdgeDRNN. Estimated results by Eq. 3 are within 7.1% relative error to measured results, so Eq. 3 is useful to estimate EdgeDRNN performance. On average, EdgeDRNN can run all tested networks with less than 0.54ms latency, which corresponds to 20.2 GOp/s effective throughput for the 2L-768H-DeltaGRU.

C. Power Measurement

Fig. 5 shows the power breakdown of the MiniZed system. The total power is measured by a USB power meter; the PS, PL and static power is estimated by the Xilinx Power Analyzer. The whole system active burns at most




Fig. 6: (Top) Audio spectrogram filter bank features with annotated labels and (bottom) measured hardware latency per frame of a sample (2589604A.WAV) from the TIDIGITS test set benchmarked on different hardware platforms.

TABLE IV: Comparison of EdgeDRNN with previous work and commercial products (the 5 W Google Edge TPU does not support RNNs).

Platform	FPGA			ASIC			GPU		
	This Work	DeepStore	ESE	NCS2	Jetson Nano	Jetson TX2	GTX 1080		
Chip	XC7Z007S	XC7Z045	XCKU060	Myriad X	Tegra X1	Tegra X2	GP104		
Dev. Kit Cost	\$89	\$2,495	\$3,295	\$69	\$99	\$411	\$500+PC		
Bit Precision (A/W)	INT 16/8	INT 16/16	INT 16/12	FP 16/16	FP 32/32	FP 32/32	FP 32/32		
Test Network	DeltaGRU	LSTM	Google LSTM	LSTM	GRU	GRU	GRU		
Network Size	2L-768H	2L-128H	1L-1024H	2L-664H	2L-768H	2L-768H	2L-768H		
#Parameters	5.40 M	0.26 M	3.25 M	5.40 M	5.40 M	5.40 M	5.40 M		
WER on TIDIGITS	$\Theta = 0x00$	$\Theta = 0x08$	$\Theta = 0x40$	-	-	1.07%	0.77%	0.77%	0.77%
	0.69%	0.75%	1.30%	-	-	3,588	5,327	3,240	715
Latency (μ s)	2633	1673	536	-	-	3,588	5,327	3,240	715
Batch-1 Throughput (GOp/s)	4.10	6.46	20.16	1.04	79.20	3.01	2.03	3.33	15.10
On-Chip Power (W)	1.48			2.30	-	-	-	-	-
Batch-1 On-Chip Power Efficiency (GOp/s/W)	3.20	4.36	13.62	0.45	-	-	-	-	-
Wall Plug Power (W)	2.42			-	41.00+PC	1.74	7.56	11.70	92.43+PC
Batch-1 System Power Efficiency (GOp/s/W)	1.70	2.68	8.35	-	1.93	1.73	0.27	0.28	0.16

2.416 W. The EdgeDRNN logic burns only 87 mW. Thus the wall plug and incremental power efficiency are 8.4 GOp/s/W and 231.7 GOp/s/W respectively. Varying modes of operation allows inferring EdgeDRNN DRAM memory power of 358 mW, resulting in EdgeDRNN+DRAM power efficiency of 38.3 GOp/s/W. We used the wall plug power efficiency for the following comparisons.

V. CONCLUSION

Table IV compares EdgeDRNN with other platforms. The same task (first 10,000 timesteps of the test set) was benchmarked on EdgeDRNN, ASIC and GPUs. The Intel Compute Stick 2 (NCS2) does not support GRU and was benchmarked with an LSTM network with similar parameter count and trained on the same dataset and hyperparameters. For benchmark of GPUs, we used the cuDNN implementation of GRU that achieved 715 μ s latency on NVIDIA GTX 1080, which is 2.4X quicker than the DeltaGRU using the NVIDIA cuSPARSE library. We also compare this work to reported specifications of DeepStore [14], which has similar power consumption as EdgeDRNN, and ESE [6], which is a sparse matrix-vector multiplication accelerator for LSTM.

The power efficiency results show that EdgeDRNN achieves over 4.8X higher system power efficiency compared to com-

mercial ASIC and GPU products, 30X higher on-chip power efficiency compared to [14] and 4.3X higher system power efficiency than ESE.

Fig. 6 compares the latencies on a test set sample. EdgeDRNN is as quick as 1080 GPU and 6X quicker than the other platforms. EdgeDRNN latency is lower during the silent or quieter periods (e.g. between 120s and 140s).

The delta threshold Θ allows instantaneous tradeoff of accuracy versus latency. Using sparsity in delta vectors allows the arithmetic units on this task to effectively compute ten times more operations.

The throughput of commercial edge devices on batch-1 RNNs are a factor of more than 100X less than the claimed peak performance offered by these platforms, which range from 500 GOp/s for Jetson Nano [15] up to nearly 10TOP/s for GTX 1080 [16] It shows that an optimized RNN platform can do better in throughput and especially power efficiency.

VI. ACKNOWLEDGEMENT

This work was partially funded by the Samsung Advanced Institute of Technology and the Swiss National Science Foundation, HEAR-EAR, 200021_172553 grant.

REFERENCES

- [1] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997. [Online]. Available: <http://dx.doi.org/10.1162/neco.1997.9.8.1735>
- [2] K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder–decoder for statistical machine translation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Oct. 2014, pp. 1724–1734. [Online]. Available: <http://www.aclweb.org/anthology/D14-1179>
- [3] Amodei, Dario et al., "Deep Speech 2: End-to-end Speech Recognition in English and Mandarin," in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ser. ICML'16. JMLR.org, 2016, pp. 173–182. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3045390.3045410>
- [4] M. Ravanelli, T. Parcollet, and Y. Bengio, "The Pytorch-Kaldi speech recognition toolkit," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 6465–6469.
- [5] J. Chen and X. Ran, "Deep Learning With Edge Computing: A Review," *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1655–1674, Aug. 2019.
- [6] S. Han, J. Kang, H. Mao, Y. Hu, X. Li, Y. Li, D. Xie, H. Luo, S. Yao, Y. Wang et al., "ESE: Efficient speech recognition engine with sparse LSTM on FPGA," in *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, 2017, pp. 75–84.
- [7] S. Cao, C. Zhang, Z. Yao, W. Xiao, L. Nie, D. Zhan, Y. Liu, M. Wu, and L. Zhang, "Efficient and effective sparse LSTM on FPGA with bank-balanced sparsity," ACM, Feb. 2019, pp. 63–72. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3289602.3293898>
- [8] S. Wang, Z. Li, C. Ding, B. Yuan, Q. Qiu, Y. Wang, and Y. Liang, "C-LSTM: Enabling efficient LSTM using structured compression techniques on FPGAs," in *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, ser. FPGA '18. New York, NY, USA: ACM, 2018, pp. 11–20. [Online]. Available: <http://doi.acm.org/10.1145/3174243.3174253>
- [9] C. Gao, D. Neil, E. Ceolini, S.-C. Liu, and T. Delbruck, "DeltaRNN: A power-efficient recurrent neural network accelerator," in *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, ser. FPGA '18. New York, NY, USA: ACM, 2018, pp. 21–30. [Online]. Available: <http://doi.acm.org/10.1145/3174243.3174261>
- [10] D. Neil, J. Lee, T. Delbrück, and S. Liu, "Delta networks for optimized recurrent network computation," in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, 2017, pp. 2584–2593. [Online]. Available: <http://proceedings.mlr.press/v70/neil17a.html>
- [11] AVNET, "Minized." [Online]. Available: <http://zedboard.org/product/minized>
- [12] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd International Conference on Machine Learning*, ser. ICML '06. New York, NY, USA: ACM, 2006, pp. 369–376. [Online]. Available: <http://doi.acm.org/10.1145/1143844.1143891>
- [13] A. Mishra, E. Nurvitadhi, J. J. Cook, and D. Marr, "WRPN: Wide reduced-precision networks," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=B1ZvaeAZ>
- [14] A. X. M. Chang and E. Culurciello, "Hardware accelerators for recurrent neural networks on FPGA," in *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2017, pp. 1–4.
- [15] "Jetson Nano Developer Kit," Mar. 2019. [Online]. Available: <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>
- [16] "GeForce GTX 1080 Graphics Cards from NVIDIA GeForce." [Online]. Available: <https://www.nvidia.com/en-us/geforce/products/10series/geforce-gtx-1080/>