

Schrödinger's FP: Dynamic Adaptation of Floating-Point Containers for Deep Learning Training

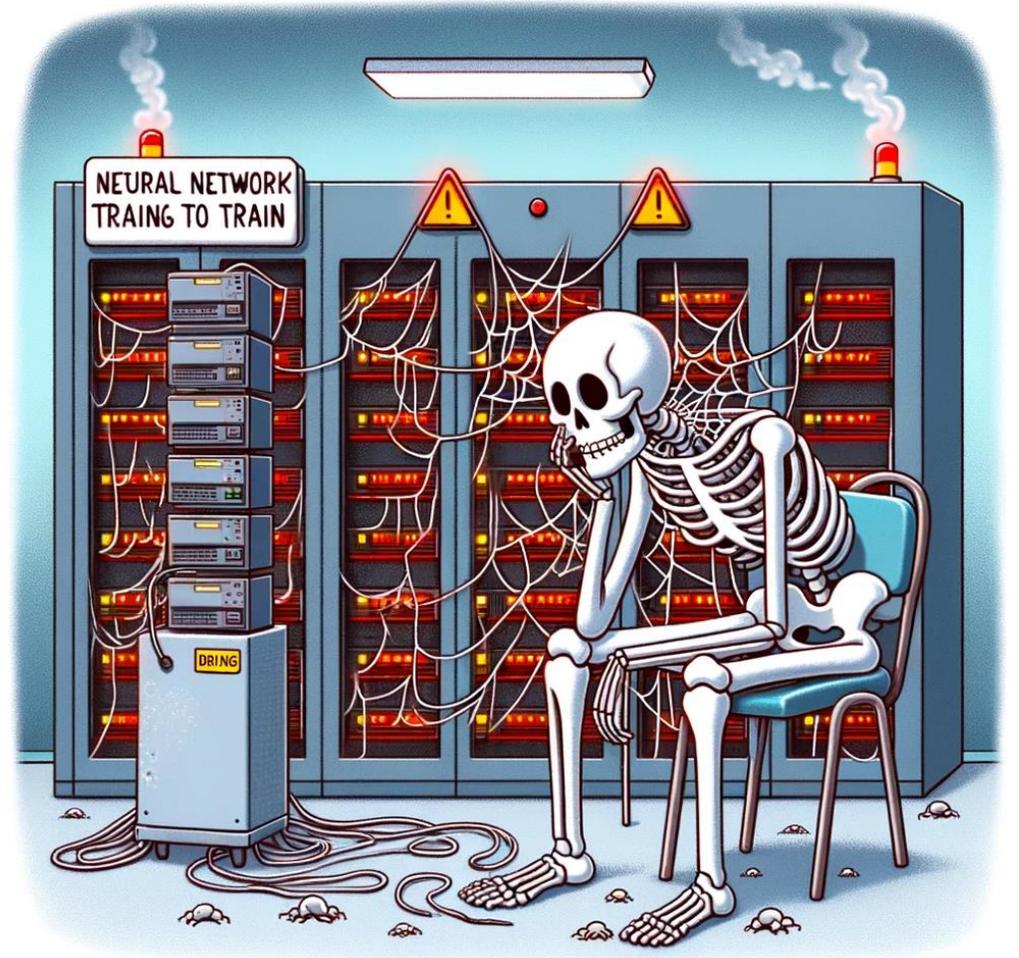
Miloš Nikolić, Enrique Torres Sanchez, Jiahui Wang, Ali Hadi Zadeh,
Mostafa Mahmoud, Ameer Abdelhadi, Kareem Ibrahim, and Andreas Moshovos



The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

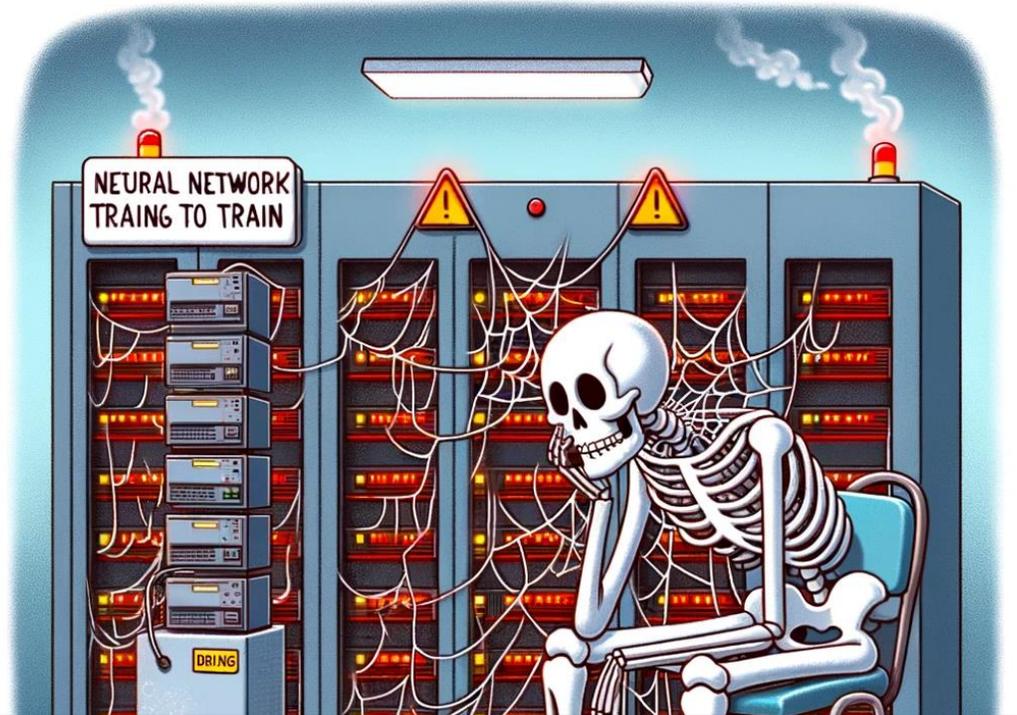
Training is Expensive!

- Neural Networks are becoming ubiquitous
- But they are expensive
 - Energy
 - Time



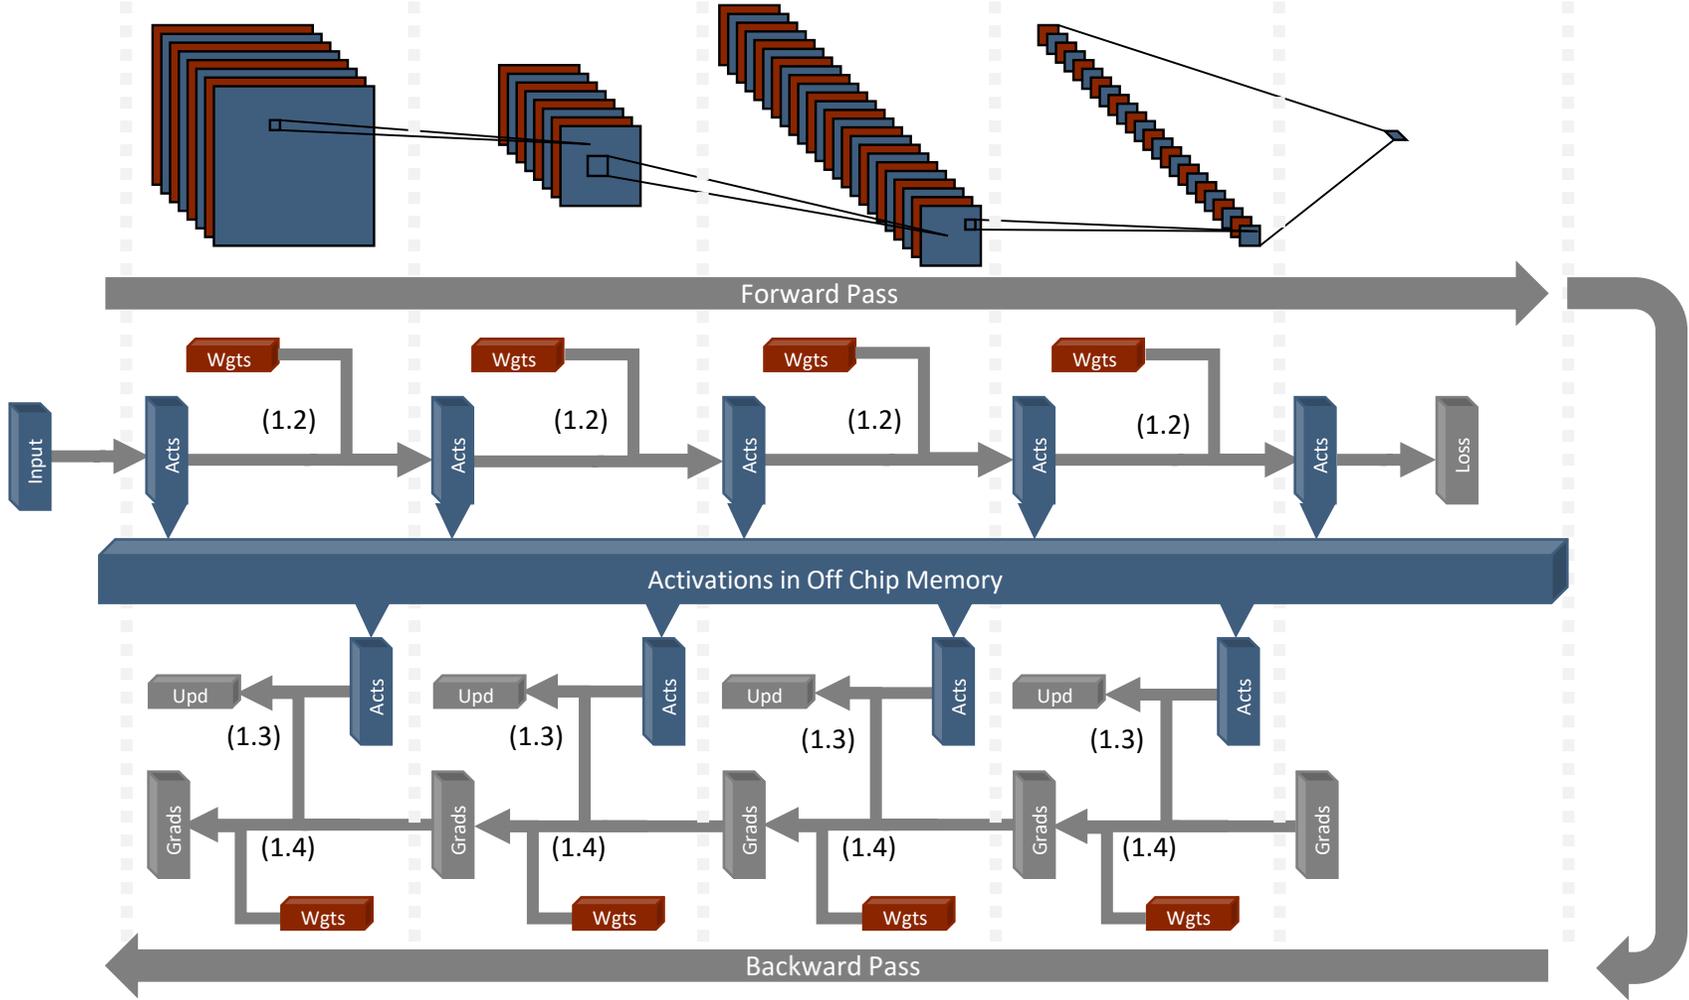
Training is Expensive!

- Neural Networks are becoming ubiquitous
- But they are expensive
 - Energy
 - Time



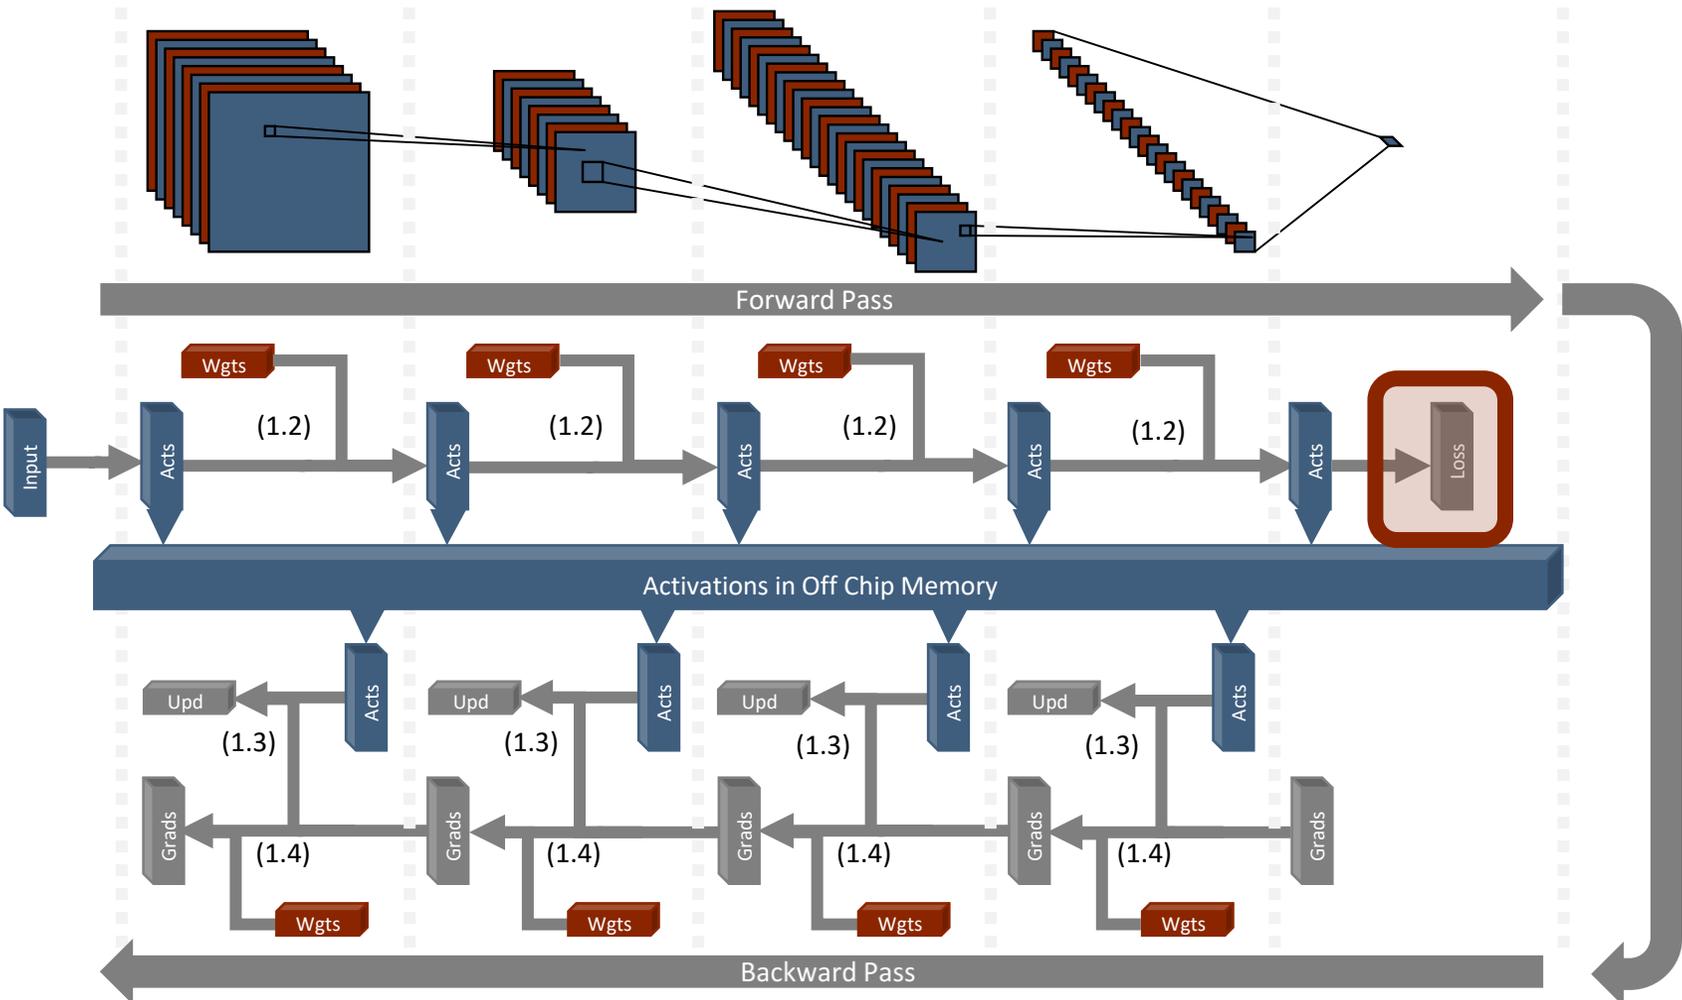
Goal: Improve energy efficiency and performance of Training

Gradient Descent – Overview



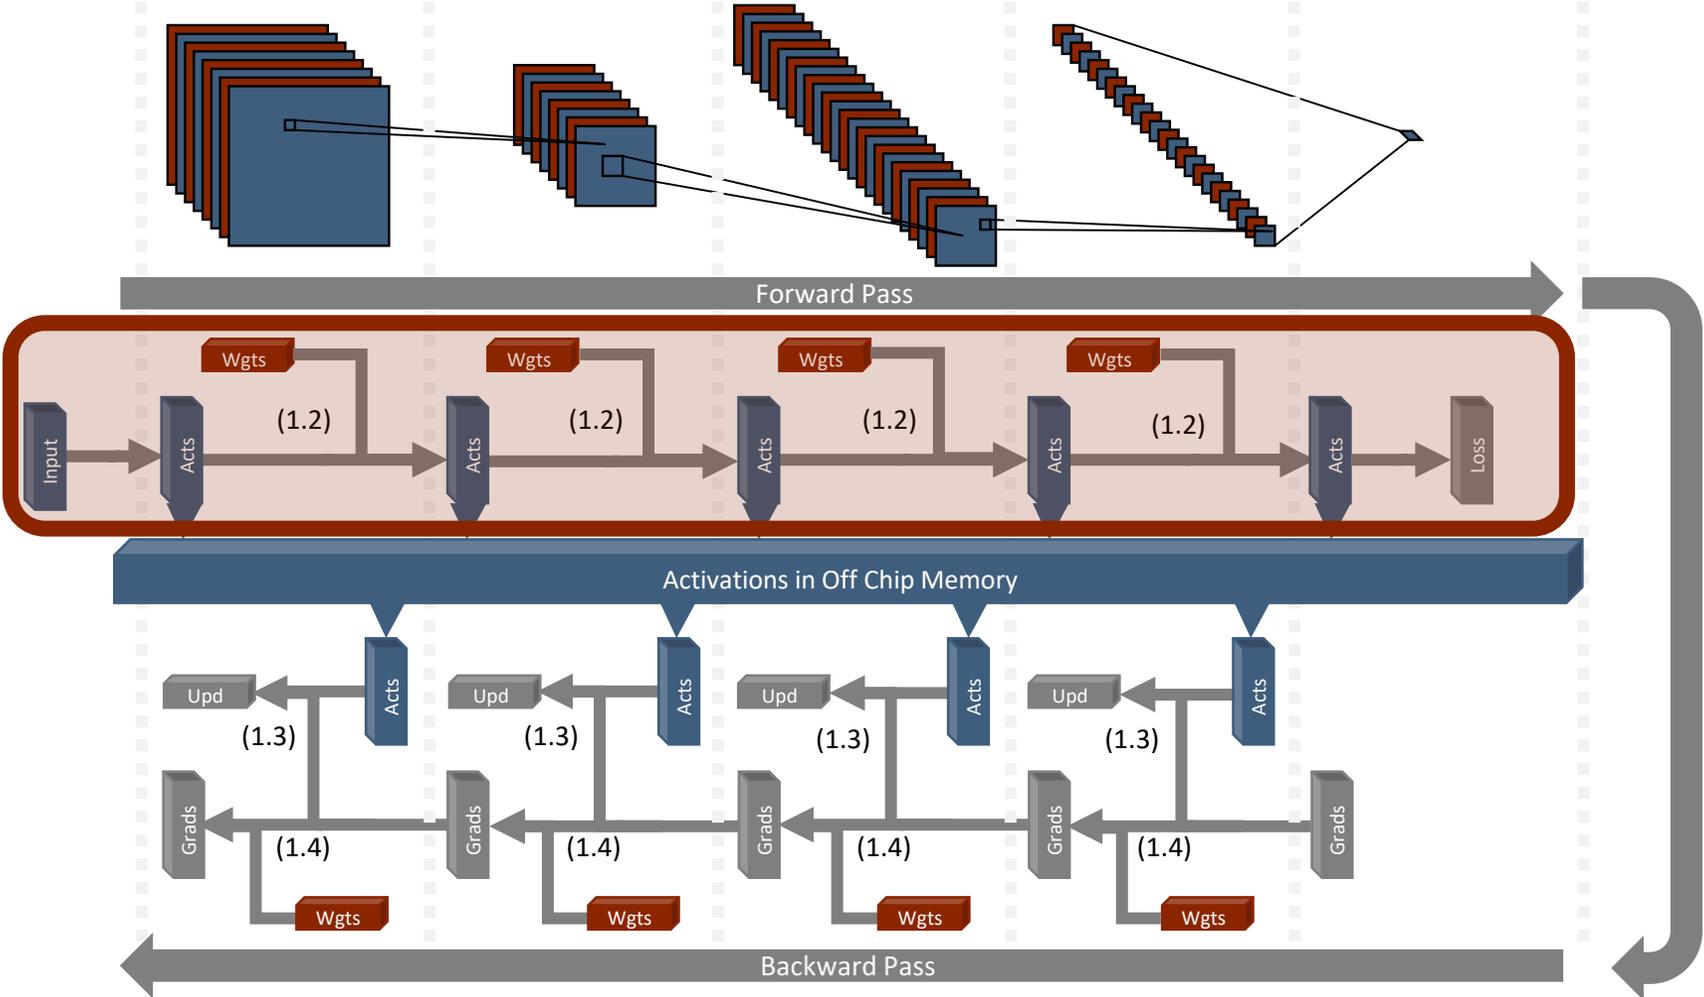
Gradient Descent – Overview

- Loss function



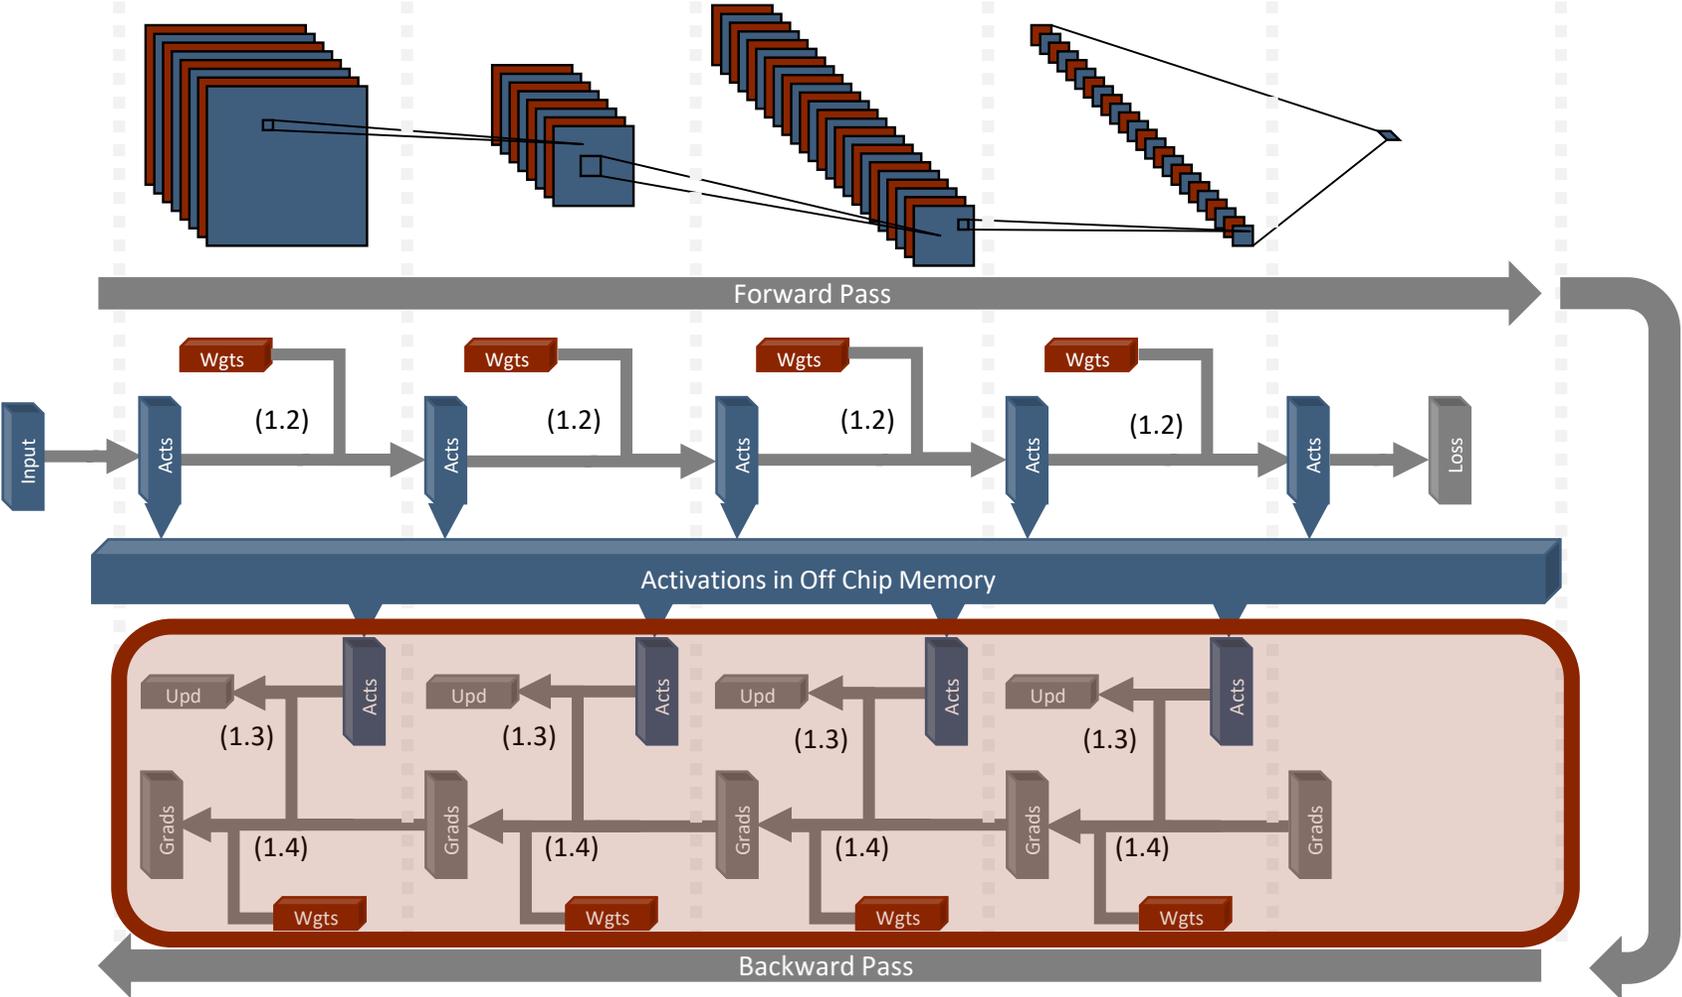
Gradient Descent – Overview

- Loss function
- Forward pass



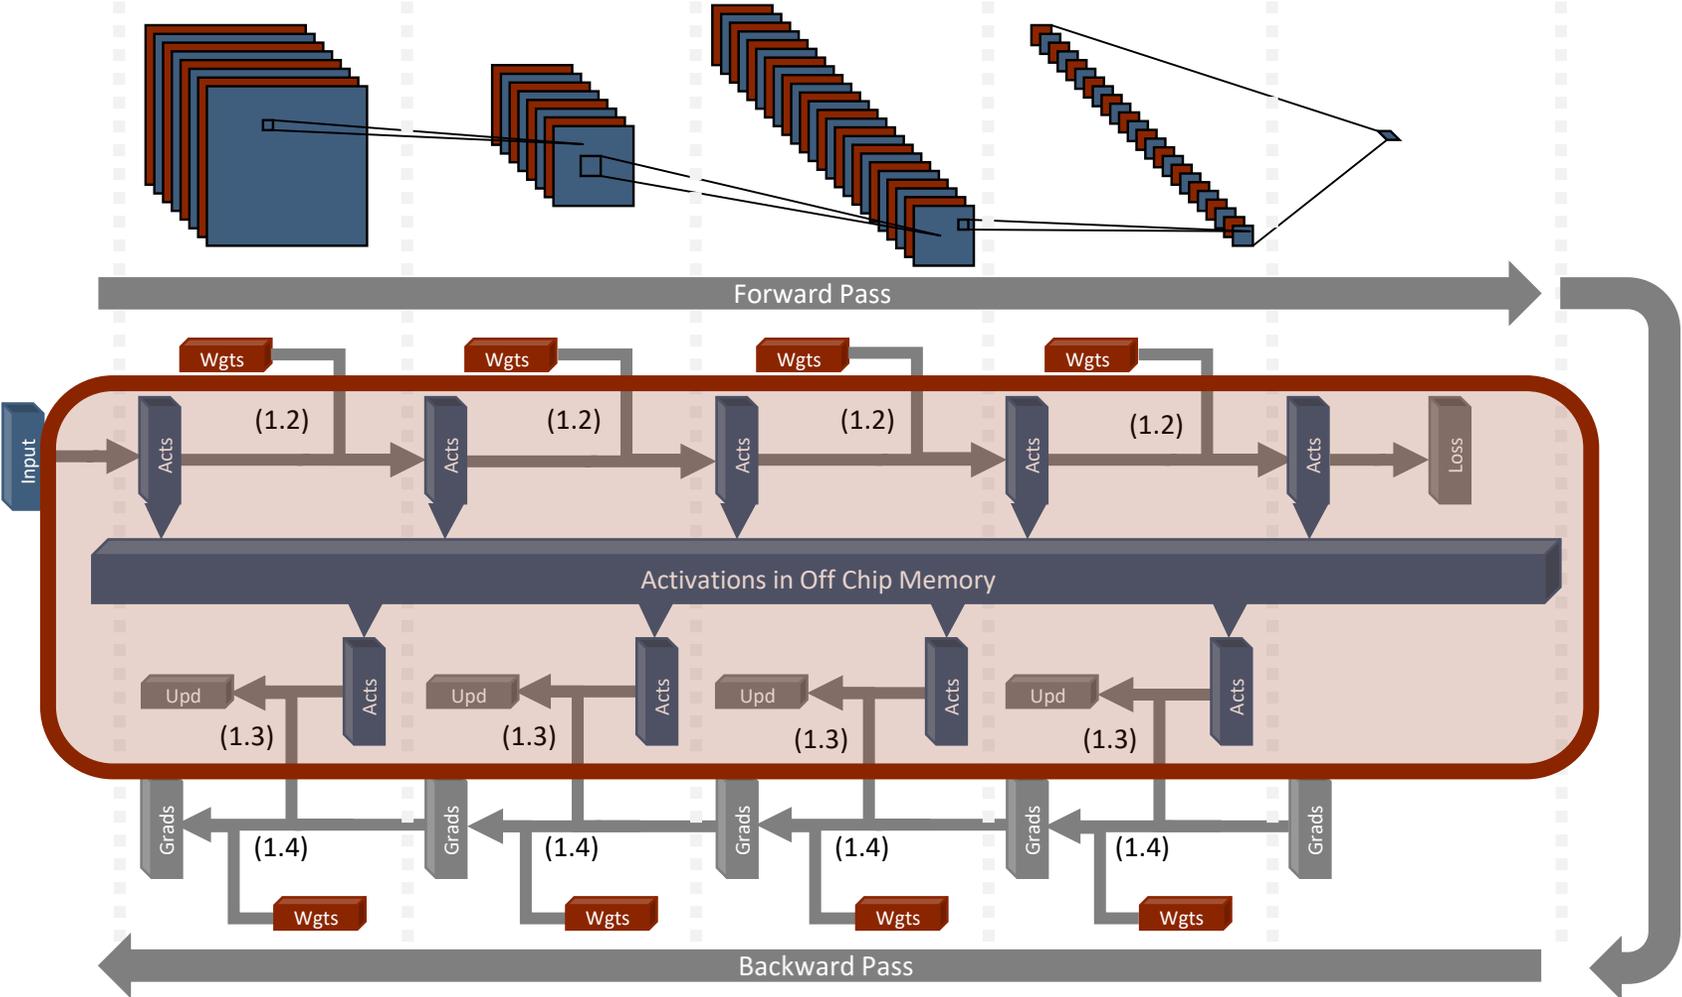
Gradient Descent – Overview

- Loss function
- Forward pass
- Backward pass
 - Update
 - Gradient



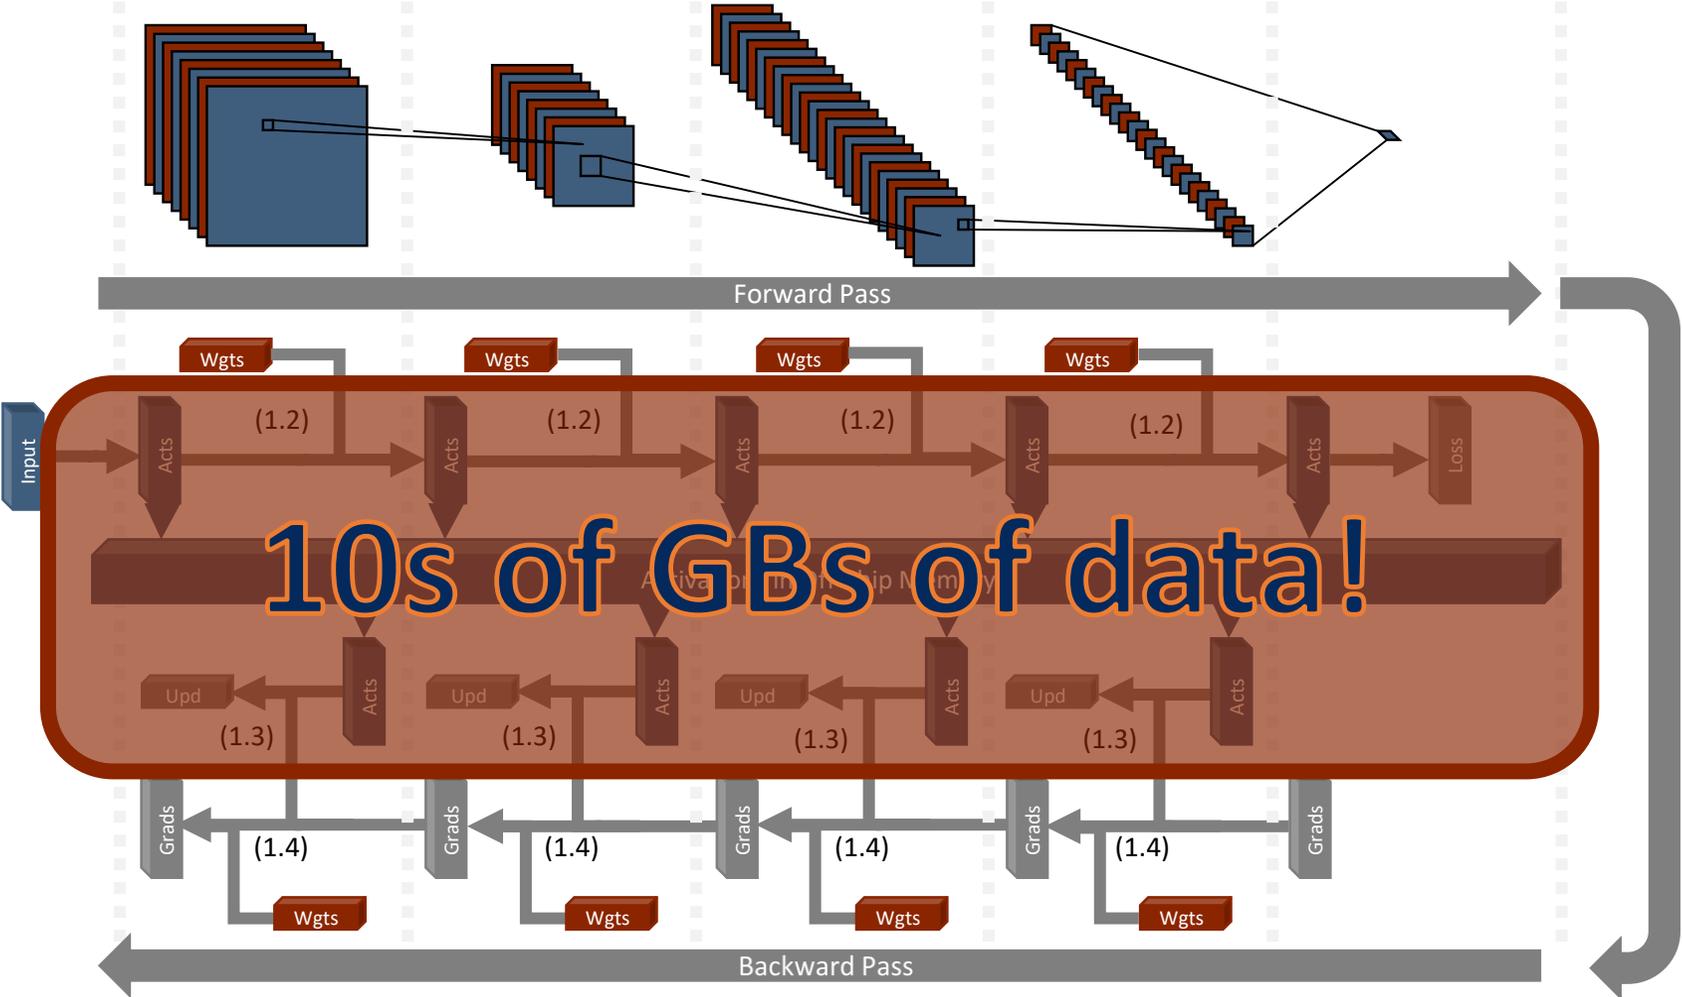
Gradient Descent – Overview

- Loss function
- Forward pass
- Backward pass
 - Update
 - Gradient



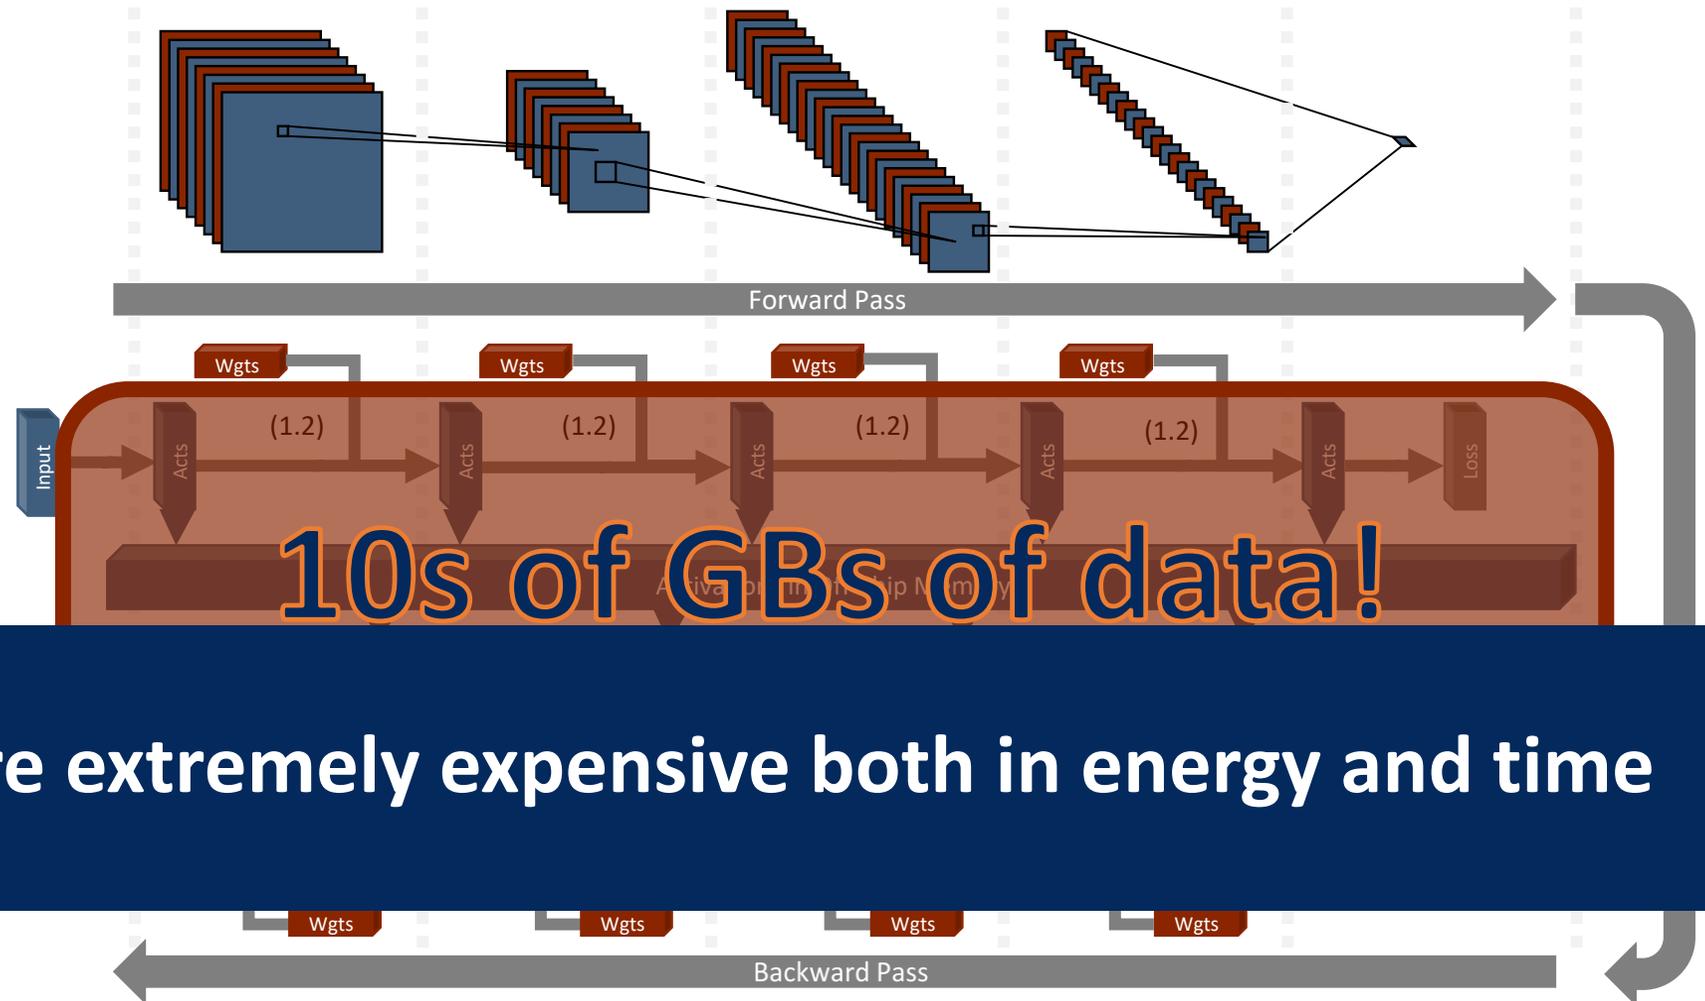
Gradient Descent – Overview

- Loss function
- Forward pass
- Backward pass
 - Update
 - Gradient



Gradient Descent – Overview

- Loss function
- Forward pass
- Backward pass
 - Update
 - Gradient



Off-chip accesses are extremely expensive both in energy and time

Narrow Datatypes can help

- Efficient Datatype reduce costs!
 - Memory
 - Compute

Narrow Datatypes can help

- Efficient Datatype reduce costs!
 - Memory
 - Compute



Narrow Datatypes can help

- Efficient Datatype reduce costs!
 - Memory
 - Compute



Narrow Datatypes can help

- Efficient Datatype reduce costs!
 - Memory
 - Compute



Narrow Datatypes can help

- Efficient Datatype reduce costs!
 - Memory
 - Compute



Which Datatype? Where? When?

- There are many great options



9.05433v2 [cs.LG] 29 Sep 2022

Which Datatype? Where? When?

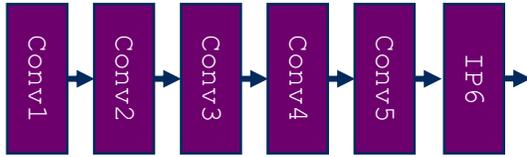
- There are many great options
- But it's a trade-off
 - Accuracy vs. Efficiency
 - All have edge cases



Which Datatype? Where? When?

- There are many great options
- But it's a trade-off
 - Accuracy vs. Efficiency
 - All have edge cases

FP32						
BF16	✓					✓
FP8		✓		✓		
FP6			✓		✓	



MIXED PRECISION TRAINING
Sharan Narang, Gregory Diamos, Erich Elsen, Blaise Research, {sharan, gdiamos}@baidu.com

Training DNNs with Hybrid Block Floating Point
Mario Drummond, EcoCloud, EPFL, mario.drummond@epfl.ch
Tao Lin, EcoCloud, EPFL, tao.lin@epfl.ch
Martin Jaggi, EcoCloud, EPFL, martin.jaggi@epfl.ch
Babak Falsafi, EcoCloud, EPFL, babak.falsafi@epfl.ch

Training Deep Neural Networks with 8-bit Floating Point Numbers
Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen and Kailash Gopalakrishnan, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA, {nwang, choij, danbrand, cchen, kailash}@us.ibm.com

A Study of BFLOAT16 for Deep Learning Training
Dhiraj Kalamkar¹, Dhruvraj Mudigere², Naveen Mellempudi¹, Dipankar Das¹, Komal Banerjee¹, Saikrishna Avacha¹, Dharna Teja Vooturi¹, Nataraj Jummalamadaka¹, Anshul Kumar¹, Jyayu Huang¹, Hector Yuen¹, Jigang Yang¹, Jongsan Park¹, Alexander Heinecke¹, Evangelos Georgantas¹, Sudarshan Srinivasan¹, Abhishek Kundu¹, Mihir Smolyansky¹, Bharat Kaul¹, and Pradeep Dubey¹

Ultra-Low Precision 4-bit Training of Deep Neural Networks
Xiao Sun, Naigang Wang, Chia-yu Chen, Jia-min Ni, Ankur Agrawal, Xiaodong Cui, Swagath Venkataramani, Kaustav El Maghraoui, Vijayalakshmi Srivisan

FAST: DNN Training Under Variable Precision Block Floating Point with Stochastic Rounding
Sui Qian Zhang¹, Bradley McDaneF¹, and H.T. Kung¹

FP8 FORMATS FOR DEEP LEARNING
Paulius Mickevičius, Dusan Stolic, Patrick Judd, John Kamuh, Stuart Oberman, Mohammad Shoeybi, {paulius, dstolic, pjudd, jkamalu, soberman, mshoeybi, stou, skyu}@nvidia.com

With Shared Microexponents, A Little Shifting Goes a Long Way
Zhaoxia (Summer) Deng, Sam Naghshineh, Jongsoo Park, Maxin Naumov

9,05433v2 [cs.LG] 29 Sep 2022

Which Datatype? Where? When?

- There are many great options
- But it's a trade-off
 - Accuracy vs. Efficiency
 - All have edge cases

FP32						
BF16	✓					✓
FP8		✓		✓		
FP6			✓			✓

CONV1	→	CONV2	→	CONV3	→	CONV4	→	CONV5	→	IP6
-------	---	-------	---	-------	---	-------	---	-------	---	-----



9,05433v2 [cs.LG] 29 Sep 2022

Which Datatype? Where? When?

- There are many great options
- But it's a trade-off
 - Accuracy vs. Efficiency
 - All have edge cases
- Guess
 - Confirm
 - Repeat

MIXED PRECISION TRAINING

Sharan Narang¹, Gregory Diamos, Erich Elsen¹,
Blaise Research
{sharan, gdiamos}@baidu.com

Paulius Mikelicivicius¹, Jonah Alben, David Garcia, Boris Ginsburg, Michael Houston,
Oleksii Kuchalev, Ganeesh Venkatesh, Hao Wu¹,
NVIDIA
{paulium, alben, dagarcia, bginsburg, mhouston,
okuchalev, gavenkatesh, skyw}@nvidia.com

ABSTRACT
Typically improves accuracy
for training the model
using half-

Training Deep Neural Networks with 8-bit Floating Point Numbers

Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen and Kailash Gopalakrishnan
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598, USA
{nwang, choij, danbrand, cchen, kailash}@us.ibm.com

ABSTRACT
The state-of-the-art hardware platforms for training Deep Neural Networks (DNNs) are moving from traditional single precision (32-bit) floating point to 16-bit and 8-bit precision – in large part due to the high energy cost of floating point arithmetic associated with using reduced-precision data types. Unlike inference, training with reduced-precision data types is challenging due to the need for backpropagation through the network.

Ultra-Low Precision 4-bit Training of Deep Neural Networks

Xiao Sun, Naigang Wang, Chia-yu Chen, Jia-min Ni
Ankur Agrawal, Xiaodong Cui, Swagath Venkataramani
Kautsar El Maghraoui, Vijayalakshmi Srivivasan
Kailash Gopalakrishnan
IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA
{sun, nwang, cchen, jkamin.ni, ankurag, cui, swagath.venkataramani, kelmaghr, viji,kailash}@us.ibm.com

ABSTRACT
novel techniques and numerical representations to enable this advance. To enable this advance, we propose a 4-bit floating point format that addresses the challenges of training with reduced-precision data types.

FP8 FORMATS FOR DEEP LEARNING

Paulius Mikelicivicius, Dusan Stolic, Patrick Judd, John Kamalu, Stuart Oberman, Mohammad Shoeybi,
Michael Sin, Hao Wu
NVIDIA
{paulium, dstolic, pjudd, jkamalu, soberman, mshoeybi, msin, skyw}@nvidia.com

Neil Burgess, Sangwon Ha, Richard Grisenbawite,
{neil.burgess, sangwon.ha, richard.grisenbawite}@arm.com

Naveen Mellempudi, Marius Cornes, Alexander Heinecke, Pradeep Dubej
{naveen.k.mellempudi, marius.cornes, alexander.heinecke, pradeep.dubej}@intel.com

ABSTRACT
FP8 is a natural progression for accelerating deep learning training inference. This paper introduces Block Data Representations (BDRs), a framework for exploring and evaluating a wide spectrum of narrow-precision formats for deep learning. It enables comparison of popular quantization standards, such as FP8, R, new formats based on FP8, and FP8 with dynamic ranges. The study covers tasks such as image classification, object detection, and natural language processing.

Training DNNs with Hybrid Block Floating Point

Mario Drummond, EcoCloud
EPLL
mario.drummond@epfl.ch

Tao Lin, EcoCloud
EPLL
tao.lin@epfl.ch

Martin Jaggi, EcoCloud
EPLL
martin.jaggi@epfl.ch

Babak Falsafi, EcoCloud
EPLL
babak.falsafi@epfl.ch

ABSTRACT
The wide adoption of DNNs has given birth to a new generation of hardware accelerators. These accelerators typically employ arithmetic to maximize performance. This paper introduces a novel hybrid block floating point arithmetic. It is designed to support a wide range of dynamic ranges. We show that this arithmetic can be used to train DNNs with high accuracy and low energy consumption.

A Study of BFLOAT16 for Deep Learning Training

Dhiraj Kalamkar¹, Dhruvraj Mudigere², Naveen Mellempudi¹, Dipankar Das¹,
Komal Banerjee¹, Sankarsh Avacha¹, Dhruva Teja Vooturi¹, Nataraj Jummalamadaka¹,
Ajayya Huang¹, Hector Yuen¹, Jagan Yang¹, Jongsung Park¹, Alexander Heinecke¹,
Evangelos Georgakopoulos¹, Sudarshan Srivivasan¹, Aishwarya Kundu¹,
Misha Smelyanskiy¹, Bharat Lall¹, and Pradeep Dubej¹

¹Paradee Computing Lab, Intel Labs
²Facebook, 1 Hacker Way, Menlo Park, CA

ABSTRACT
This paper presents the first comprehensive empirical study of the efficacy of the Brain Floating Point (BFloat16) format for training deep neural networks. We compare BFloat16 against FP16 and FP32 across a wide range of tasks and hardware configurations. Our results show that BFloat16 can achieve comparable accuracy to FP16 and FP32 while consuming significantly less energy.

FAST: DNN Training Under Variable Precision Block Floating Point with Stochastic Rounding

Sui Qian Zhang¹, Bradley McDaneF¹, and H.T. Kung¹
¹Harvard University
²Franklin and Marshall College

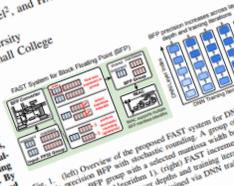


Fig. 1. (left) Overview of the proposed FAST system for DNN training. (right) Comparison of the proposed FAST system with other BFP systems. The FAST system achieves higher accuracy with lower energy consumption compared to other BFP systems.

With Shared Microexponents, A Little Shifting Goes a Long Way

Yuhan¹, Ritchie Zhao, Venmugil Elango, Rasoul Shafiqpour, Mathew Hall, Maral Mesmakhoshroshahi,
More, Levi Melnick, Maximilian Golub, Girish Varatkar, Lei Shao, Gaurav Kulkarni, Dmitry Melts,
Jasmine Klar, Renee L'Huereux, Matt Perry, Doug Burger, Eric Chung
Zhaoxia (Summer) Deng, Sam Naghshineh, Jongsoo Park, Maxin Naumov
Meta

ABSTRACT
This paper introduces Block Data Representations (BDRs), a framework for exploring and evaluating a wide spectrum of narrow-precision formats for deep learning. It enables comparison of popular quantization standards, such as FP8, R, new formats based on FP8, and FP8 with dynamic ranges. The study covers tasks such as image classification, object detection, and natural language processing.

such as Google's [2], are increasingly being adopted for DNN training. However, these formats do not fit the empirical distribution of the data. With low-precision BFP and integer formats, the accuracy of the model is significantly lower. This paper introduces a novel BFP format, FAST, which uses shared microexponents and stochastic rounding. This format achieves higher accuracy than other BFP formats while maintaining low energy consumption. We validate our FAST system on a wide range of tasks and hardware configurations. Our results show that FAST can achieve comparable accuracy to FP16 and FP32 while consuming significantly less energy.

9.05433v2 [cs.LG] 29 Sep 2022

Which Datatype? Where? When?

- There are many great options
- But it's a trade-off
 - Accuracy vs. Efficiency
 - All have edge cases
- Guess
 - Confirm
 - Repeat
- Lost Potential

MIXED PRECISION TRAINING

Sharan Narang¹, Gregory Diamos, Erich Elsen¹,
Blaise Research
{sharan, gdiamos}@baidu.com

Paulius Mickevičius¹, Jonah Alben, David Garcia, Boris Ginsburg, Michael Houston,
Oleksii Kuchalev, Gansuk Venkatesh, Hao Wu
NVIDIA
{paulius, alben, dagarcia, bginsburg, mhouston,
okuchalev, gavenkatesh, skyw}@nvidia.com

Abstract
Typically improves accuracy for training the model using half-precision FP16.

Training Deep Neural Networks with 8-bit Floating Point Numbers

Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen and Kailash Gopalakrishnan
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598, USA
{nwang, choi, danbrand, cchen, kailash}@us.ibm.com

Abstract
The state-of-the-art hardware platforms for training Deep Neural Networks (DNNs) are moving from traditional single precision (32-bit) computers to GPUs with 16-bit precision – in large part due to the high energy efficiency of GPUs. Unlike inference associated with using reduced-precision, training with reduced-precision is challenging due to the need for backpropagation through the network.

Ultra-Low Precision 4-bit Training of Deep Neural Networks

Xiao Sun, Naigang Wang, Chia-yu Chen, Jia-min Ni
Ankur Agrawal, Xiaodong Cui, Swagath Venkataramani
Kasour El Maghraoui, Vijayalakshmi Srivivasan
Kailash Gopalakrishnan
IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA
{sun, nwang, cchen, jiamin.ni, ankurag, cui, swagath.venkataramani, kelmaghr, viji,kailash}@us.ibm.com

Abstract
Novel techniques and numerical representations for training systems with ultra-low precision (4-bit) floating point numbers. To enable this advance, we propose a 4-bit floating point format (FP4) that addresses the challenges of training with reduced-precision.

FP8 FORMATS FOR DEEP LEARNING

Paulius Mickevičius, Dusan Stolic, Patrick Judd, John Kamalu, Stuart Oberman, Mohammad Shoeybi,
Michael Sin, Hao Wu
NVIDIA
{paulius, dstolic, pjudd, jkamalu, soberman, mshoeybi, msu, skyw}@nvidia.com

Neil Burgess, Sangwon Ha, Richard Grisenhwaite
{neil.burgess, sangwon.ha, richard.grisenhwaite}@arm.com

Naveen Mellempudi, Marius Cornea, Alexander Heinecke, Pradeep Dubej
{naveen.k.mellempudi, marius.cornea, alexander.heinecke, pradeep.dubej}@intel.com

Abstract
FP8 is a natural progression for accelerating deep learning training inference. This paper introduces Block Data Representations (BDRs), a framework for exploring and evaluating a wide spectrum of narrow-precision formats for deep learning. It enables comparison of popular quantization standards, such as FP8, FP4, and FP3, with new formats based on a common set of constraints. The study covers tasks such as image classification, object detection, and natural language processing.

Training DNNs with Hybrid Block Floating Point

Mario Drummond
EcoCloud
EPTL
mario.drummond@epfl.ch

Tao Lin
EcoCloud
EPTL
tao.lin@epfl.ch

Martin Jaggi
EcoCloud
EPTL
martin.jaggi@epfl.ch

Babak Falsafi
EcoCloud
EPTL
babak.falsafi@epfl.ch

Abstract
The wide adoption of DNNs has given birth to a new generation of hardware accelerators. These accelerators typically employ arithmetic to maximize performance. However, the performance of these accelerators is limited by the precision of the floating-point arithmetic. We propose a hybrid block floating point (BFP) format that allows for a wide range of precision values, enabling a more efficient use of hardware resources.

A Study of BFLOAT16 for Deep Learning Training

Dhiraj Kalamkar¹, Dhruvraj Mudigere², Naveen Mellempudi¹, Dipankar Das¹,
Komal Banerjee¹, Sankarsh Avacha¹, Dhruva Teja Vooturi¹, Nataraj Jummalamadaka¹,
Ajaykumar Bhat¹, Hector Yuen¹, Jagan Yang¹, Jusong Park¹, Alexander Heinecke¹,
Jiayu Huang¹, Hector Yuen¹, Sudarshan Srivivasan¹, Aishwarya Kundu¹,
Evangelos Georgakopoulos¹, Bharat Kaul¹, and Pradeep Dubej¹

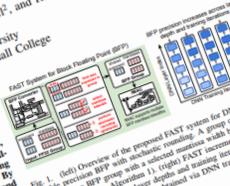
¹Parallel Computing Lab, Intel Labs
²Facebook, 1 Hacker Way, Menlo Park, CA

Abstract
This paper presents the first comprehensive empirical study of the efficacy of the Brain Floating Point (BFP) format for training Deep Neural Networks (DNNs). The BFP format is a variable-precision floating-point format that allows for a wide range of precision values, enabling a more efficient use of hardware resources.

FAST: DNN Training Under Variable Precision Block Floating Point with Stochastic Rounding

Sui Qian Zhang¹, Bradley McDaneF¹, and H.T. Kung¹
¹Harvard University
²Franklin and Marshall College

Abstract—Block Floating Point (BFP) can efficiently support adaptation for Deep Neural Network (DNN) training by providing a wide dynamic range via a shared exponent across a group of values. In this paper, we propose a Fast Floating Point (FAST) system for variable precision BFP. FAST supports stochastic rounding (SR) and variable precision BFP. FAST improves training and inference with variable precision BFP. Our FAST system is implemented on a 2.6-μm hardware accelerator. Our FAST system is implemented on a 2.6-μm hardware accelerator. Our FAST system is implemented on a 2.6-μm hardware accelerator.



With Shared Microexponents, A Little Shifting Goes a Long Way

Yuhan¹, Ritchie Zhao, Venmugil Elango, Rasoul Shafiqpour, Mathew Hall, Maral Mesmakhoshrohi,
More, Levi Melnick, Maximilian Golub, Girish Varatkar, Lei Shao, Gaurav Kulkarni, Dmitry Melts,
Jasmine Klar, Renee L'Huereux, Matt Perry, Doug Burger, Eric Chung
Zhaoxia (Summer) Deng, Sam Naghshineh, Jongsoo Park, Maxin Naumov
Meta

Abstract—This paper introduces Block Data Representations (BDRs), a framework for exploring and evaluating a wide spectrum of narrow-precision formats for deep learning. It enables comparison of popular quantization standards, such as FP8, FP4, and FP3, with new formats based on a common set of constraints. The study covers tasks such as image classification, object detection, and natural language processing.

9.05433v2 [cs.LG] 29 Sep 2022

Which Datatype? Where? When?

- There are many great options
- But it's a trade-off
 - Accuracy vs. Efficiency
 - All have edge cases
- Guess
 - Confirm
 - Repeat

Automate the datatype selection!



MIXED PRECISION TRAINING

Sharan Narang¹, Gregory Diamos, Erich Elsen¹,
Blaise Research
{sharan, gdiamos}@baidu.com

Paulius Micikevicius¹, Jonah Alben, David Garcia, Boris Ginsburg, Michael Houston,
Oleksii Kuchalev, Ganeesh Venkatesh, Hao Wu
NVIDIA
{paulius, alben, dagarcia, bginsburg, mhouston,
okuchalev, gavenkatesh, skyw}@nvidia.com

ABSTRACT
Typically improves accuracy
for training the model
using half-

Training Deep Neural Networks with 8-bit Floating Point Numbers

Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen and Kailash Gopalakrishnan
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598, USA
{nwang, choi, danbrand, cchen, kailash}@us.ibm.com

ABSTRACT
The state-of-the-art hardware platforms for training Deep Neural Networks (DNNs) are moving from traditional single precision (32-bit) floating point arithmetic to low-bit storage associated with using reduced-precision arithmetic for inference. Unlike inference, training with reduced-precision arithmetic is challenging due to the need for backpropagation through the network.

Ultra-Low Precision 4-bit Training of Deep Neural Networks

Xiao Sun, Naigang Wang, Chia-yu Chen, Jia-min Ni
Akur Agrawal, Xiaodong Cui, Swagath Venkataramani, Vijayalakshmi Srivivasan

Training DNNs with Hybrid Block Floating Point

Mario Drummond¹,
EcoCloud
EPLL
mario.drummond@epfl.ch

Tao Lin¹,
EcoCloud
EPLL
tao.lin@epfl.ch

Martin Jaggi¹,
EcoCloud
EPLL
martin.jaggi@epfl.ch

Babak Falsafi¹,
EcoCloud
EPLL
babak.falsafi@epfl.ch

ABSTRACT
The wide adoption of DNNs has given birth to a new class of accelerators. These accelerators typically employ arithmetic to maximize performance. However, this point's narrow range of values is a challenge for DNN training.

A Study of BFLOAT16 for Deep Learning Training

Dhiraaj Kalamkar¹, Dheerav Mudigere², Naveen Mellempudi¹, Dipankar Das¹,
Kunal Banerjee¹, Sankarsh Avacha¹, Dhruva Teja Vooturi¹, Nataraj Jammalamadaka¹,
Jaayuu Huang¹, Hector Yuen¹, Jigan Yang¹, Jongsoo Park¹, Alexander Heinecke¹,
Evangolos Georgantas¹, Sudarshan Srivivasan¹, Abhishek Kundu¹,
Misha Smelyanskiy¹, Bharat Kaul¹, and Pradeep Dubey¹

¹Parallel Computing Lab, Intel Labs
²Facebook, 1 Hacker Way, Menlo Park, CA

ABSTRACT
This paper presents the first comprehensive empirical study of the efficacy of the Brain Floating Point (BFP) format for training Deep Neural Networks (DNNs).

FAST: DNN Training Under Variable Precision Block Floating Point with Stochastic Rounding

Sai Qian Zhang¹, Bradley McDaneF¹, and H.T. Kung¹
¹Harvard University
Franklin and Marshall College

A Little Shifting Goes a Long Way

Yuhan¹, Ritchie Zhao, Venmugil Elango, Rasoul Shafiqpour, Mathew Hall, Maral Mesmakhoshahi,
More, Levi Melnick, Maximilian Golub, Girish Varatkar, Lei Shao, Gaurav Kulkarni, Dmitry Melts,
Jasmine Klar, Renee L'Heureux, Matt Perry, Doug Burger, Eric Chung,
Zhaoxia (Summer) Deng, Sam Naghshineh, Jongsoo Park, Maxim Naumov

ABSTRACT—This paper introduces Block Data Representations (BDR), a framework for exploring and evaluating a wide spectrum of narrow-precision formats for deep learning. It enables comparison of popular quantization standards, such as INT8, R, new formats based on...

9.05433v2 [cs.LG] 29 Sep 2018

Neil Burgess, Sangwon Ha, Richard Grisenthwaite,
{neil.burgess, sangwon.ha, richard.grisenthwaite}@arm.com

Naveen Mellempudi, Marius Cornes, Alexander Heinecke, Pradeep Dubey
{naveen.k.mellempudi, marius.cornes, alexander.heinecke, pradeep.dubey}@intel.com

ABSTRACT
BFP is a natural progression for accelerating deep learning training inference. It is a natural progression for accelerating deep learning training inference. It is a natural progression for accelerating deep learning training inference. It is a natural progression for accelerating deep learning training inference.

Our Goal

- Oracle?



Our Goal

- Oracle?
- Learn the Datatype
 - Use gradient descent



Our Goal

- Oracle?
- Learn the Datatype
 - Use gradient descent
- Monitor Loss
 - Adapt accordingly



Our Goal

- Oracle?
- Learn the Datatype
 - Use gradient descent
- Monitor Loss
 - Adapt accordingly
- Result
 - Less bits
 - Less energy
 - Less time
 - Bigger models



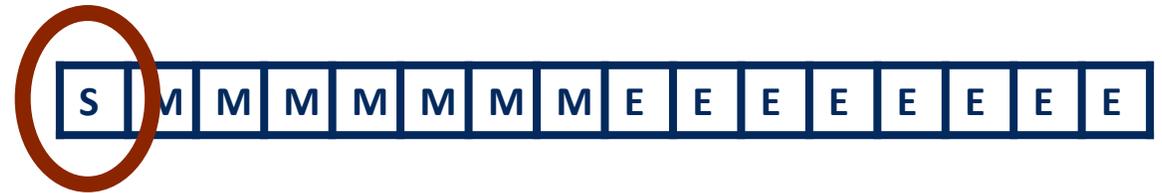
Datatype – Floating Point

- $Value = (-1)^S \times M \times 2^E$



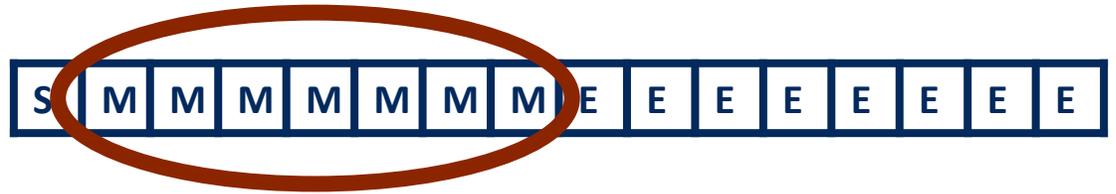
Datatype – Floating Point

- $Value = (-1)^S \times M \times 2^E$
- Sign is trivial



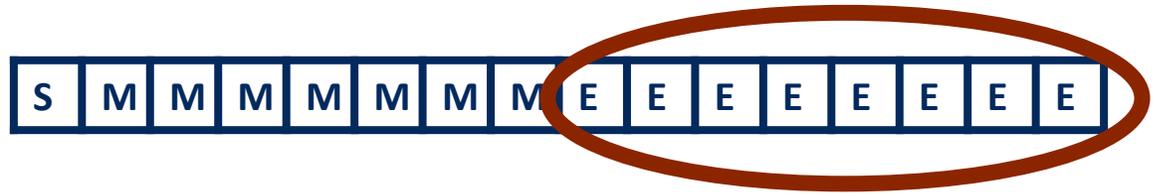
Datatype – Floating Point

- $Value = (-1)^S \times M \times 2^E$
- Sign is trivial
- Mantissa
 - Precision



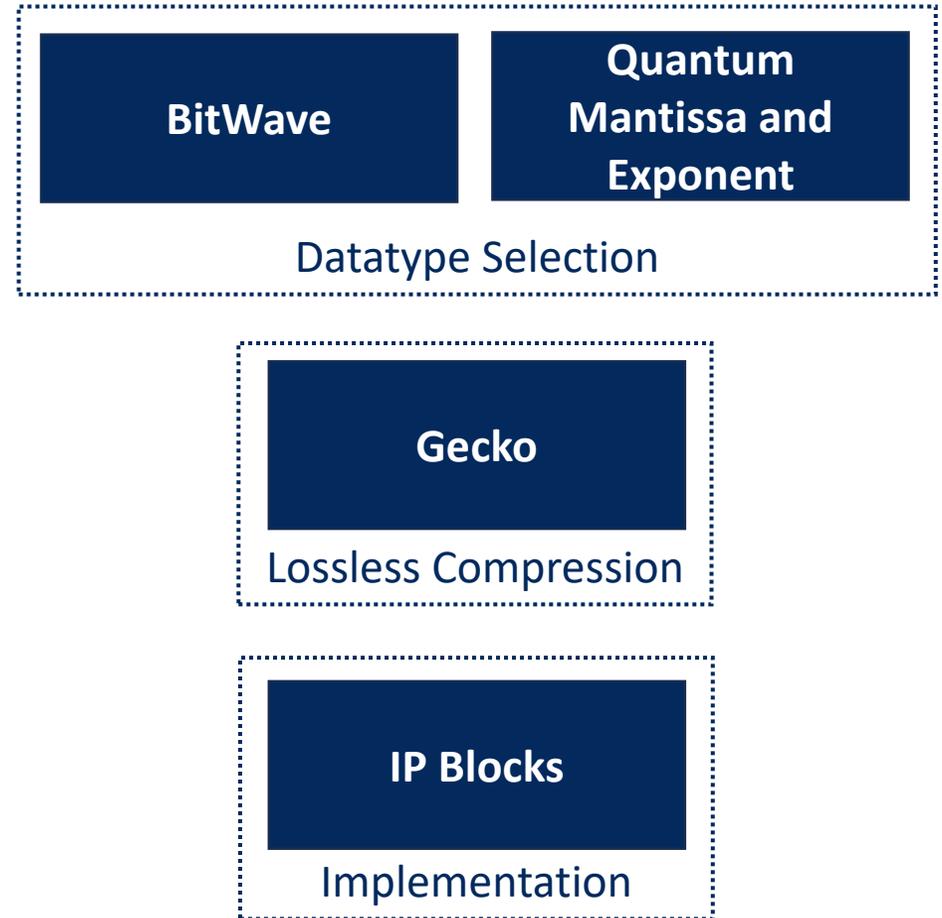
Datatype – Floating Point

- $Value = (-1)^S \times M \times 2^E$
- Sign is trivial
- Mantissa
 - Precision
- Exponent
 - Range



Schrödinger's FP

- Machine Learning
 - *Quantum Mantissa and Exponent* (**4.7x** reduction)
- Black Box Sampling
 - *BitWave* (**3.2x** reduction)
- Exponent compression
 - *Gecko* (boost to **5.6x** and **4.6x**)
- Hardware IP blocks



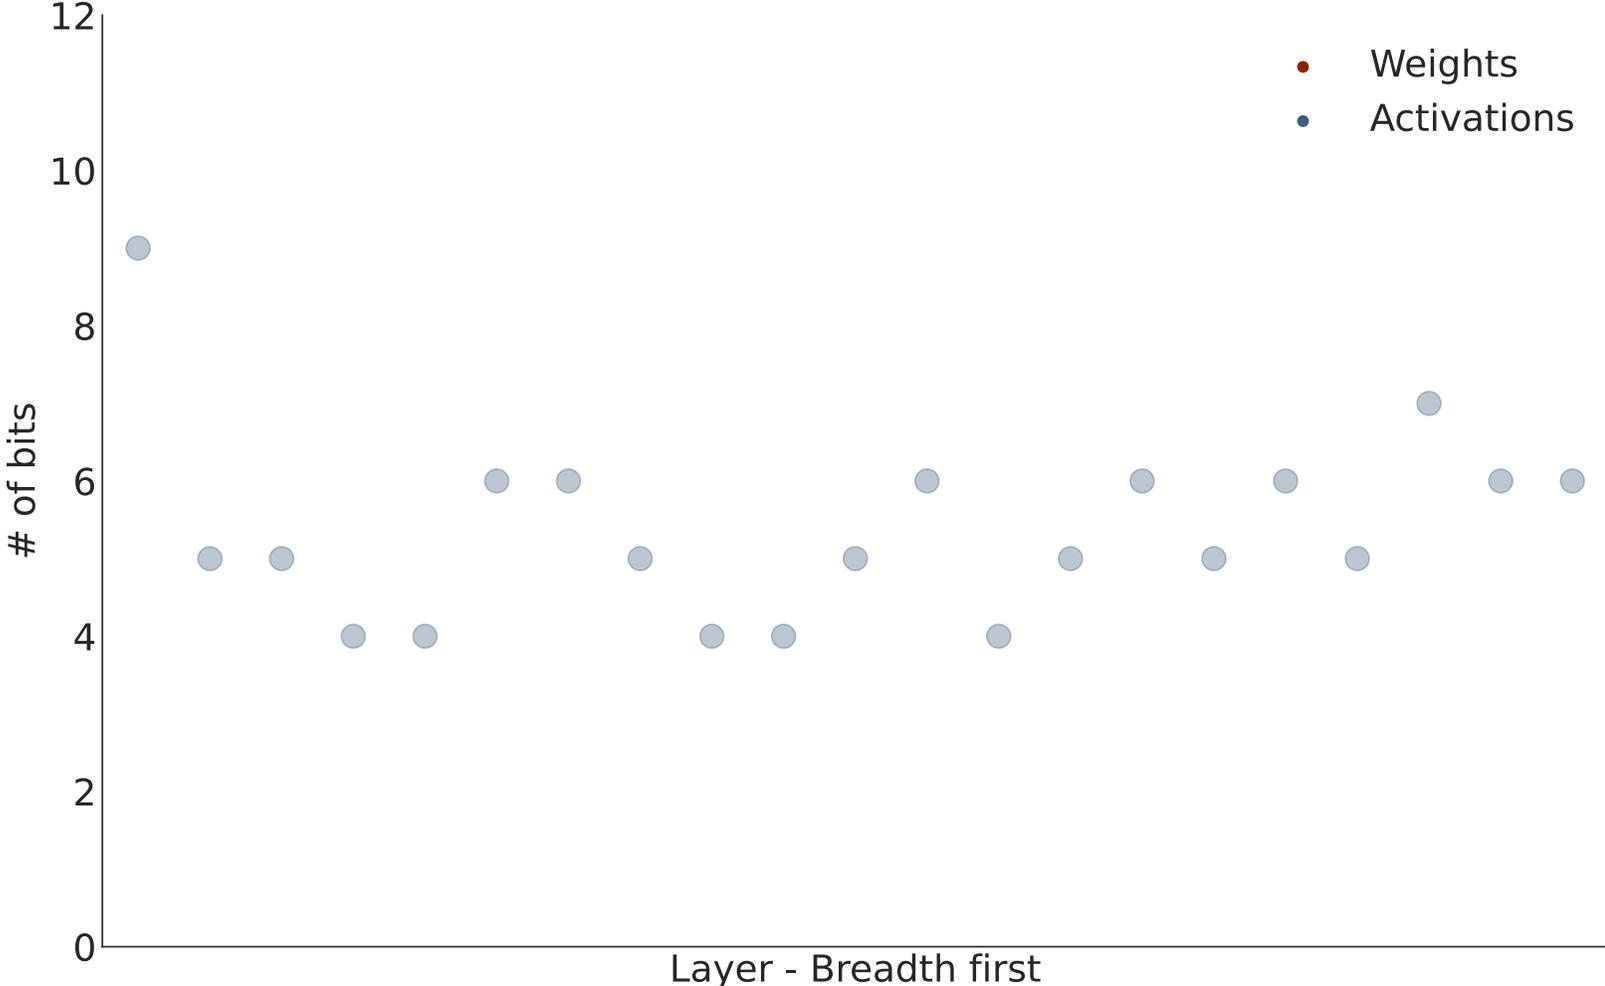
Schrödinger's FP

- Machine Learning
 - *Quantum Mantissa and Exponent*
- Black Box Sampling
 - *BitWave*
- Exponent compression
 - *Gecko*
- Hardware IP blocks



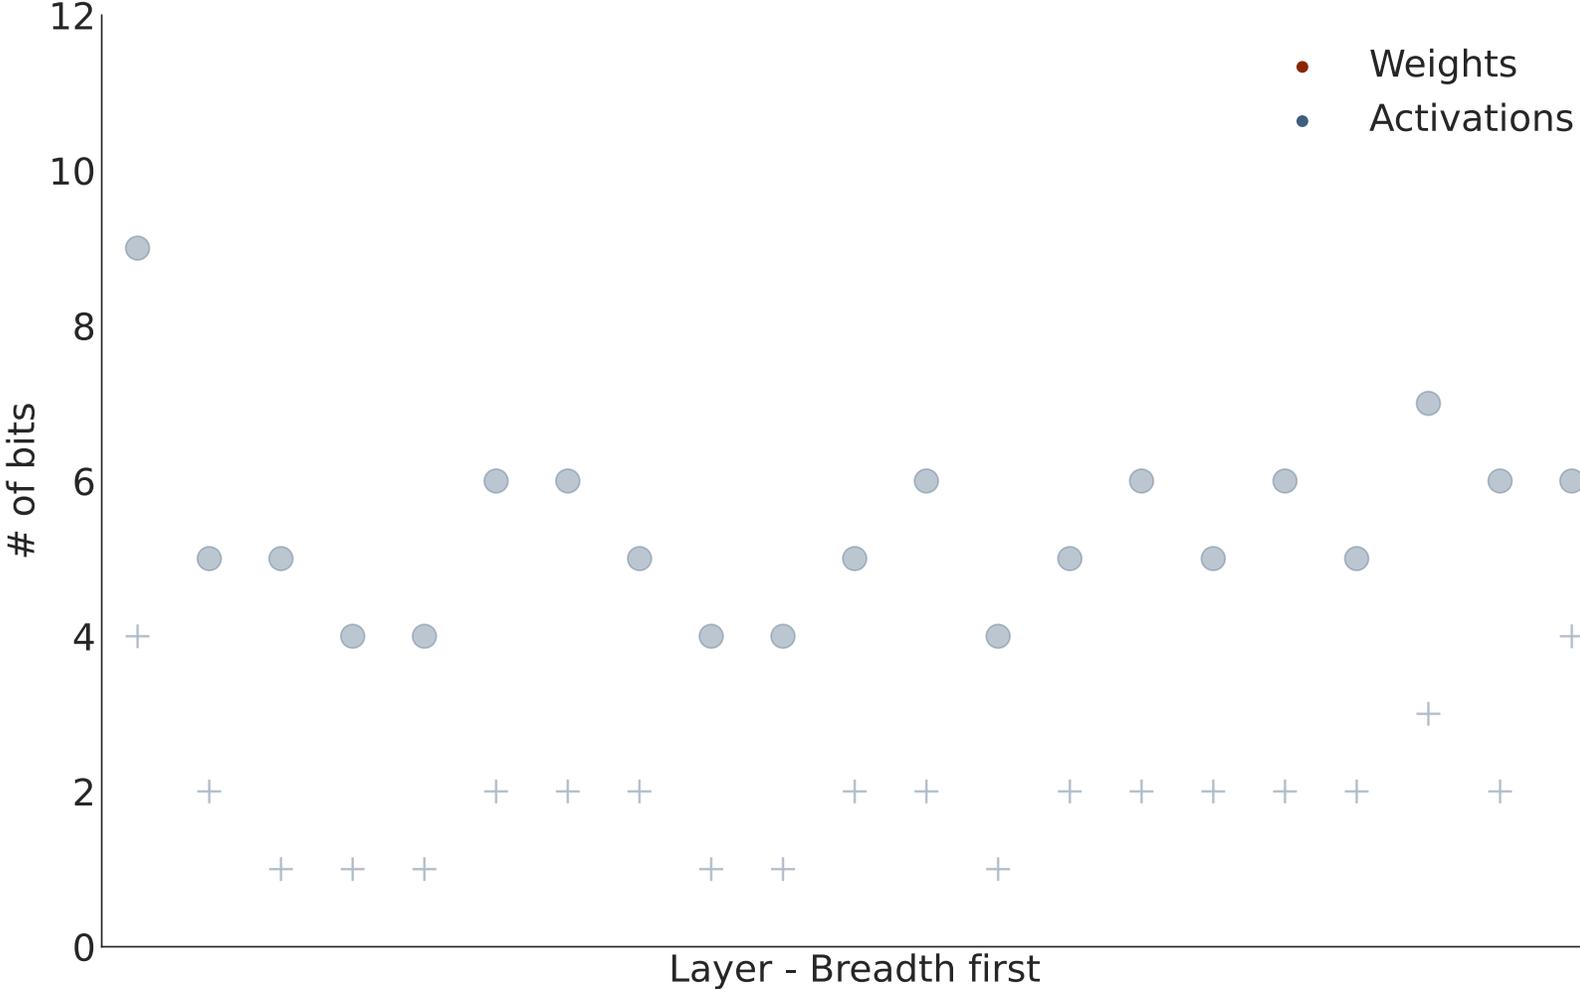
Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?



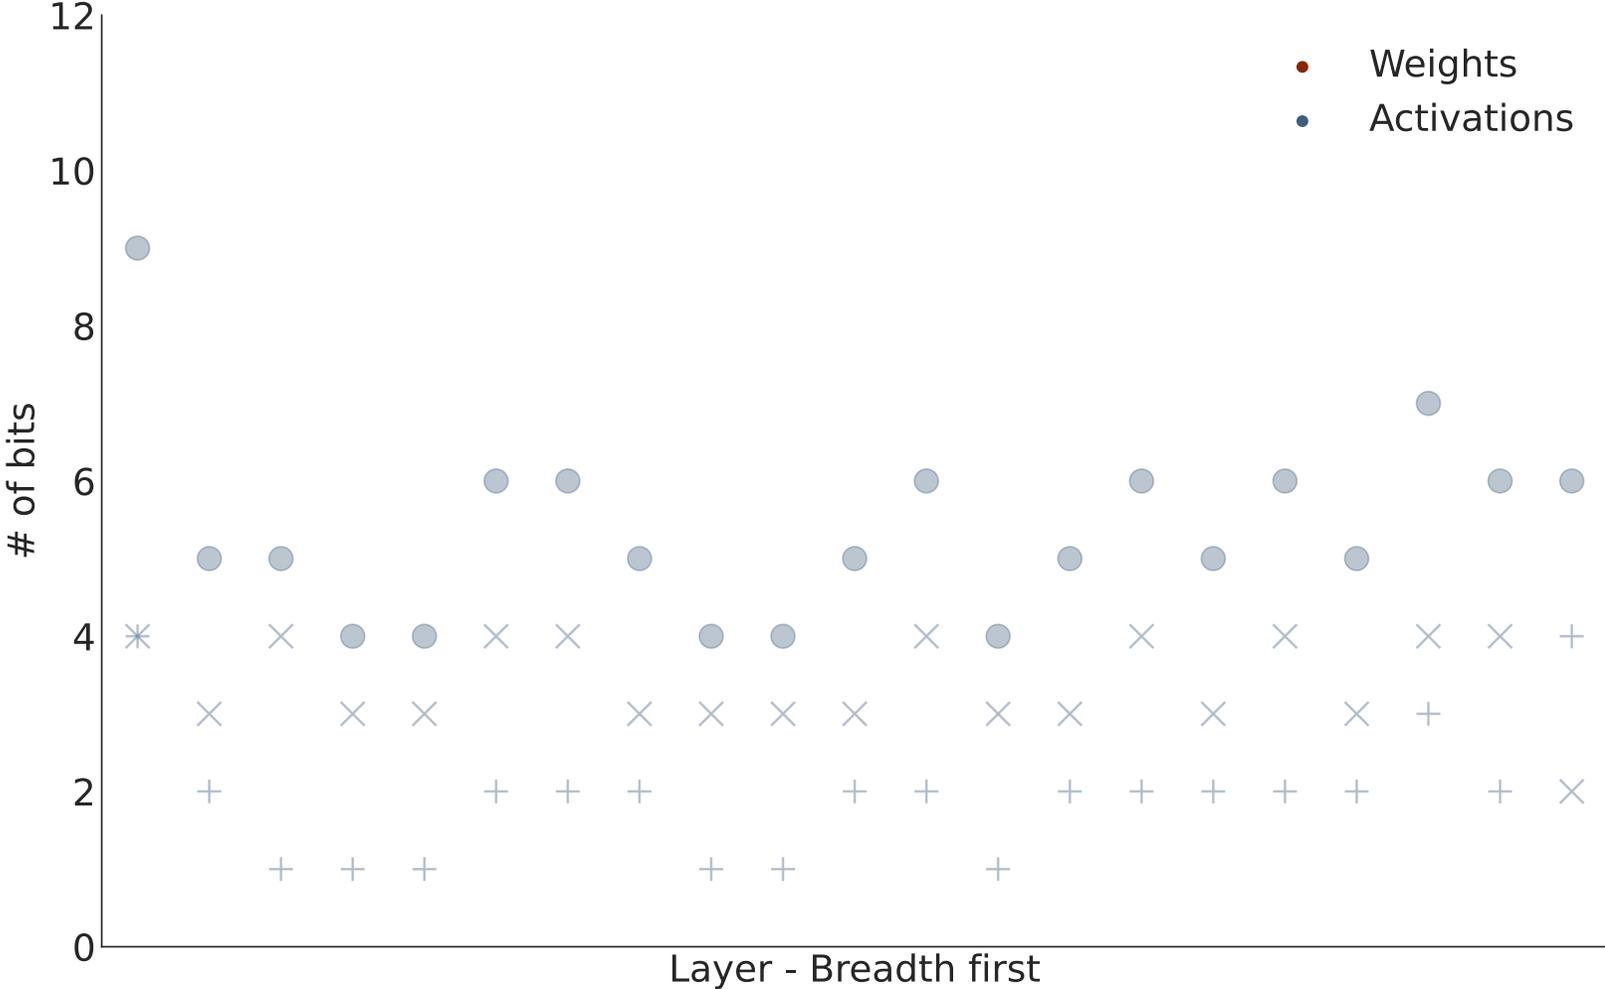
Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?



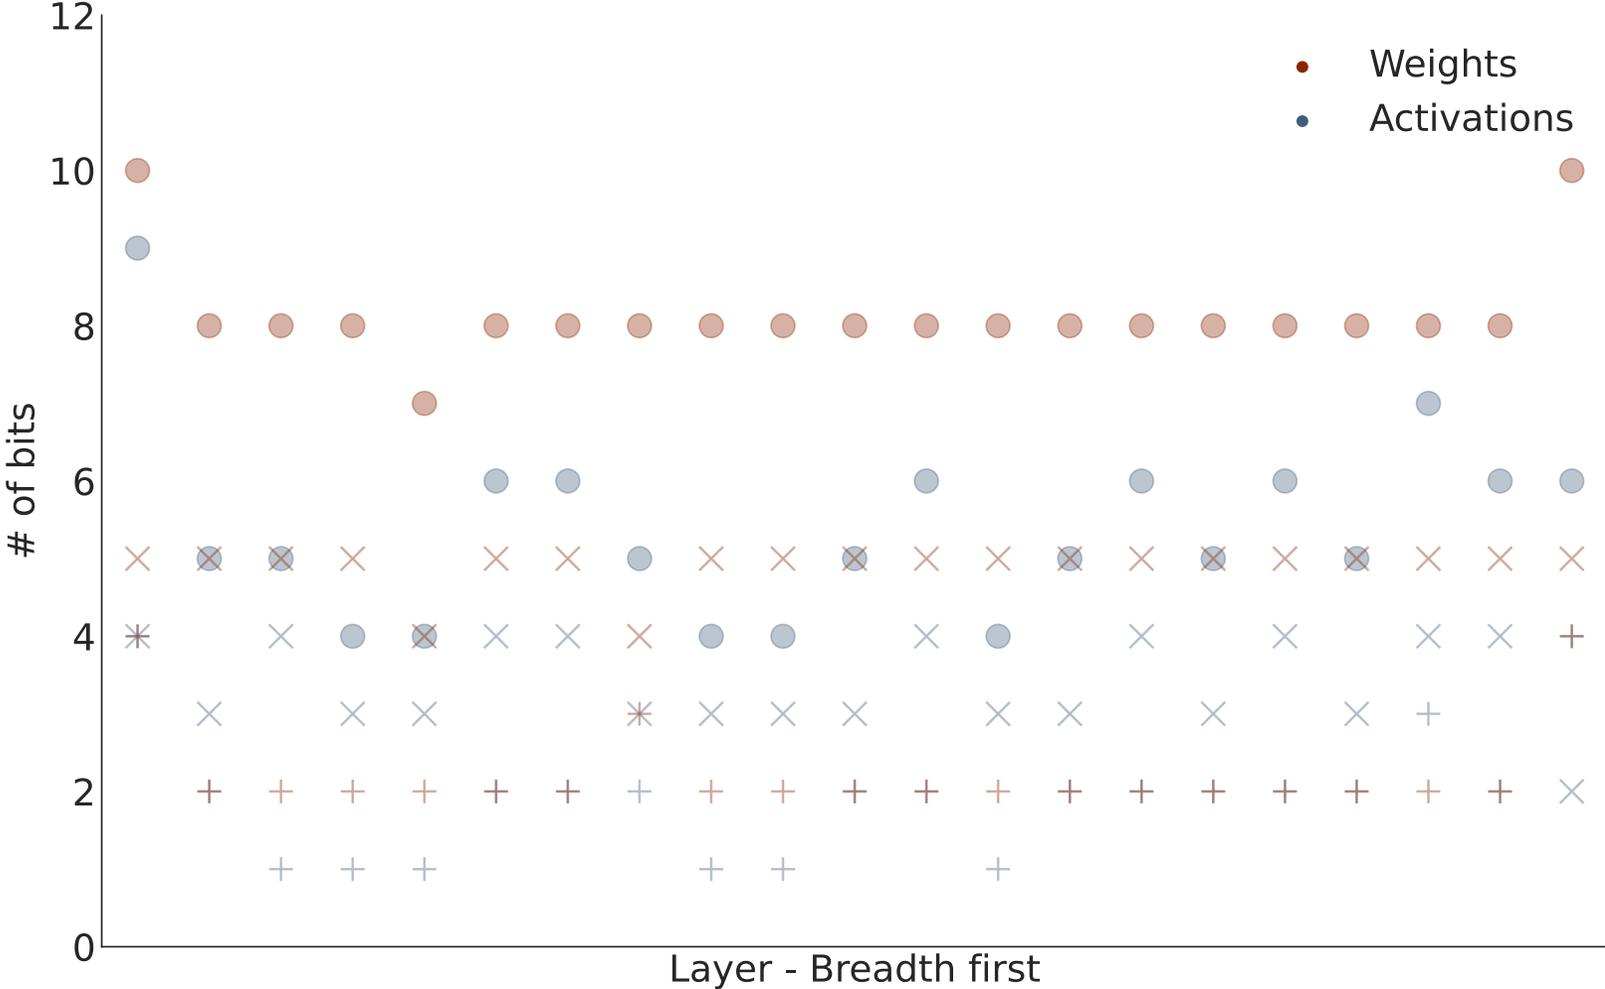
Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?



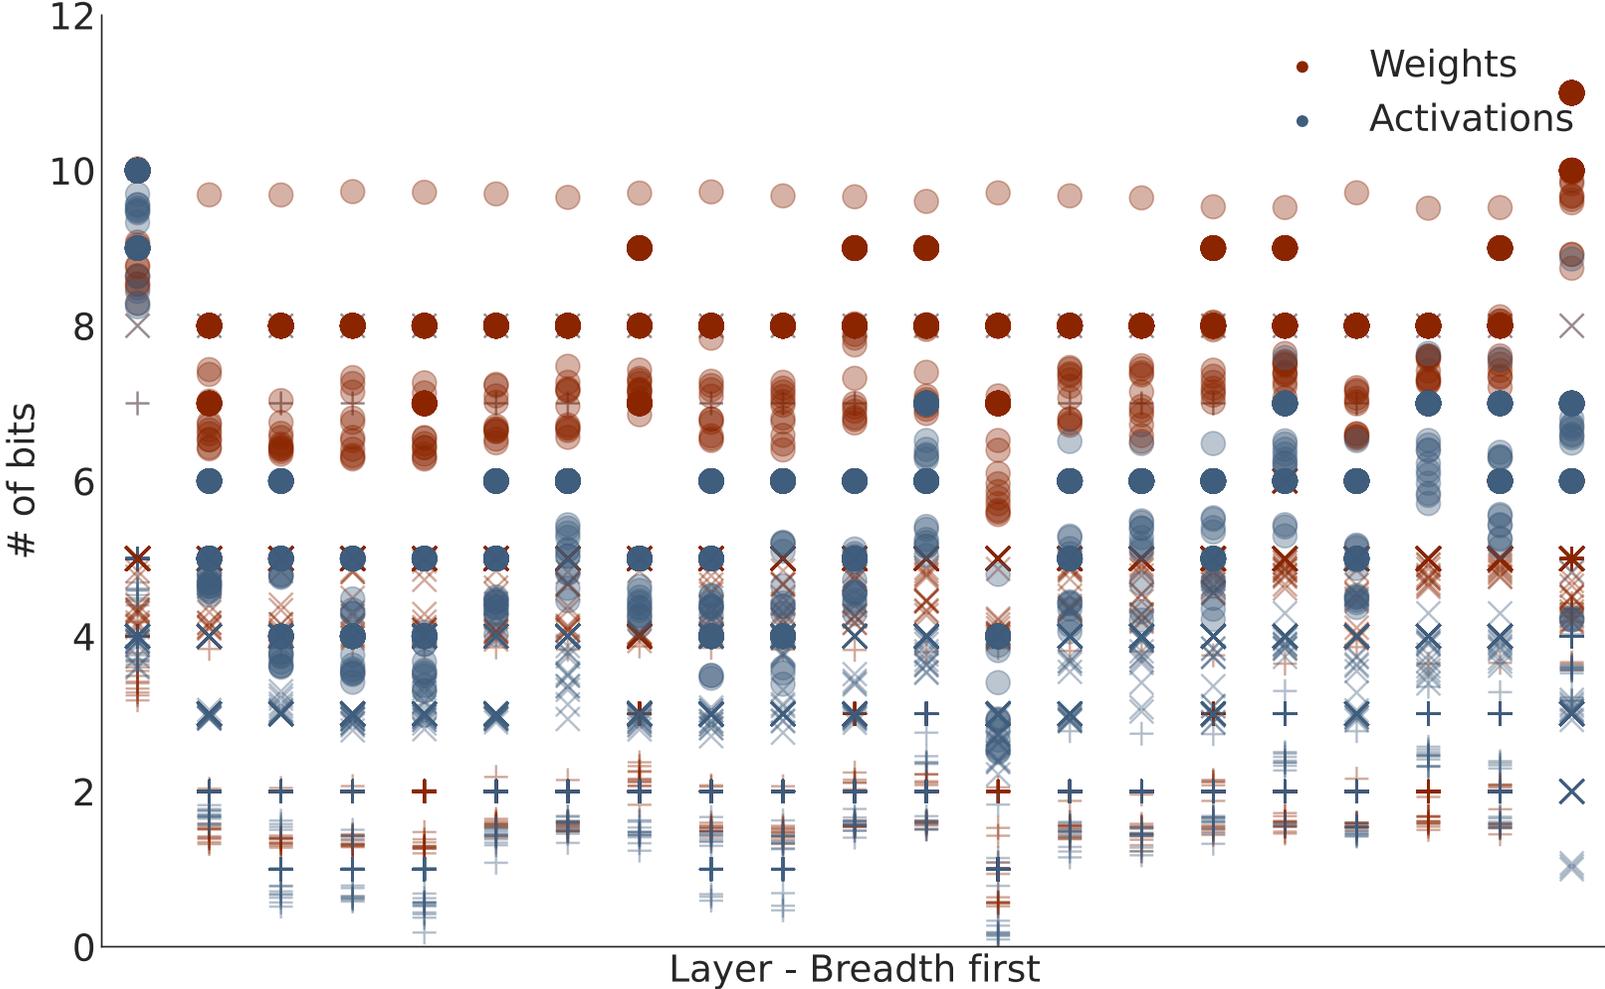
Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?



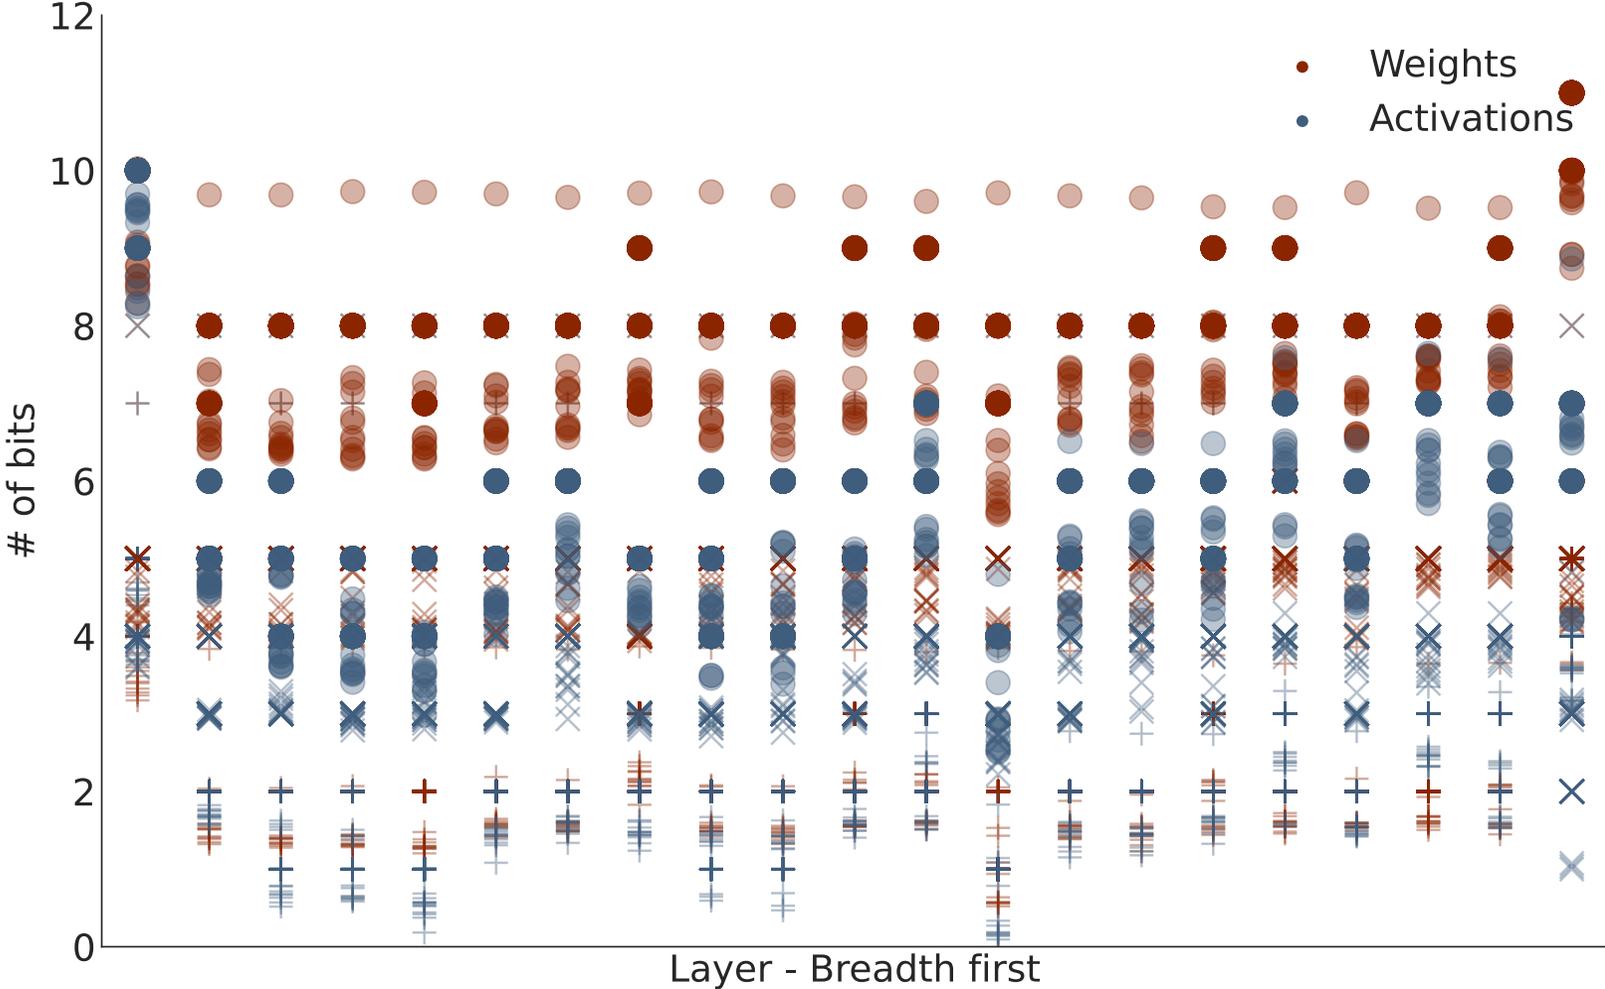
Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?



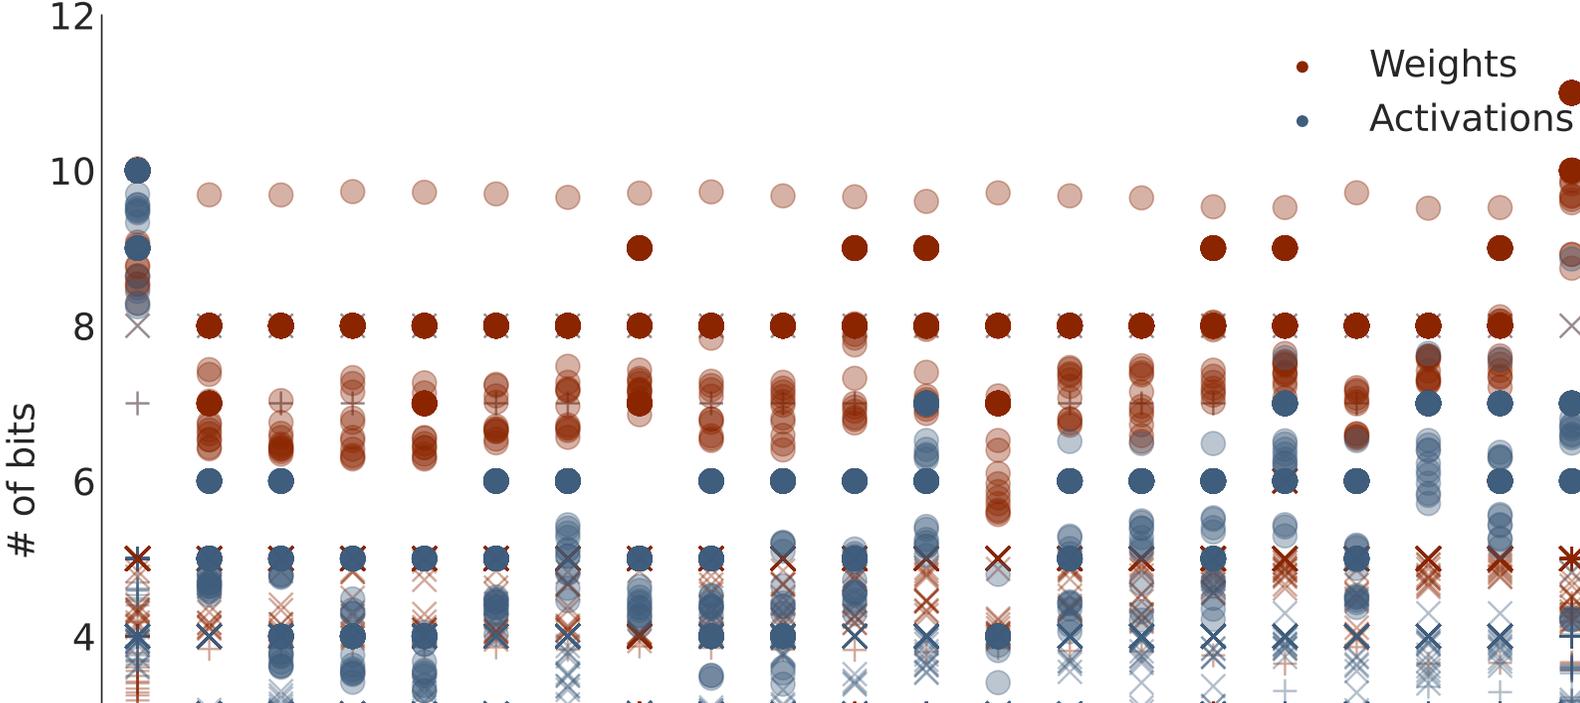
Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?
- Less than 6b on average!
 - Much better than network wide!



Per Tensor Datatype

- ResNet18
 - Per Tensor?
 - Single datatype?
- Less than 6b on average!
 - Much better than network wide!



No chance to guess effectively and expect to train well!

Quantum Mantissa Loss

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according to priority
 - Tensor footprint
 - # of operations

Quantum Mantissa Loss

- Our Loss  **Original Loss**
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according to priority
 - Tensor footprint
 - # of operations

Quantum Mantissa Loss

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according to priority
 - Tensor footprint
 - # of operations

Mantissa Lengths



Quantum Mantissa Loss

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Mantissa Weights
 - Weighted according to priority
 - Tensor footprint
 - # of operations

Quantum Mantissa Loss

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - **Regularizer Strength** 
 - Weighted according to priority
 - Tensor footprint
 - # of operations

Quantum Mantissa Bitlength

- Int Datatype



Quantum Mantissa Bitlength

- Int Datatype
- Non-Int Datatype



Quantum Mantissa Bitlength

- Int Datatype
- Non-Int Datatype



Quantum Mantissa Bitlength

- Int Datatype
- Non-Int Datatype



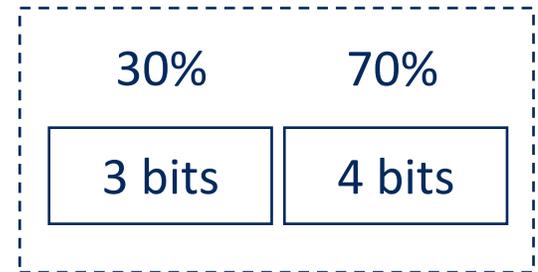
Quantum Mantissa Bitlength

- Int Datatype
- Non-Int Datatype



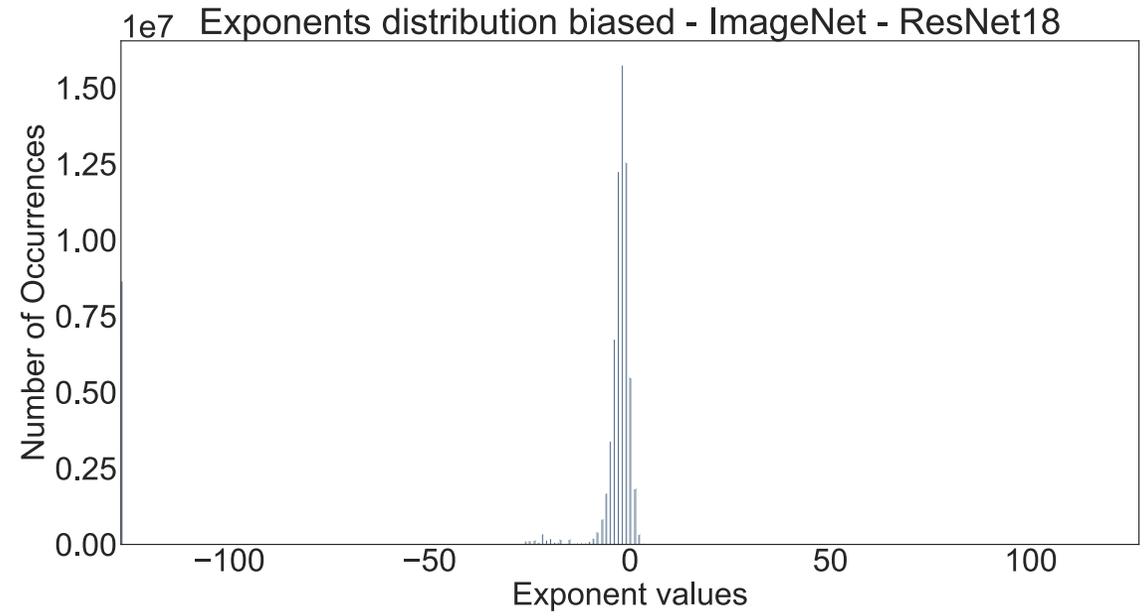
3.7 bits?

=



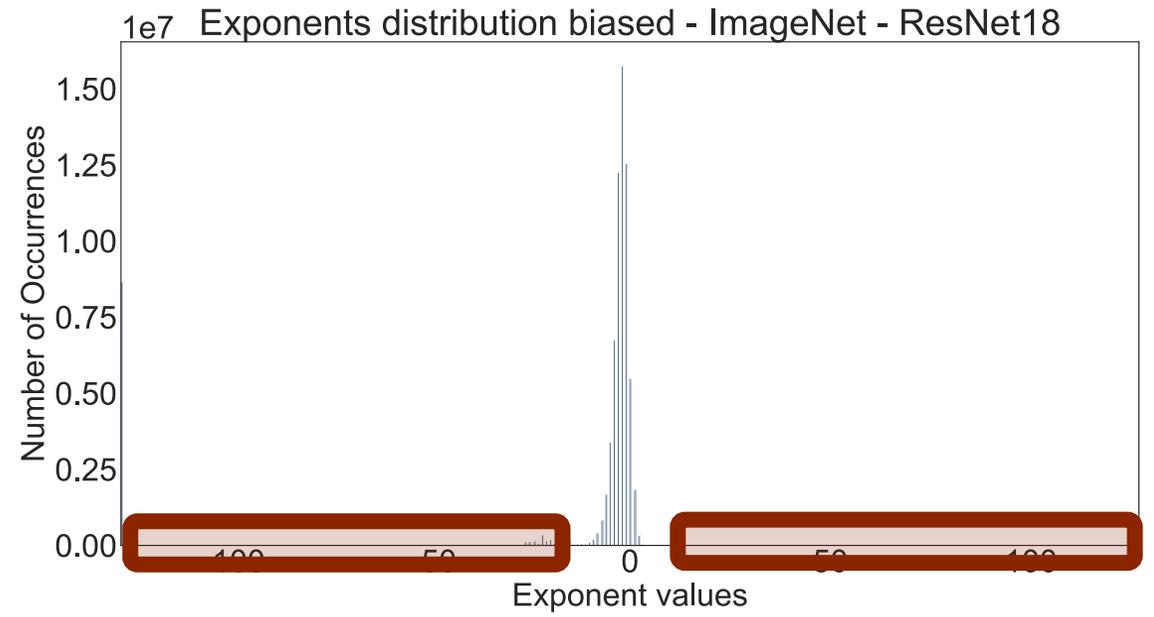
Quantum Exponent Range

- Distribution



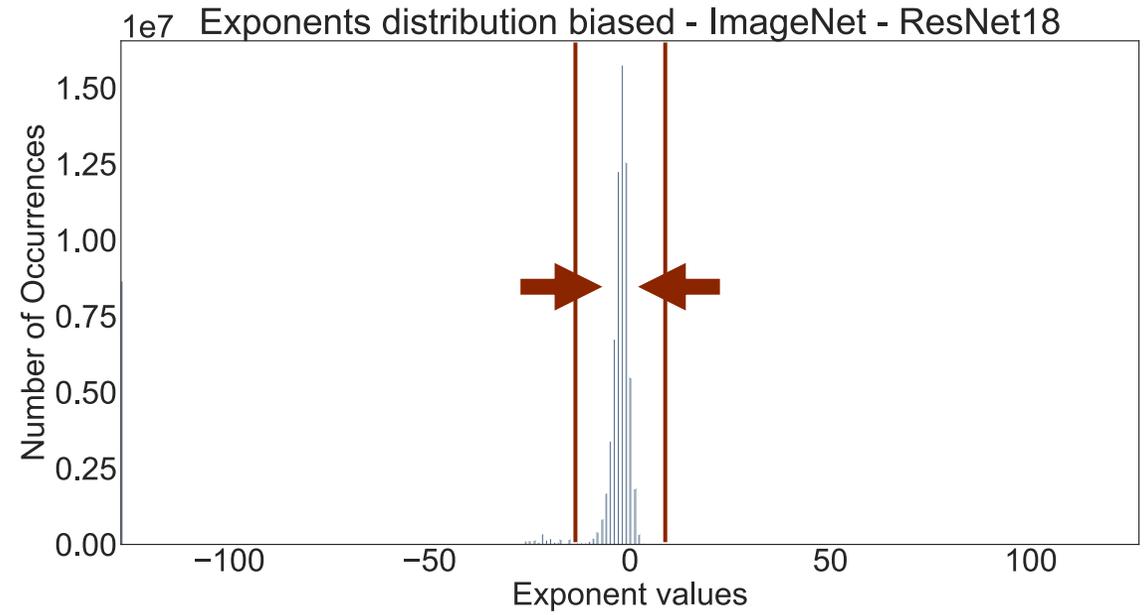
Quantum Exponent Range

- Distribution



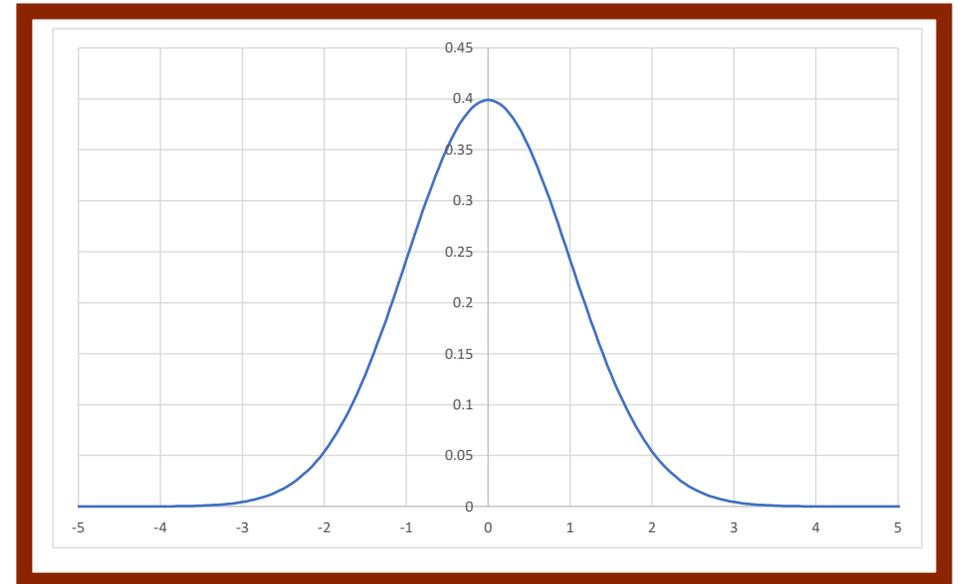
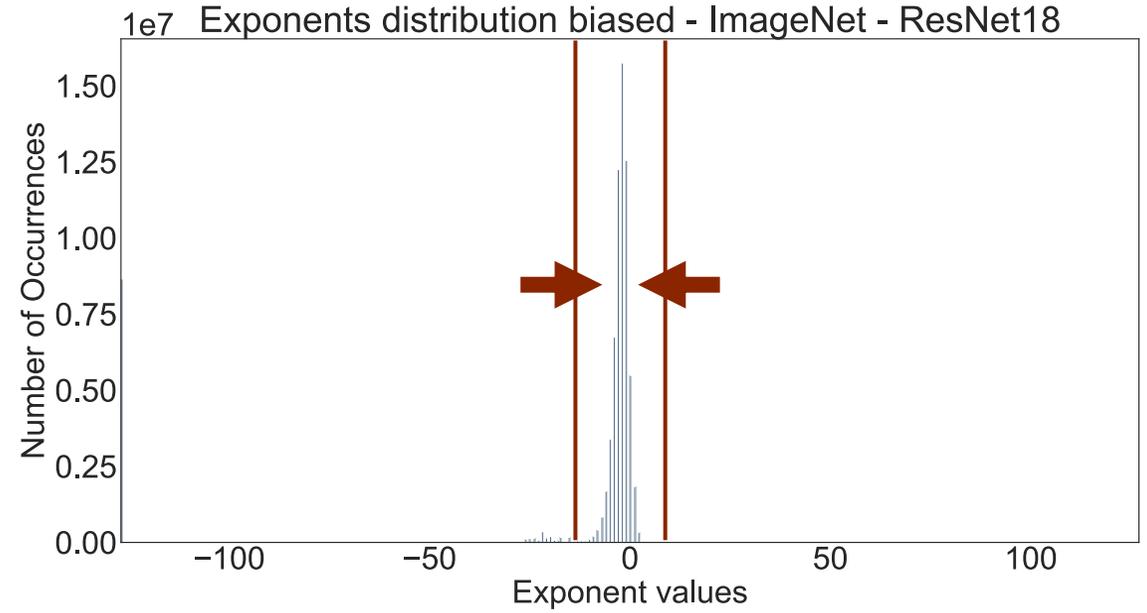
Quantum Exponent Range

- Distribution
- Compress range



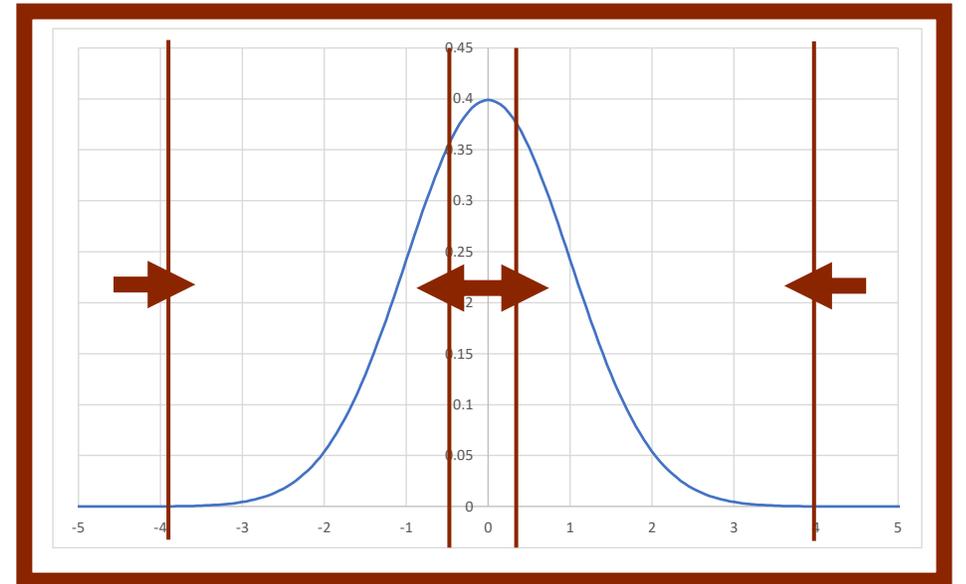
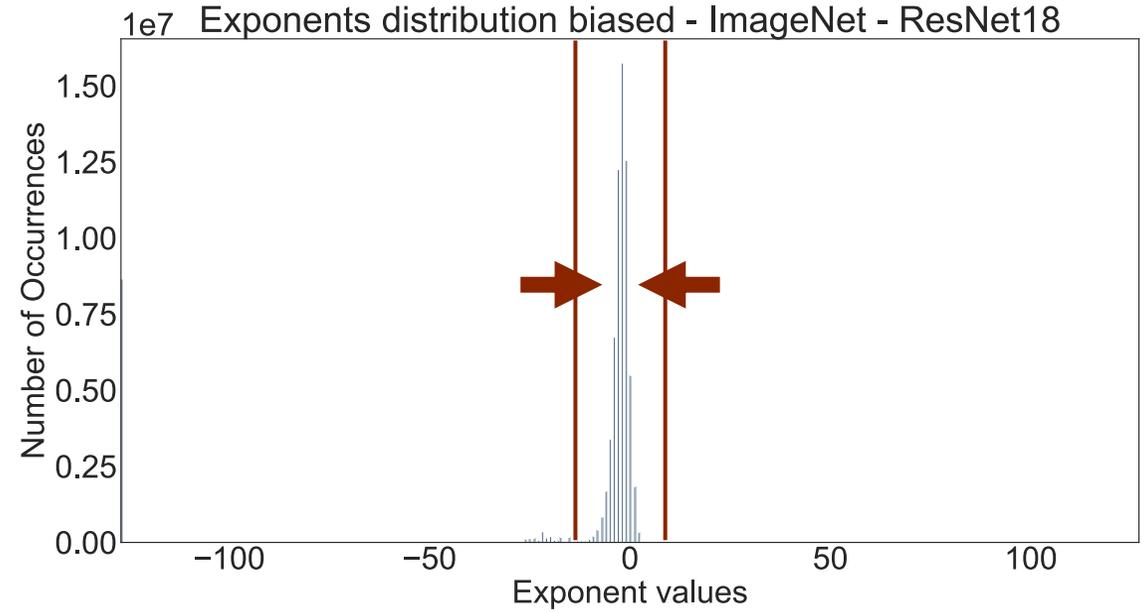
Quantum Exponent Range

- Distribution
- Compress range
- $Value = (-1)^S \times M \times 2^E$



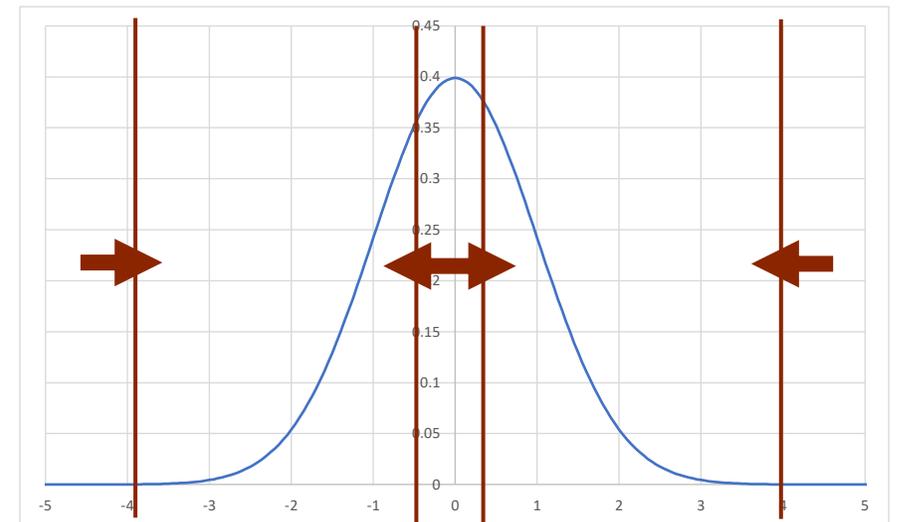
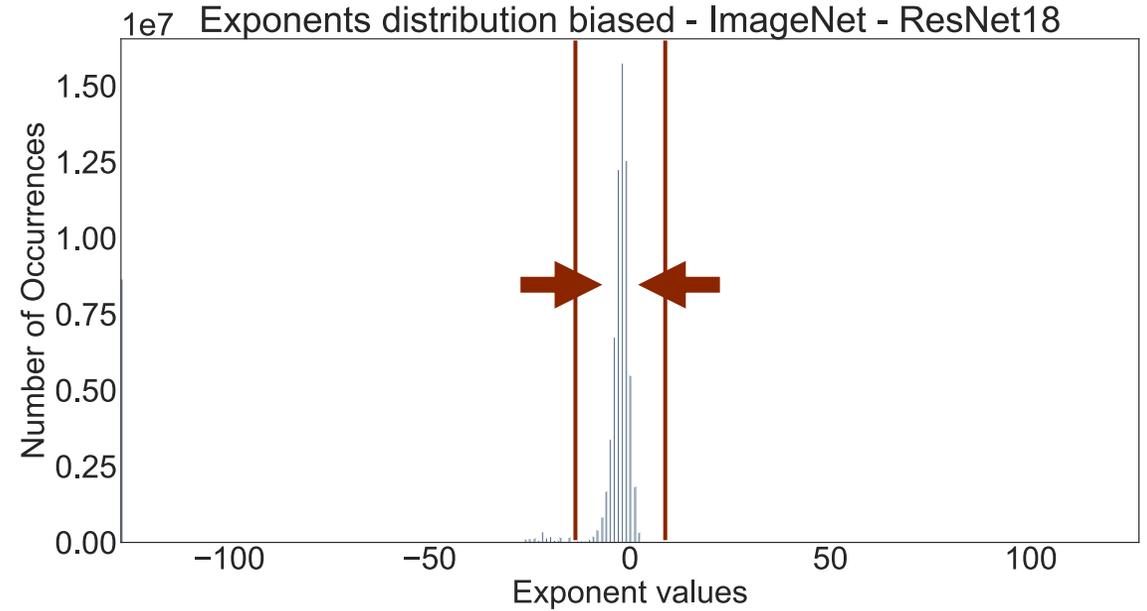
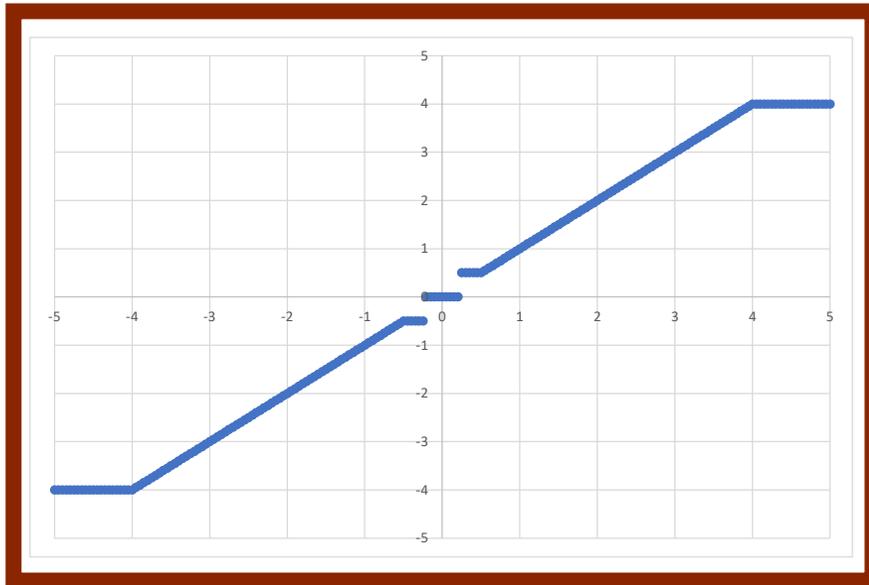
Quantum Exponent Range

- Distribution
- Compress range
- $Value = (-1)^S \times M \times 2^E$



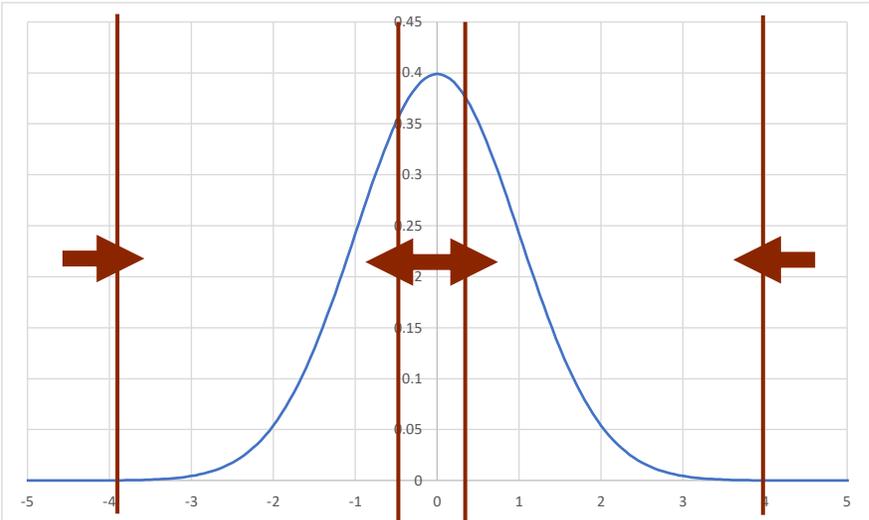
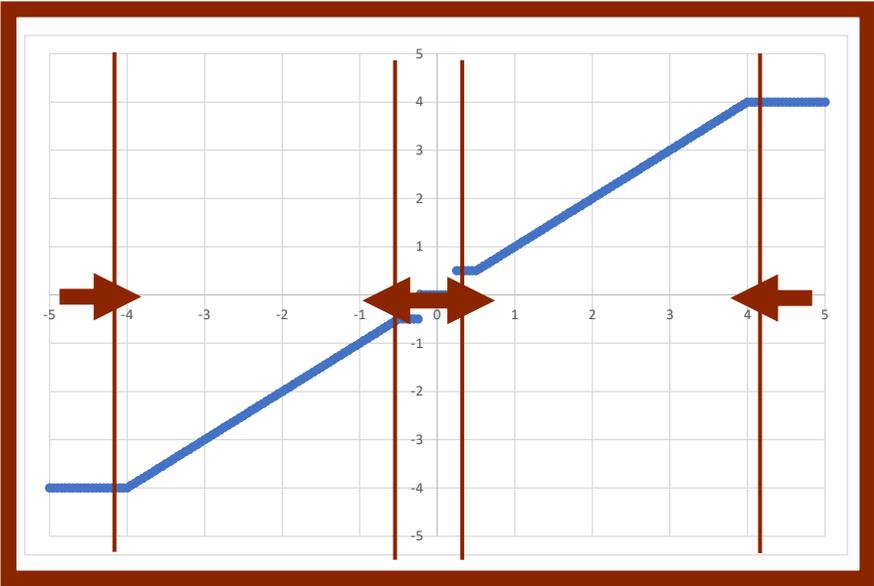
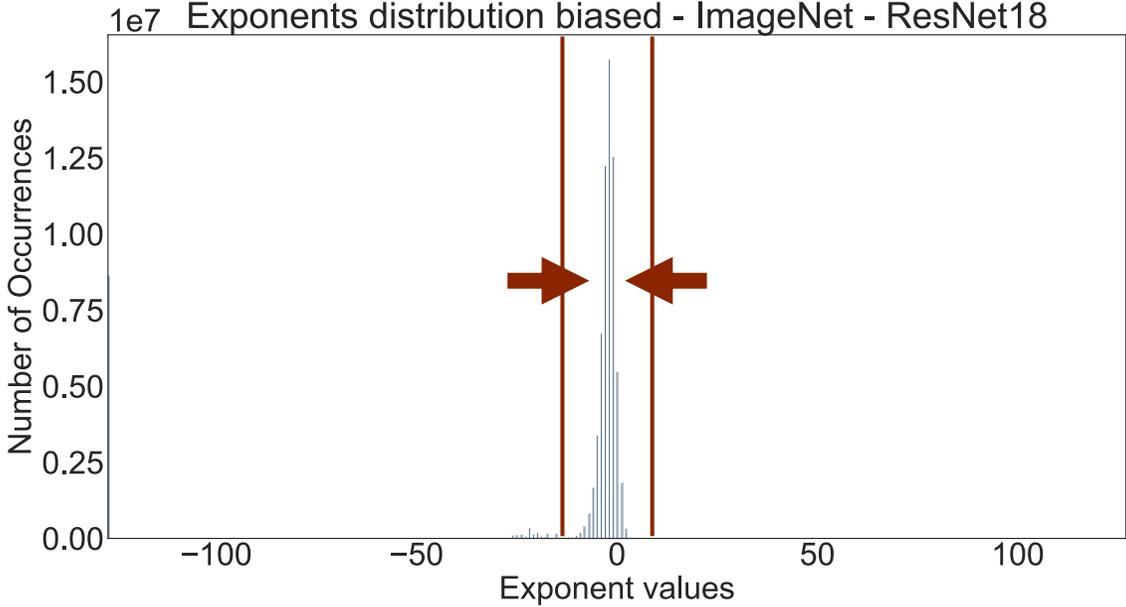
Quantum Exponent Range

- Distribution
- Compress range
- $Value = (-1)^S \times M \times 2^E$
- Range(x) overlay



Quantum Exponent Range

- Distribution
- Compress range
- $Value = (-1)^S \times M \times 2^E$
- Range(x) overlay



Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority

Quantum Exponent

- Our Loss  **Original Loss**
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority

Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority

Exponent Lengths



Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority

Exponent Weights



Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority

Regularizer Strength



Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority
- Int Datatype



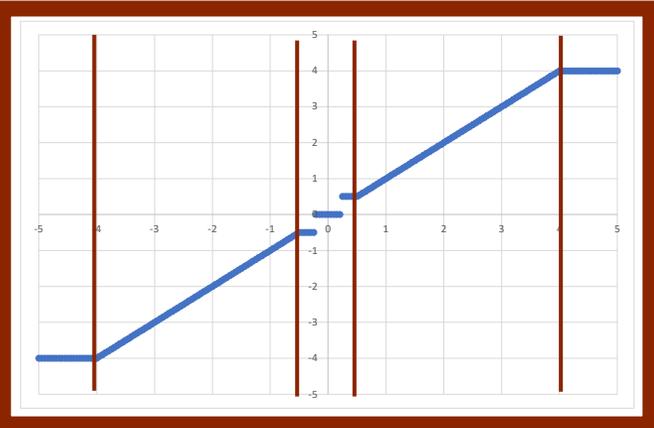
Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority
- Int Datatype
- Non-Int Datatype



Quantum Exponent

- Our Loss
 - $L = L_0 + \gamma \times (\sum \lambda_i \times \alpha_i)$
 - Weighted according priority
- Int Datatype
- Non-Int Datatype
- Determine boundaries



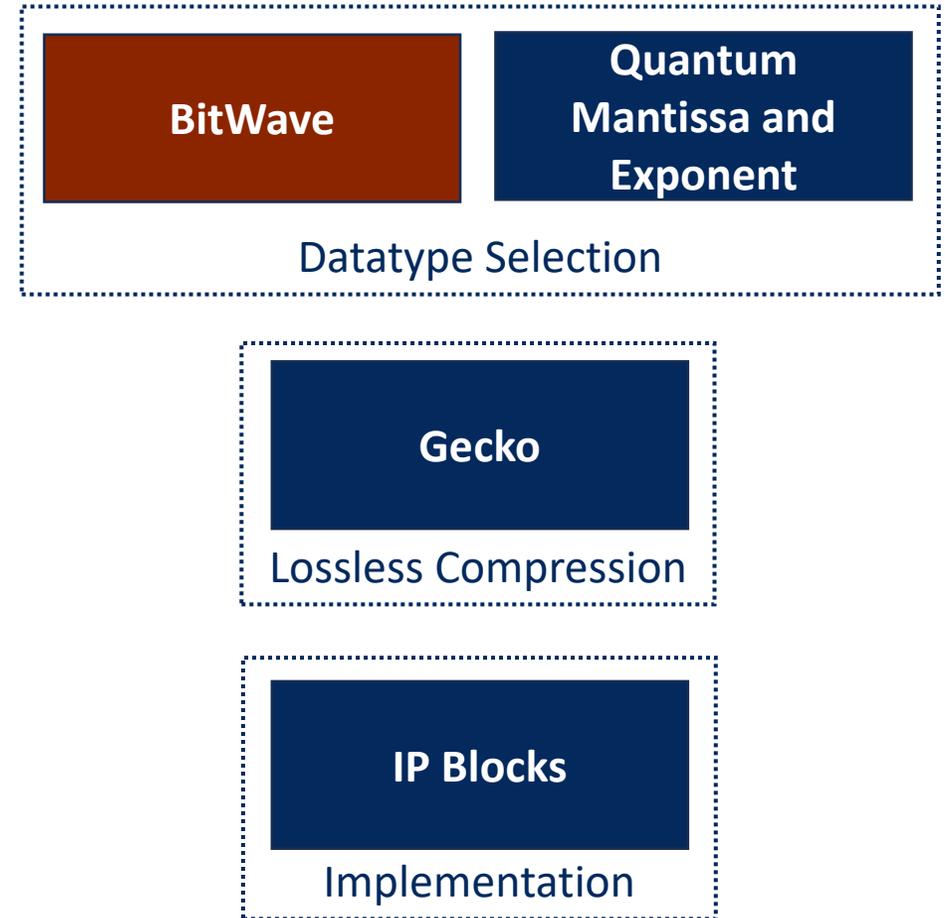
3.7 bits?

=



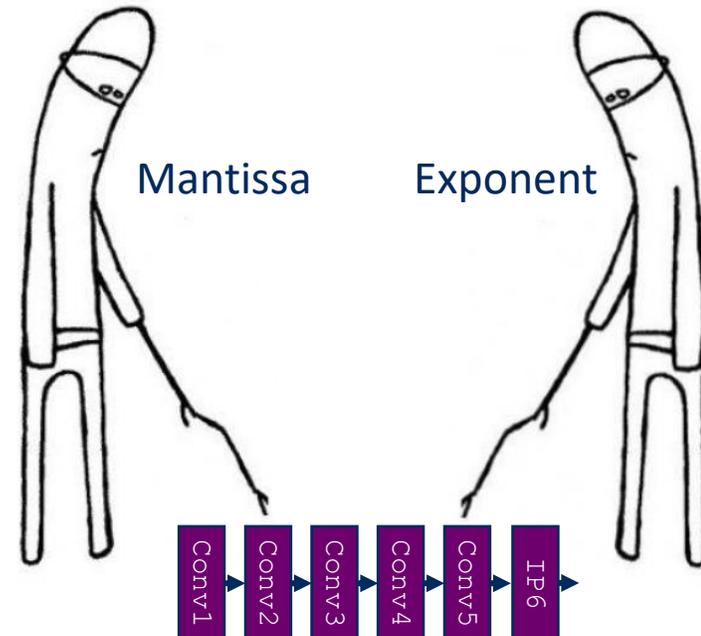
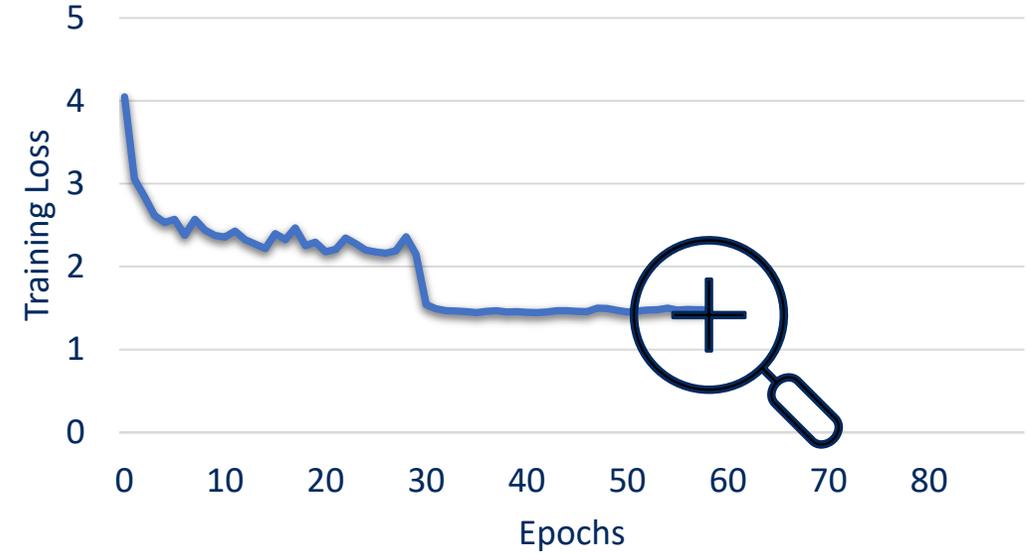
Schrödinger's FP

- Machine Learning
 - *Quantum Mantissa and Exponent*
- Black Box Sampling
 - *BitWave*
- Exponent compression
 - *Gecko*
- Hardware IP blocks



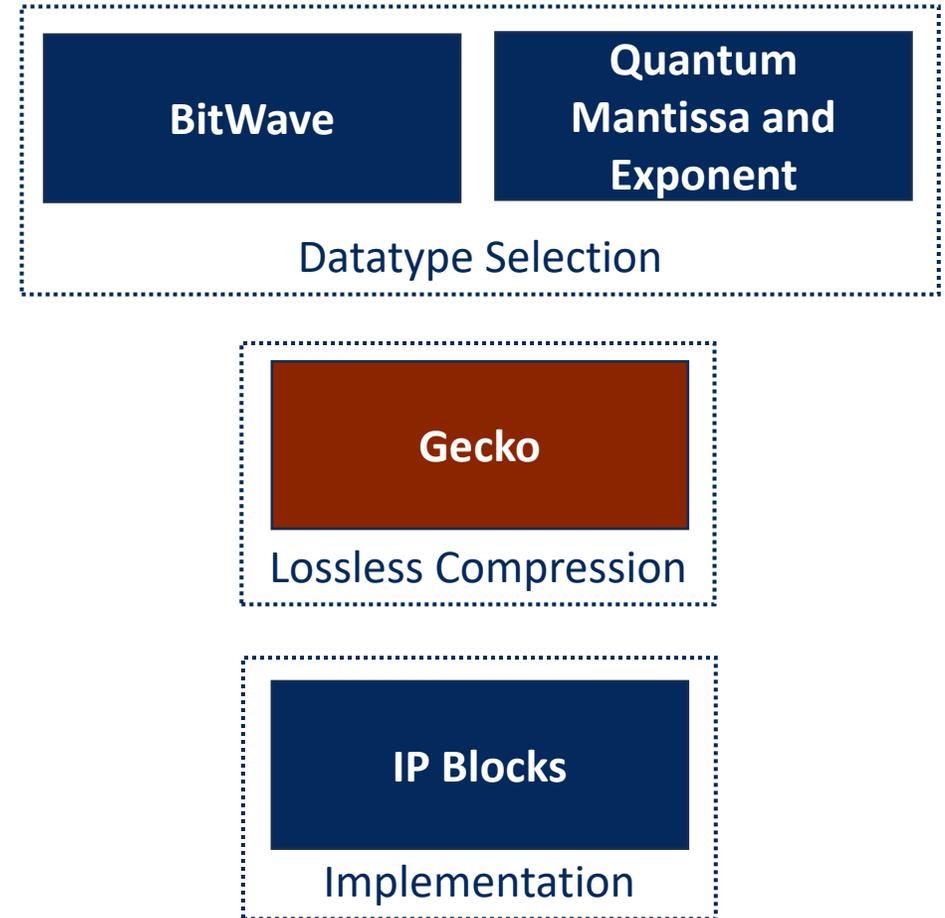
BitWave

- Observe Loss
 - Adjust bitlengths
 - Mantissa
 - Exponent
- Training is forgiving
- Network wide



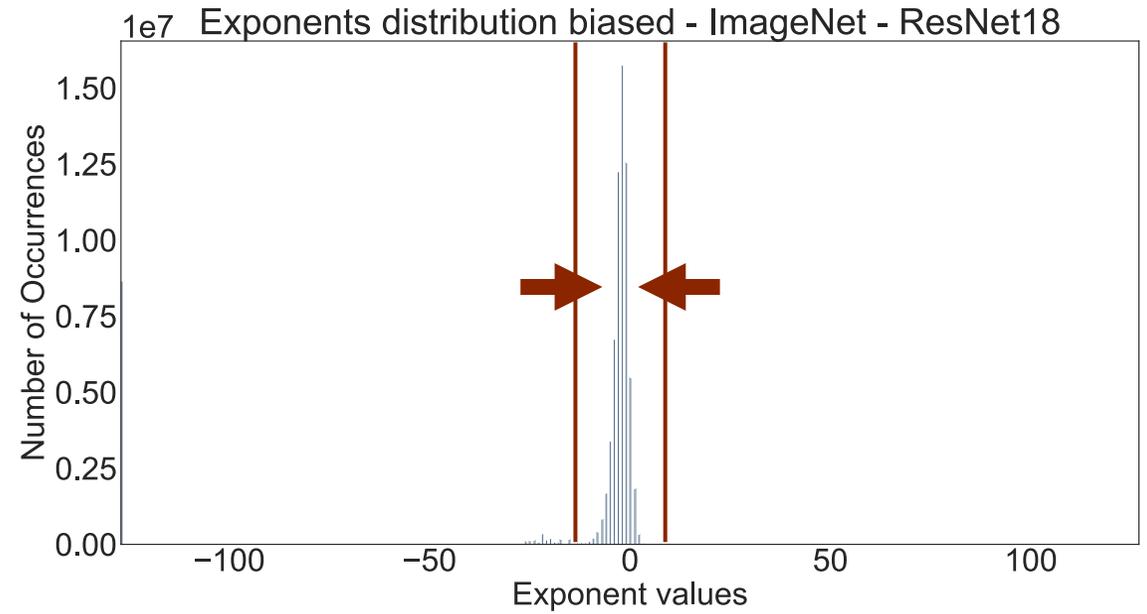
Schrödinger's FP

- Machine Learning
 - *Quantum Mantissa and Exponent*
- Black Box Sampling
 - *BitWave*
- Exponent compression
 - *Gecko*
- Hardware IP blocks



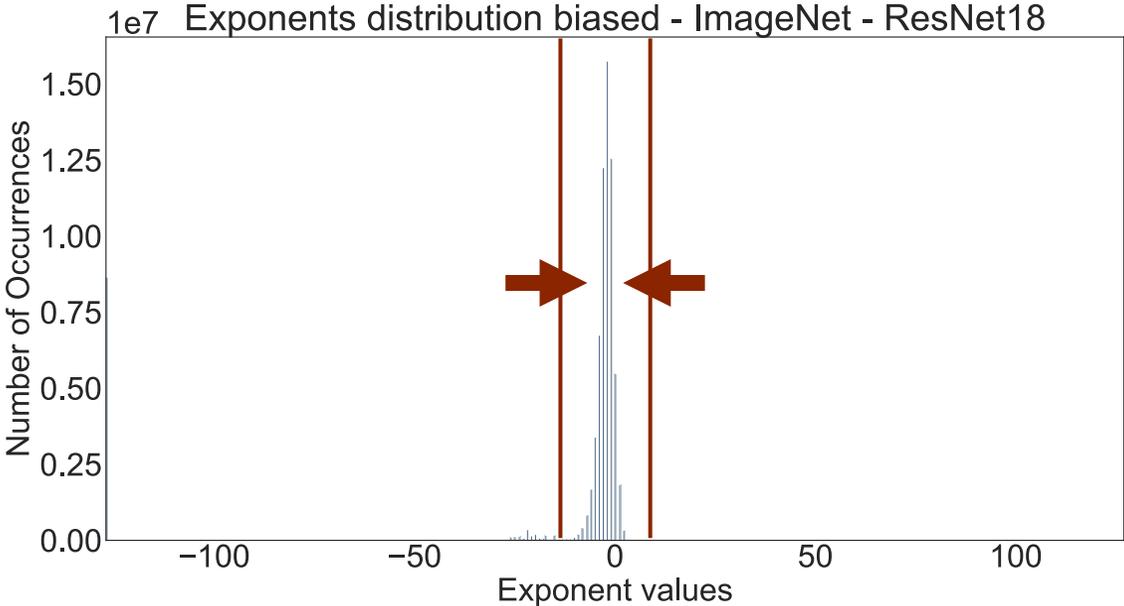
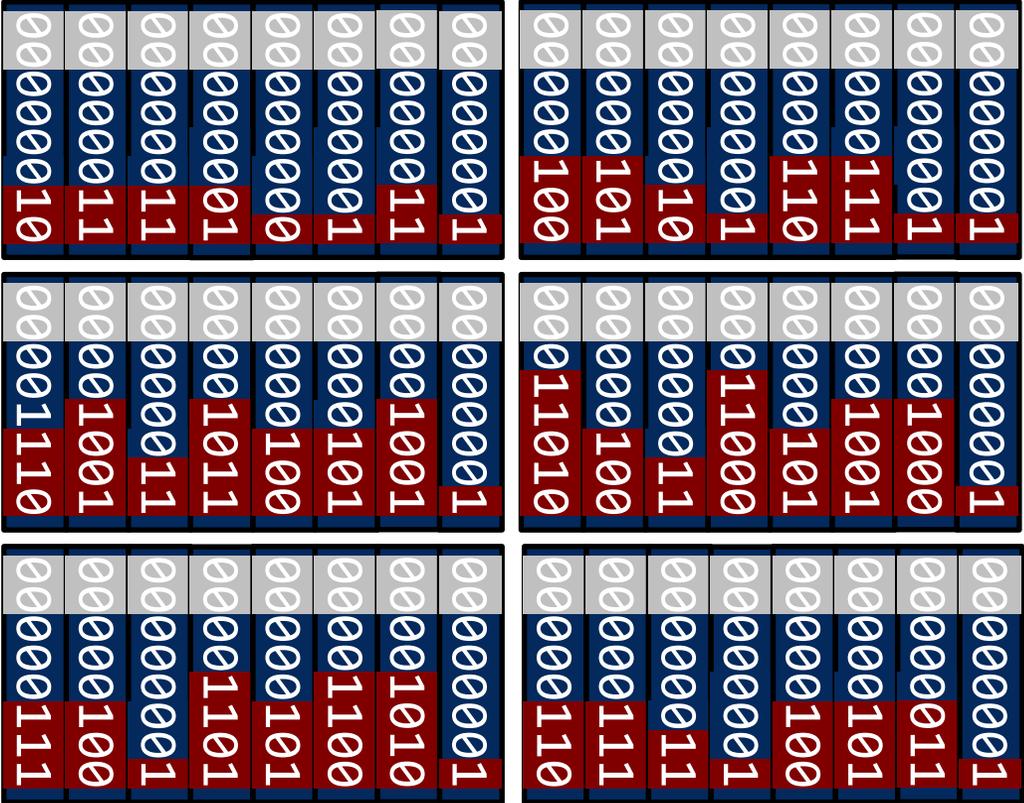
Exponent – *Gecko*

- Distribution



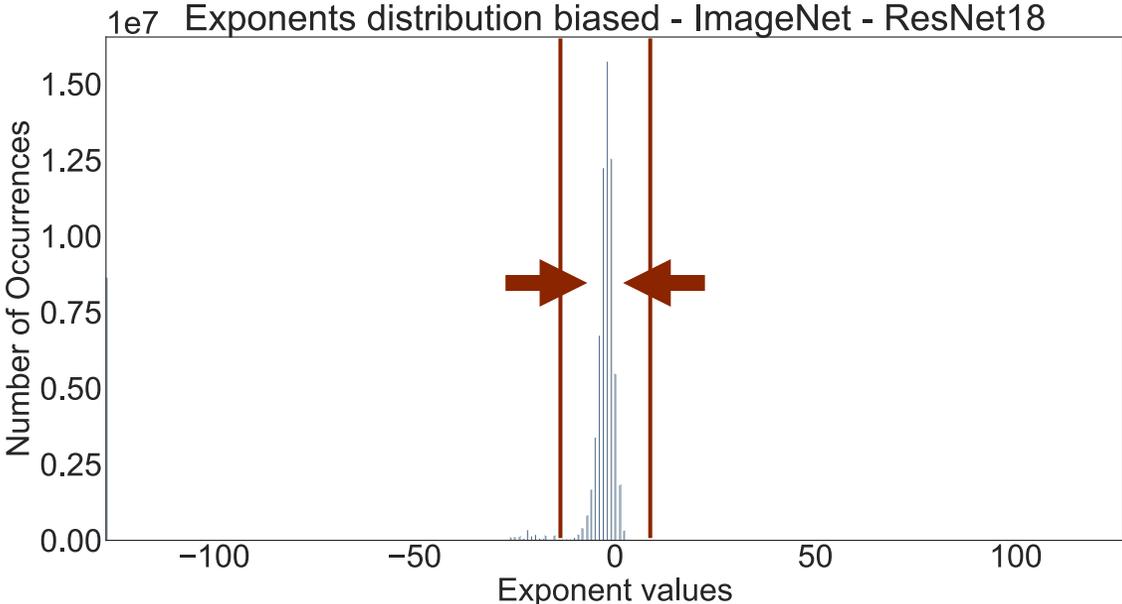
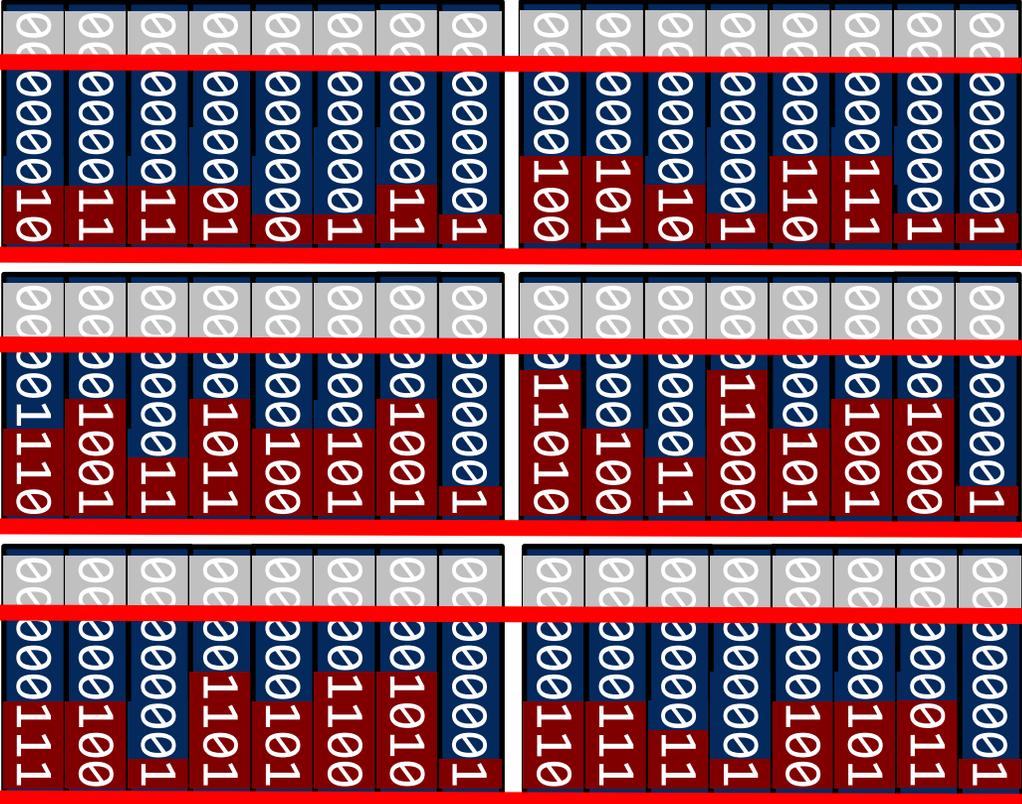
Exponent – *Gecko*

- Distribution
- Group



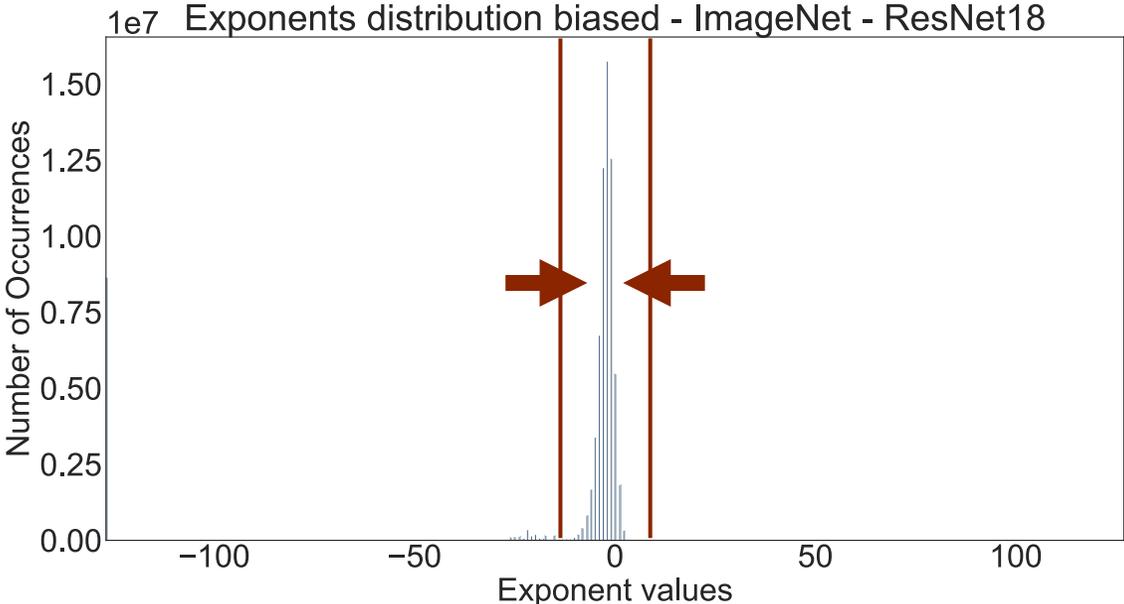
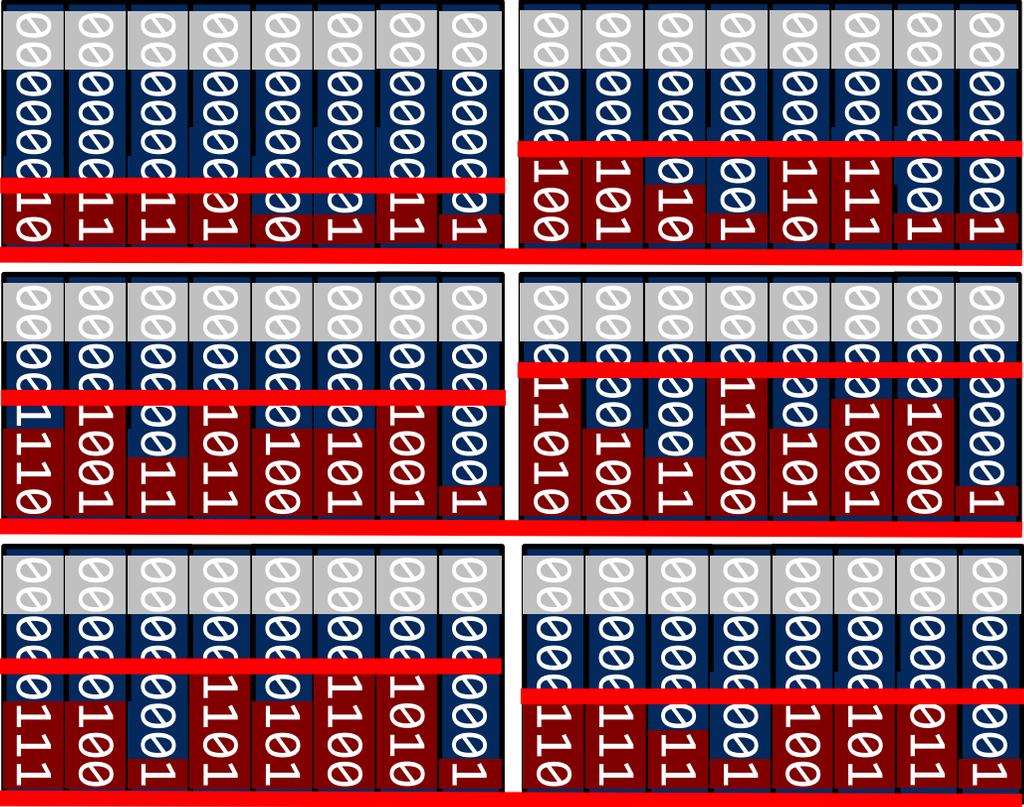
Exponent – *Gecko*

- Distribution
- Group

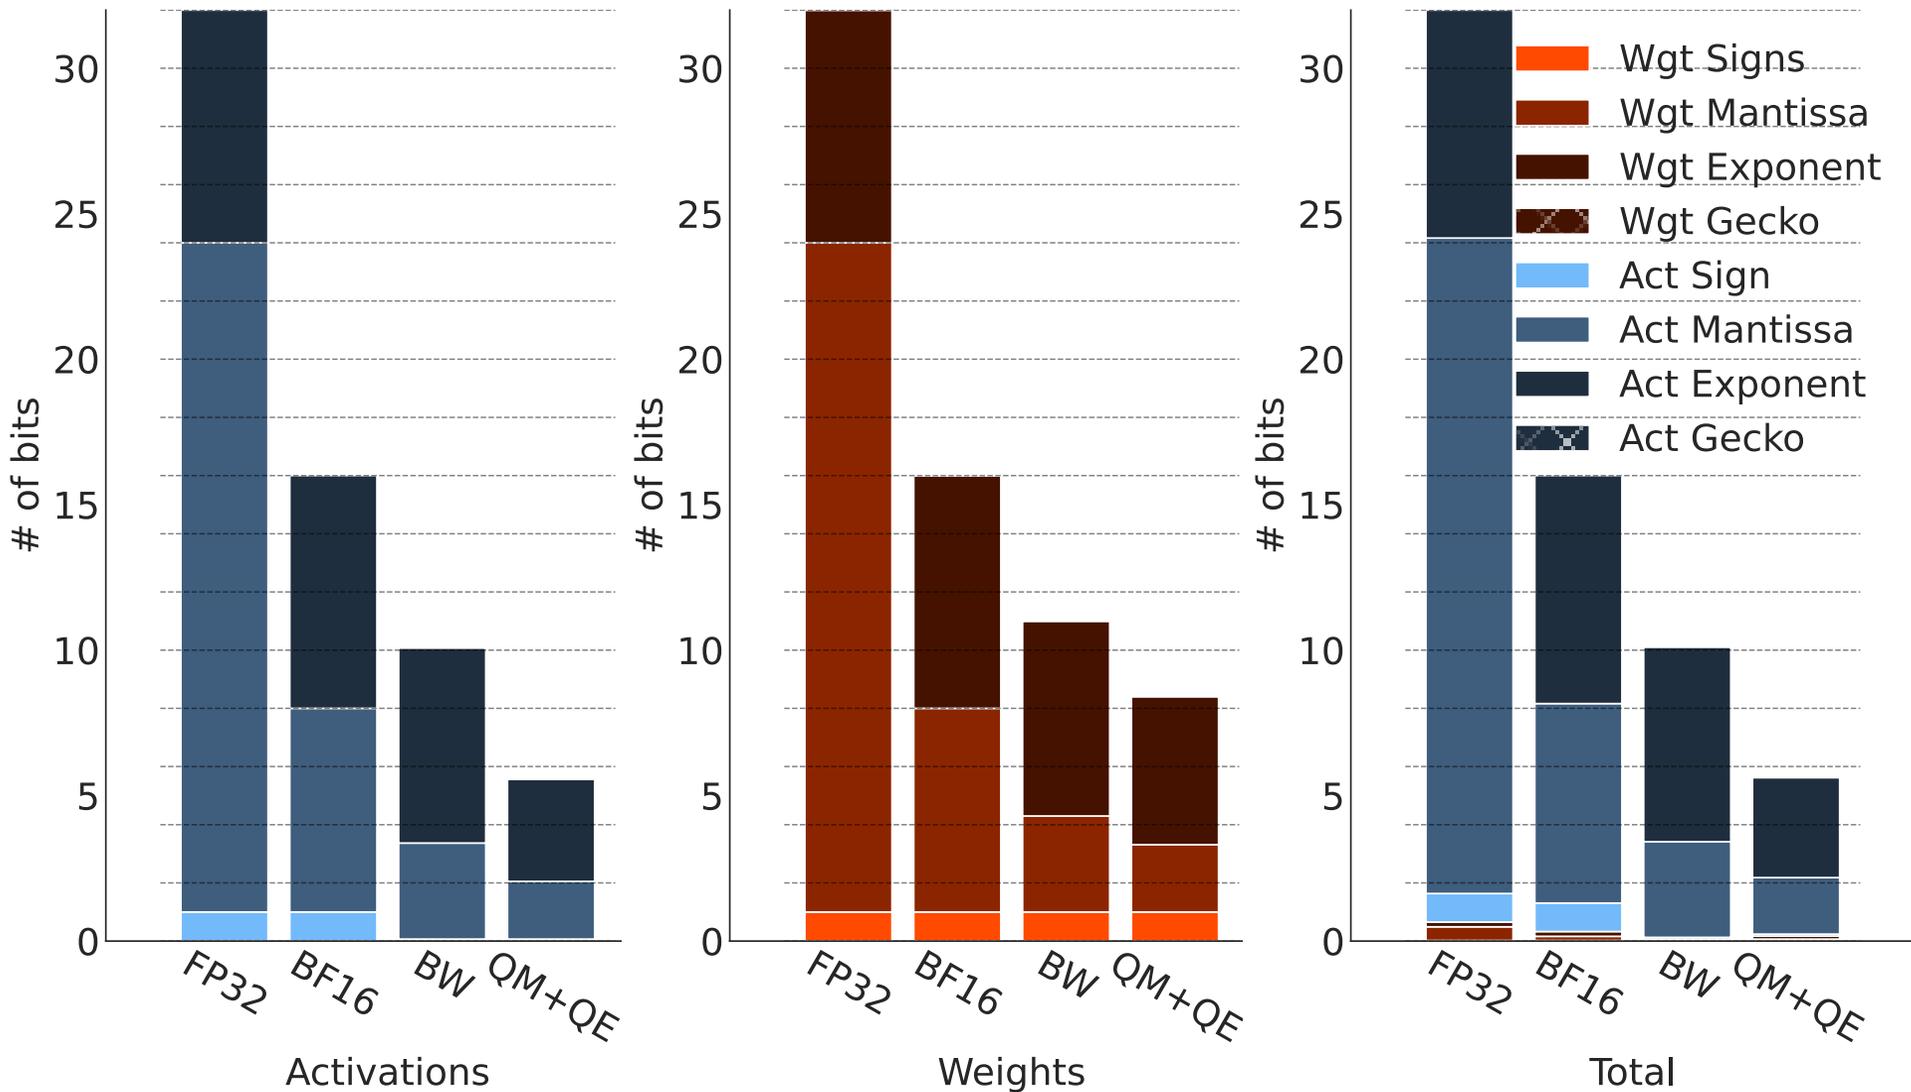


Exponent – *Gecko*

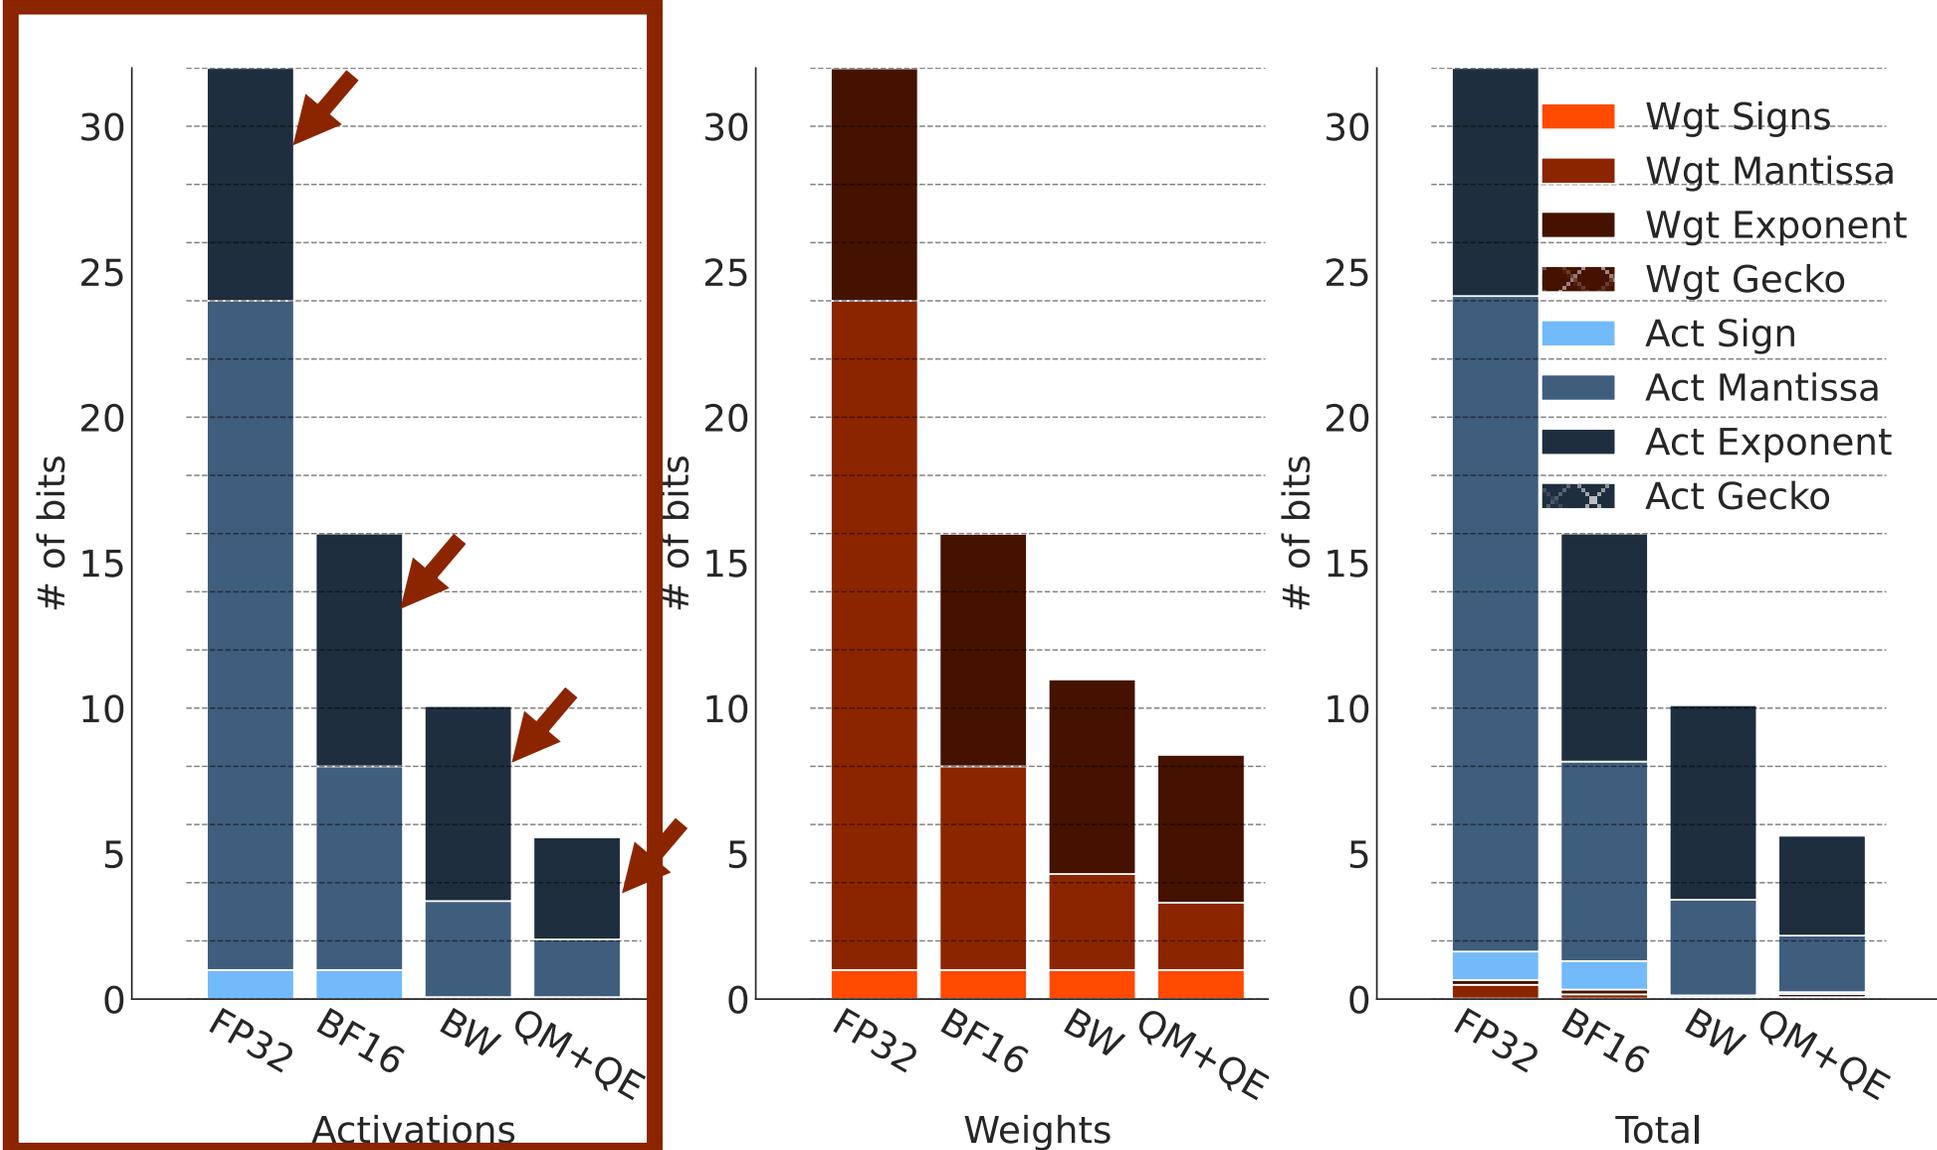
- Distribution
- Group



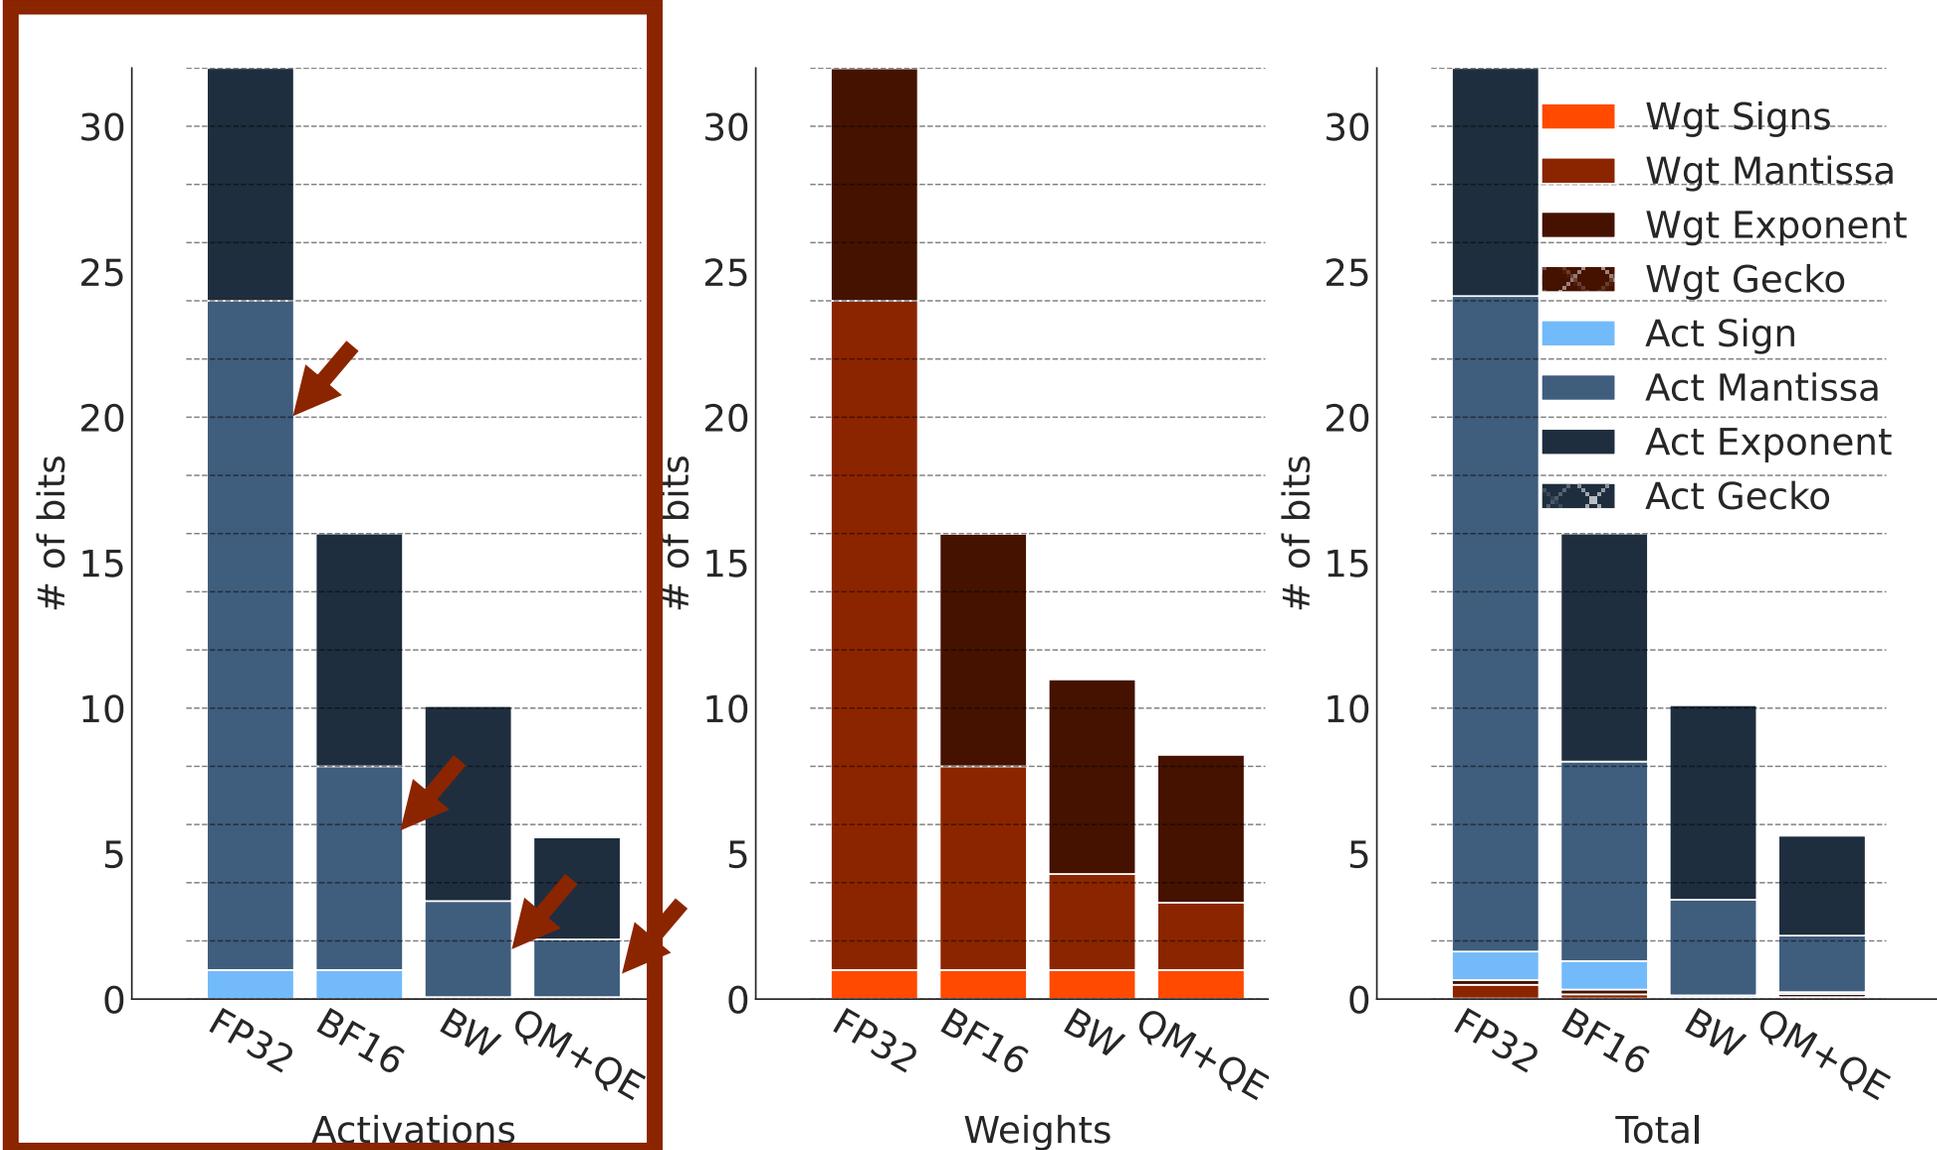
Massive footprint reduction on ResNet18!



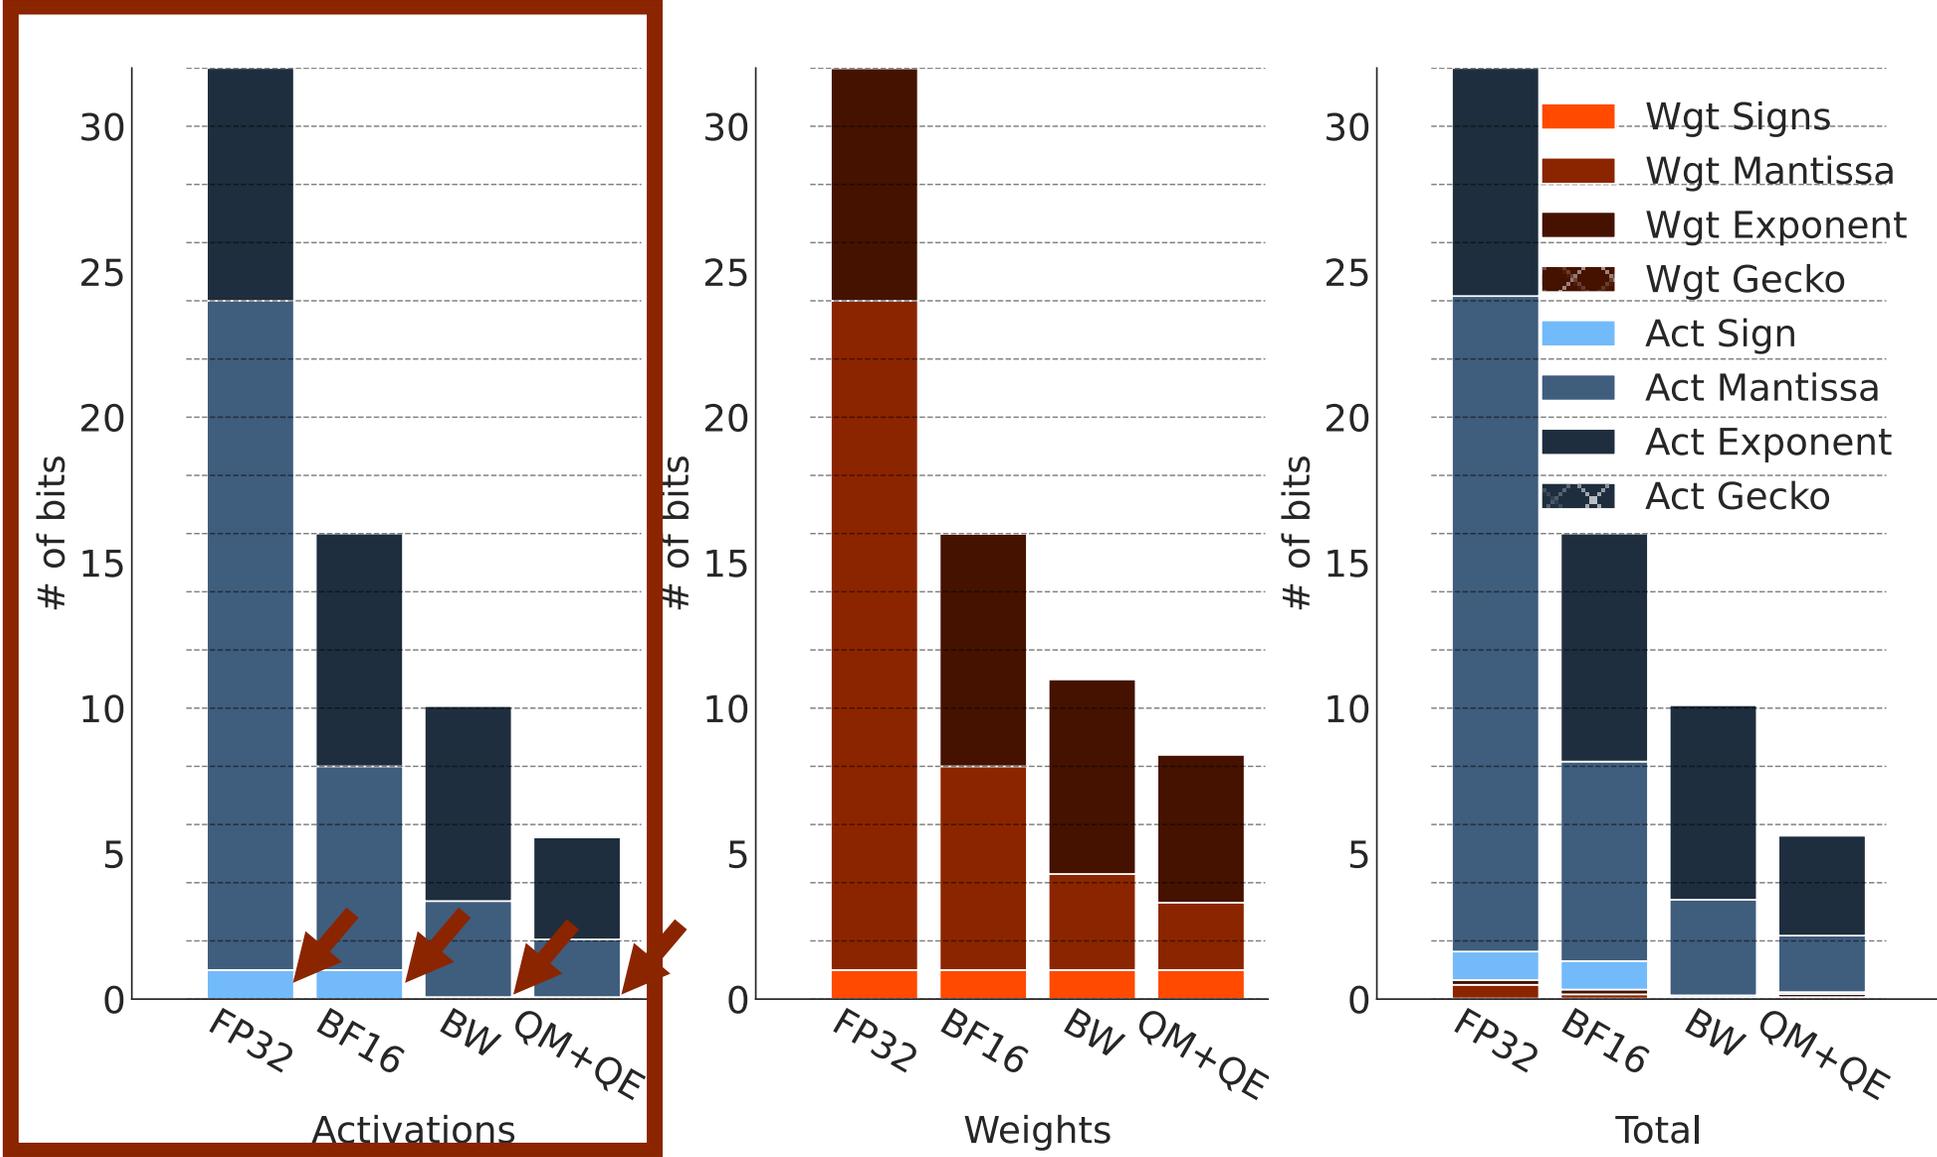
Massive footprint reduction on ResNet18!



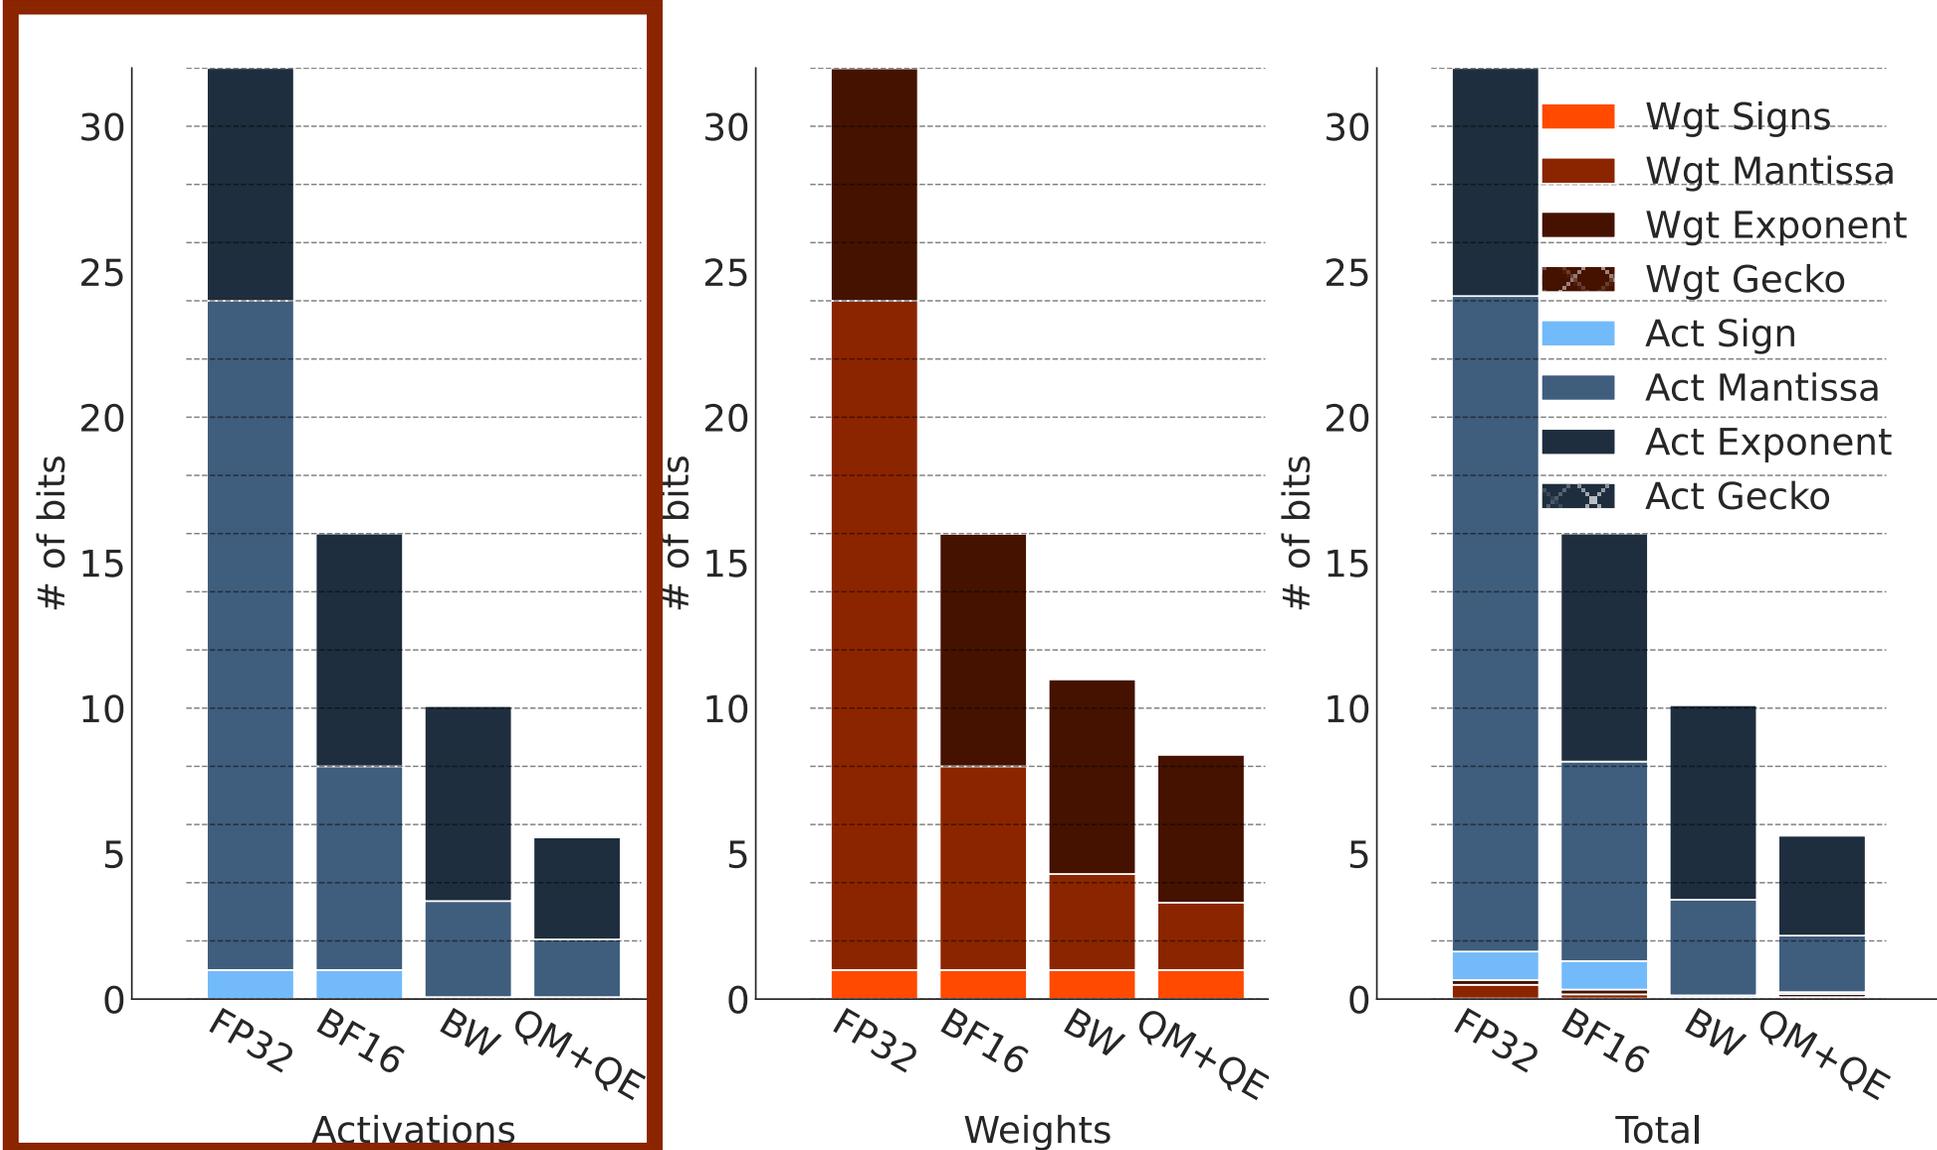
Massive footprint reduction on ResNet18!



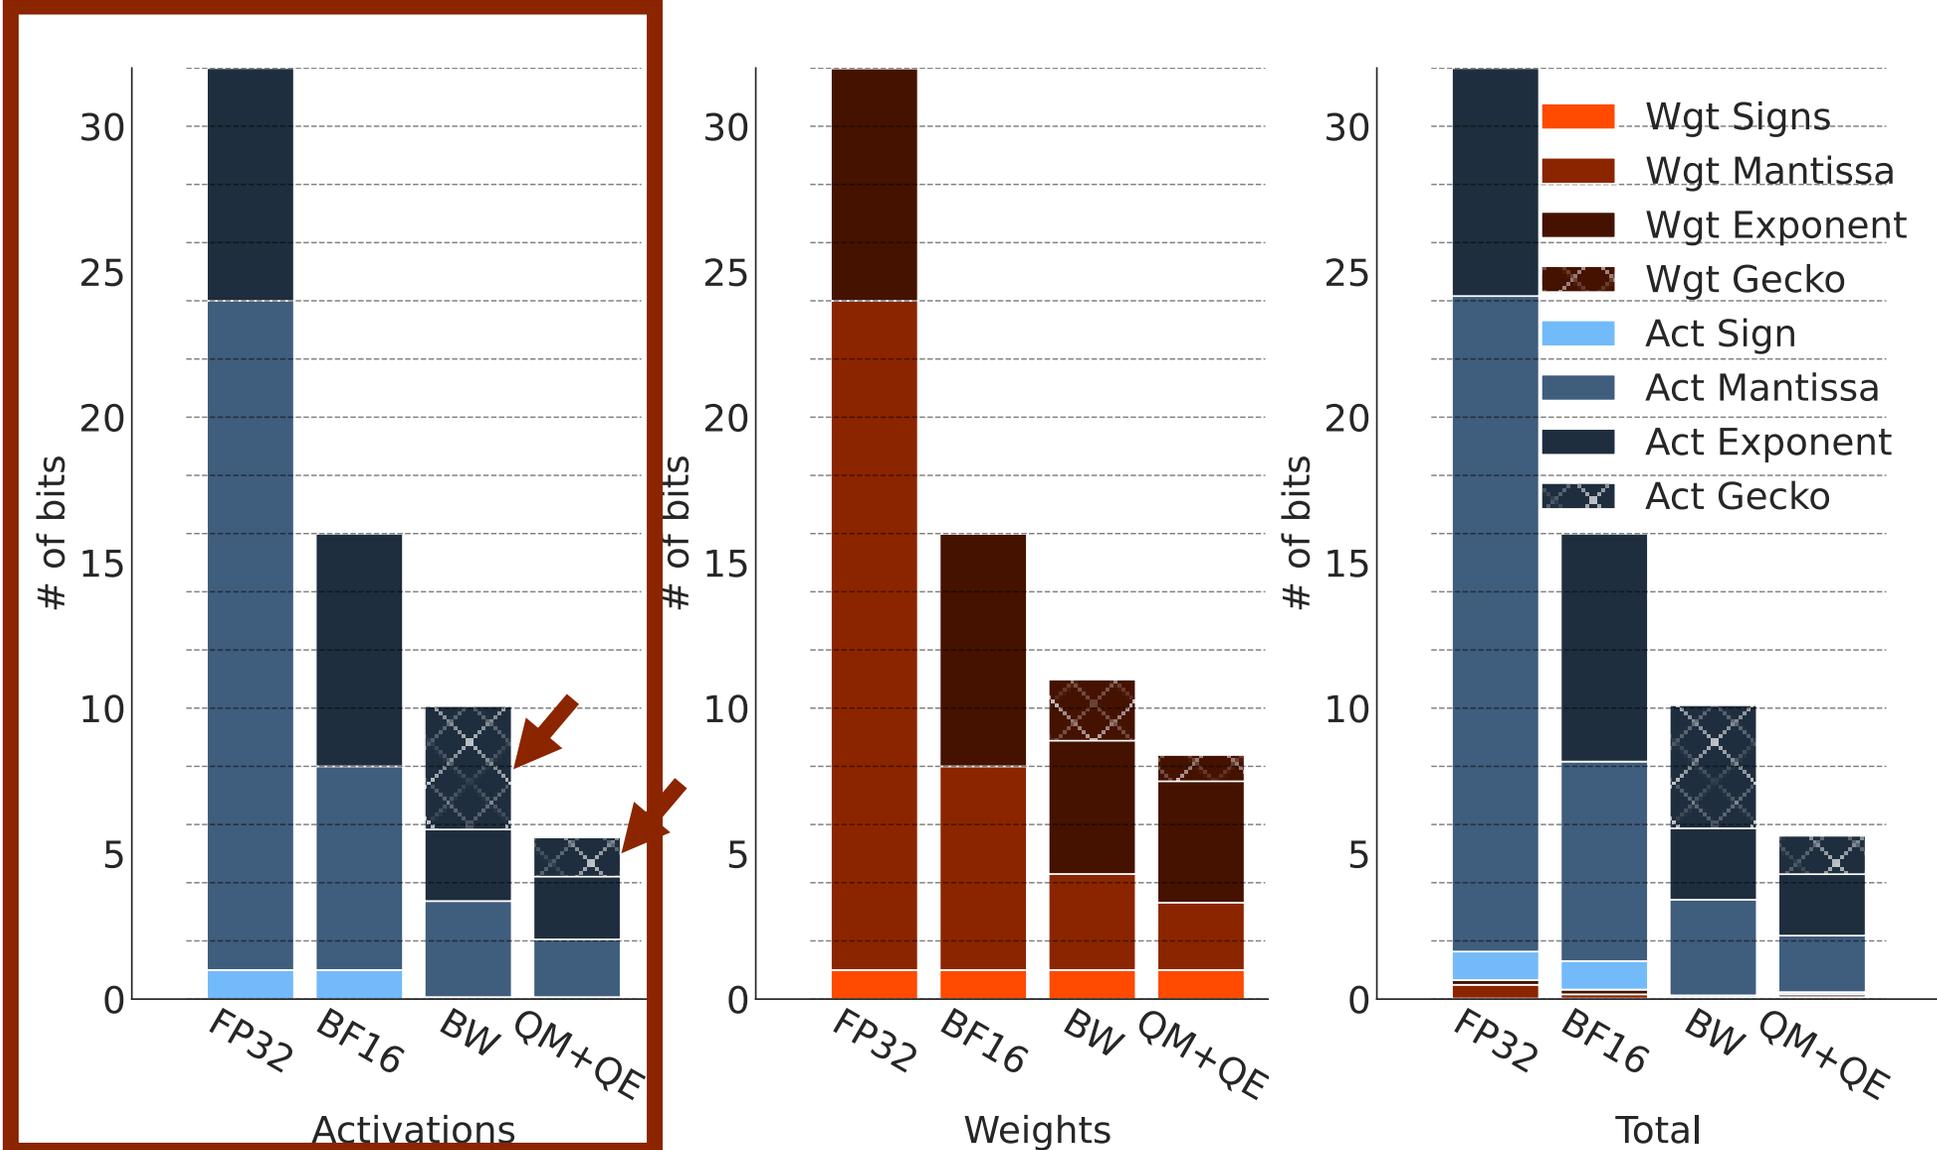
Massive footprint reduction on ResNet18!



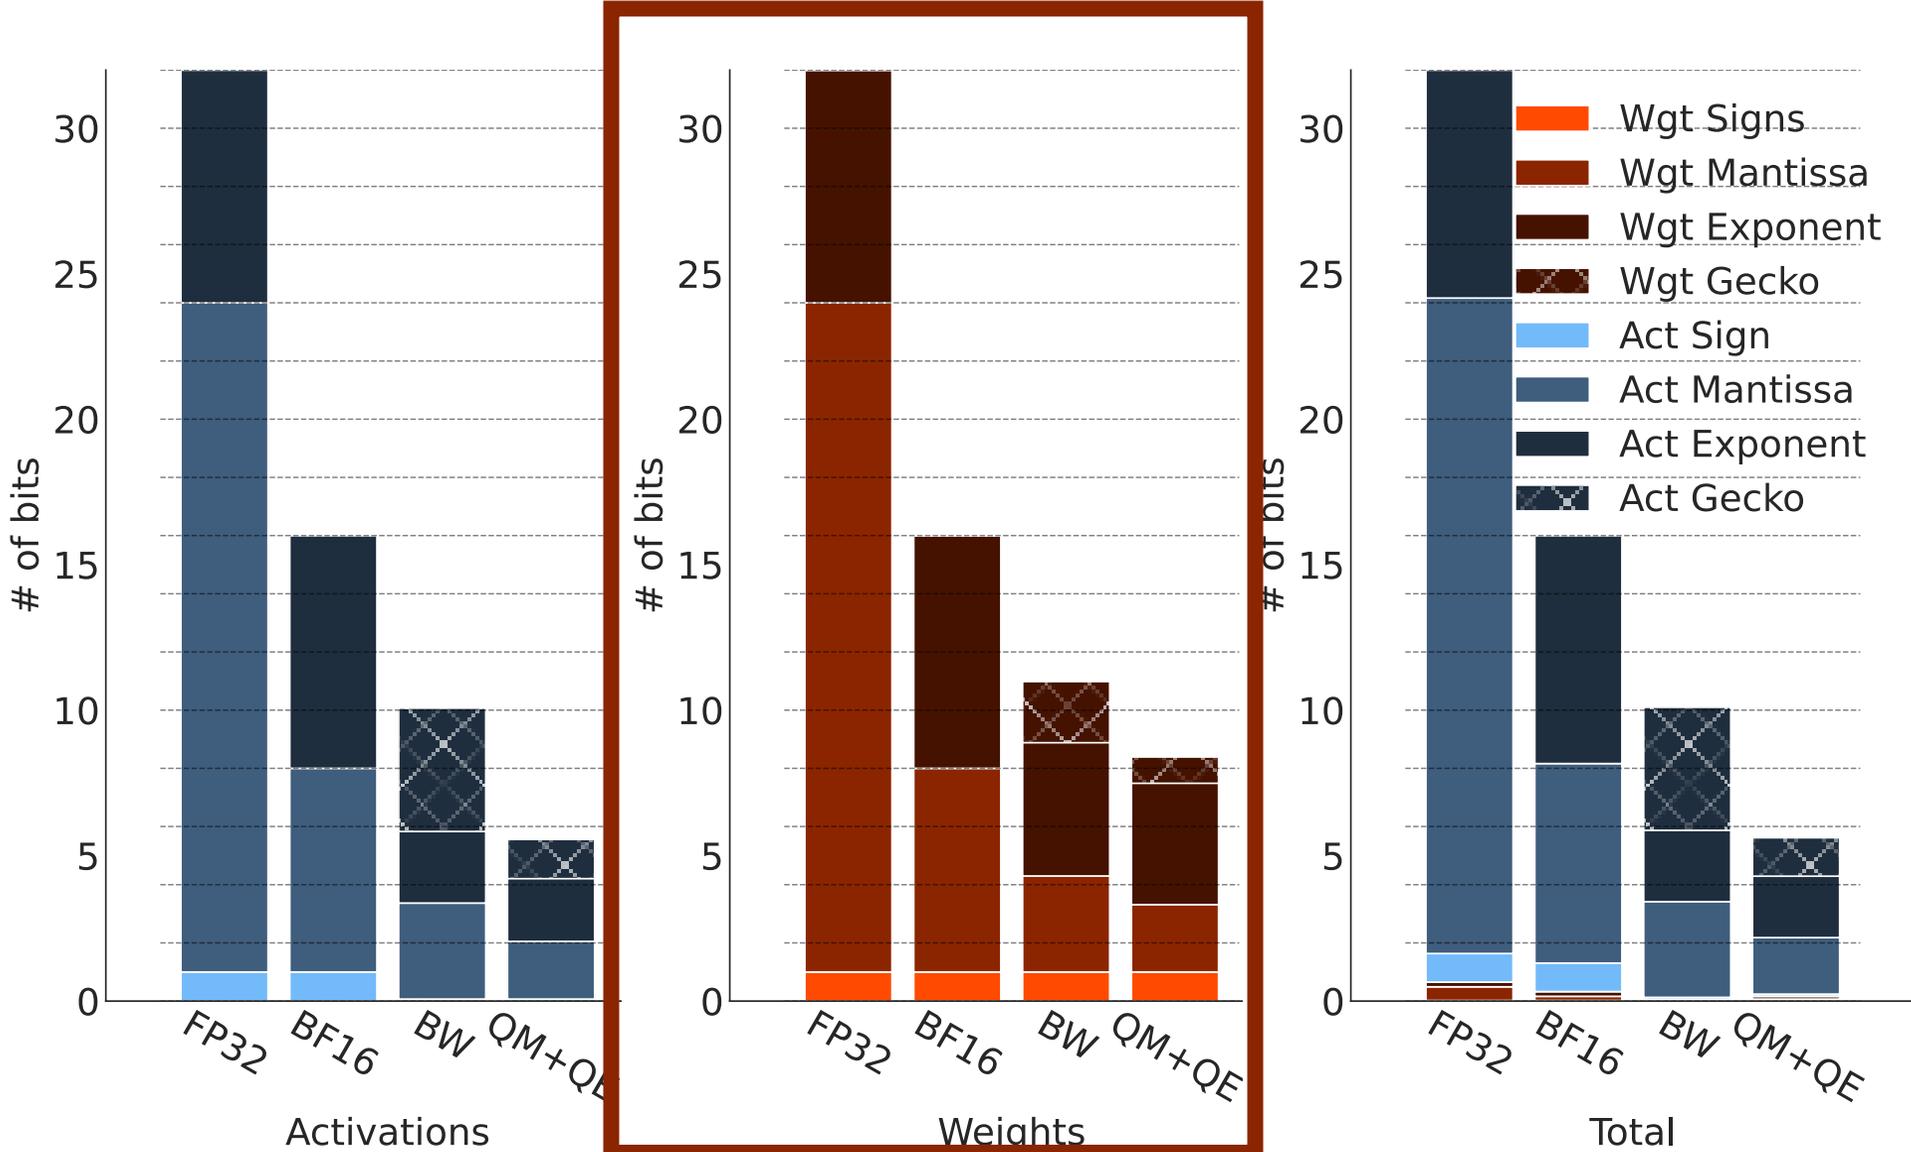
Massive footprint reduction on ResNet18!



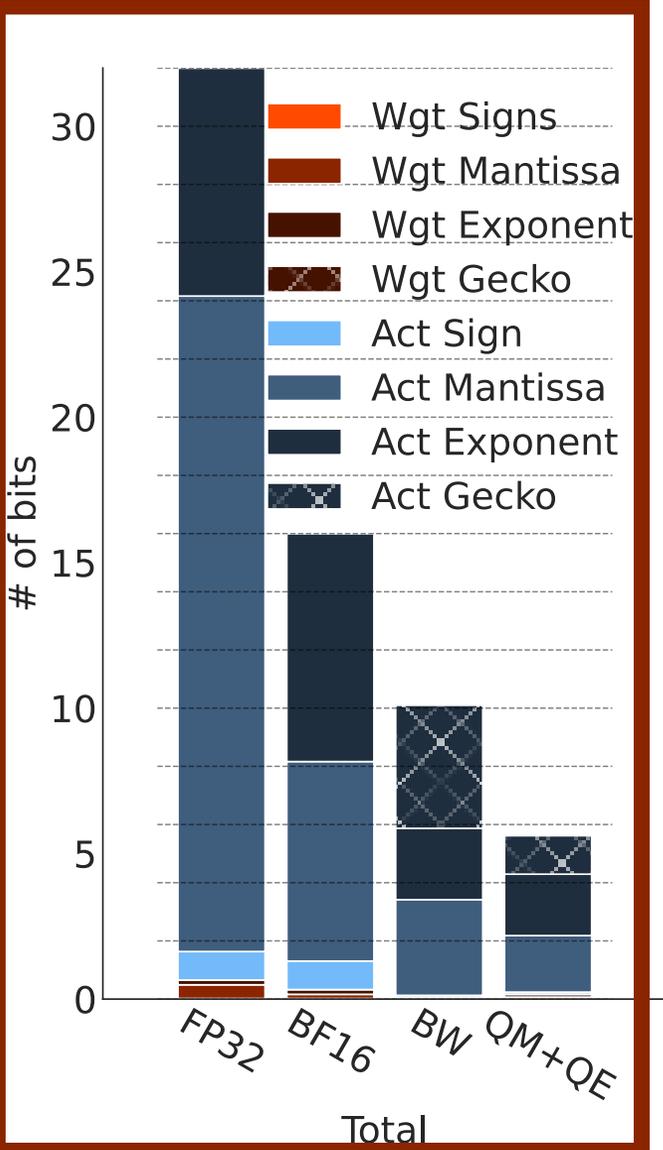
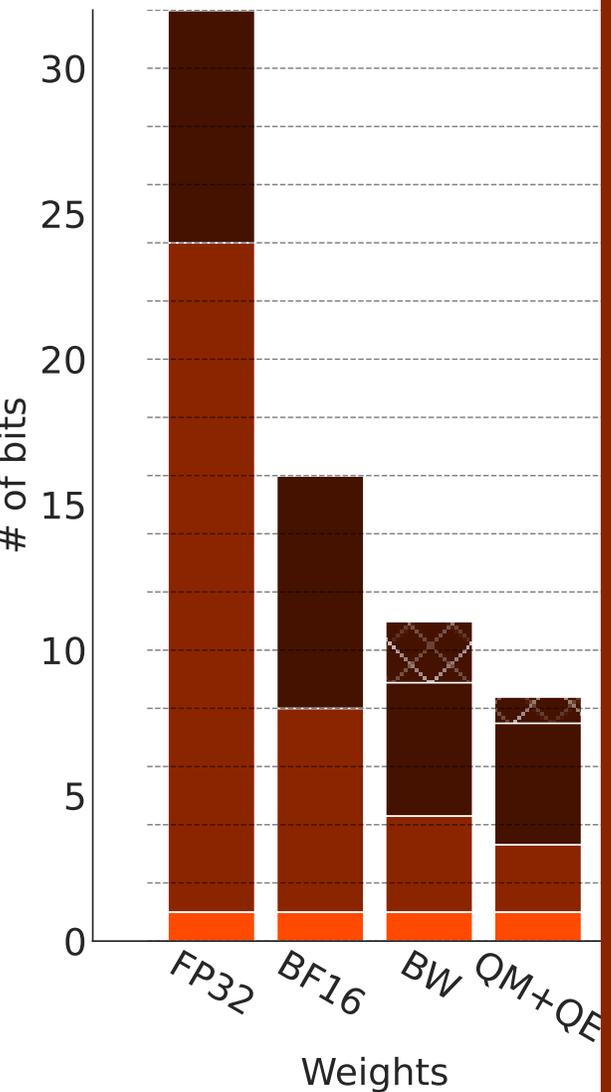
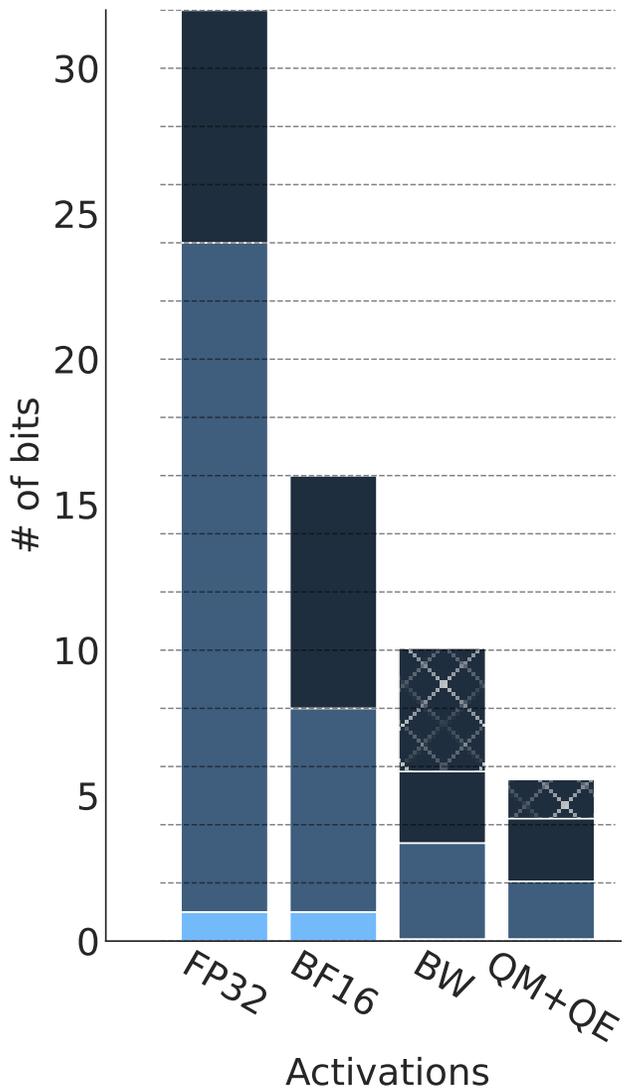
Massive footprint reduction on ResNet18!



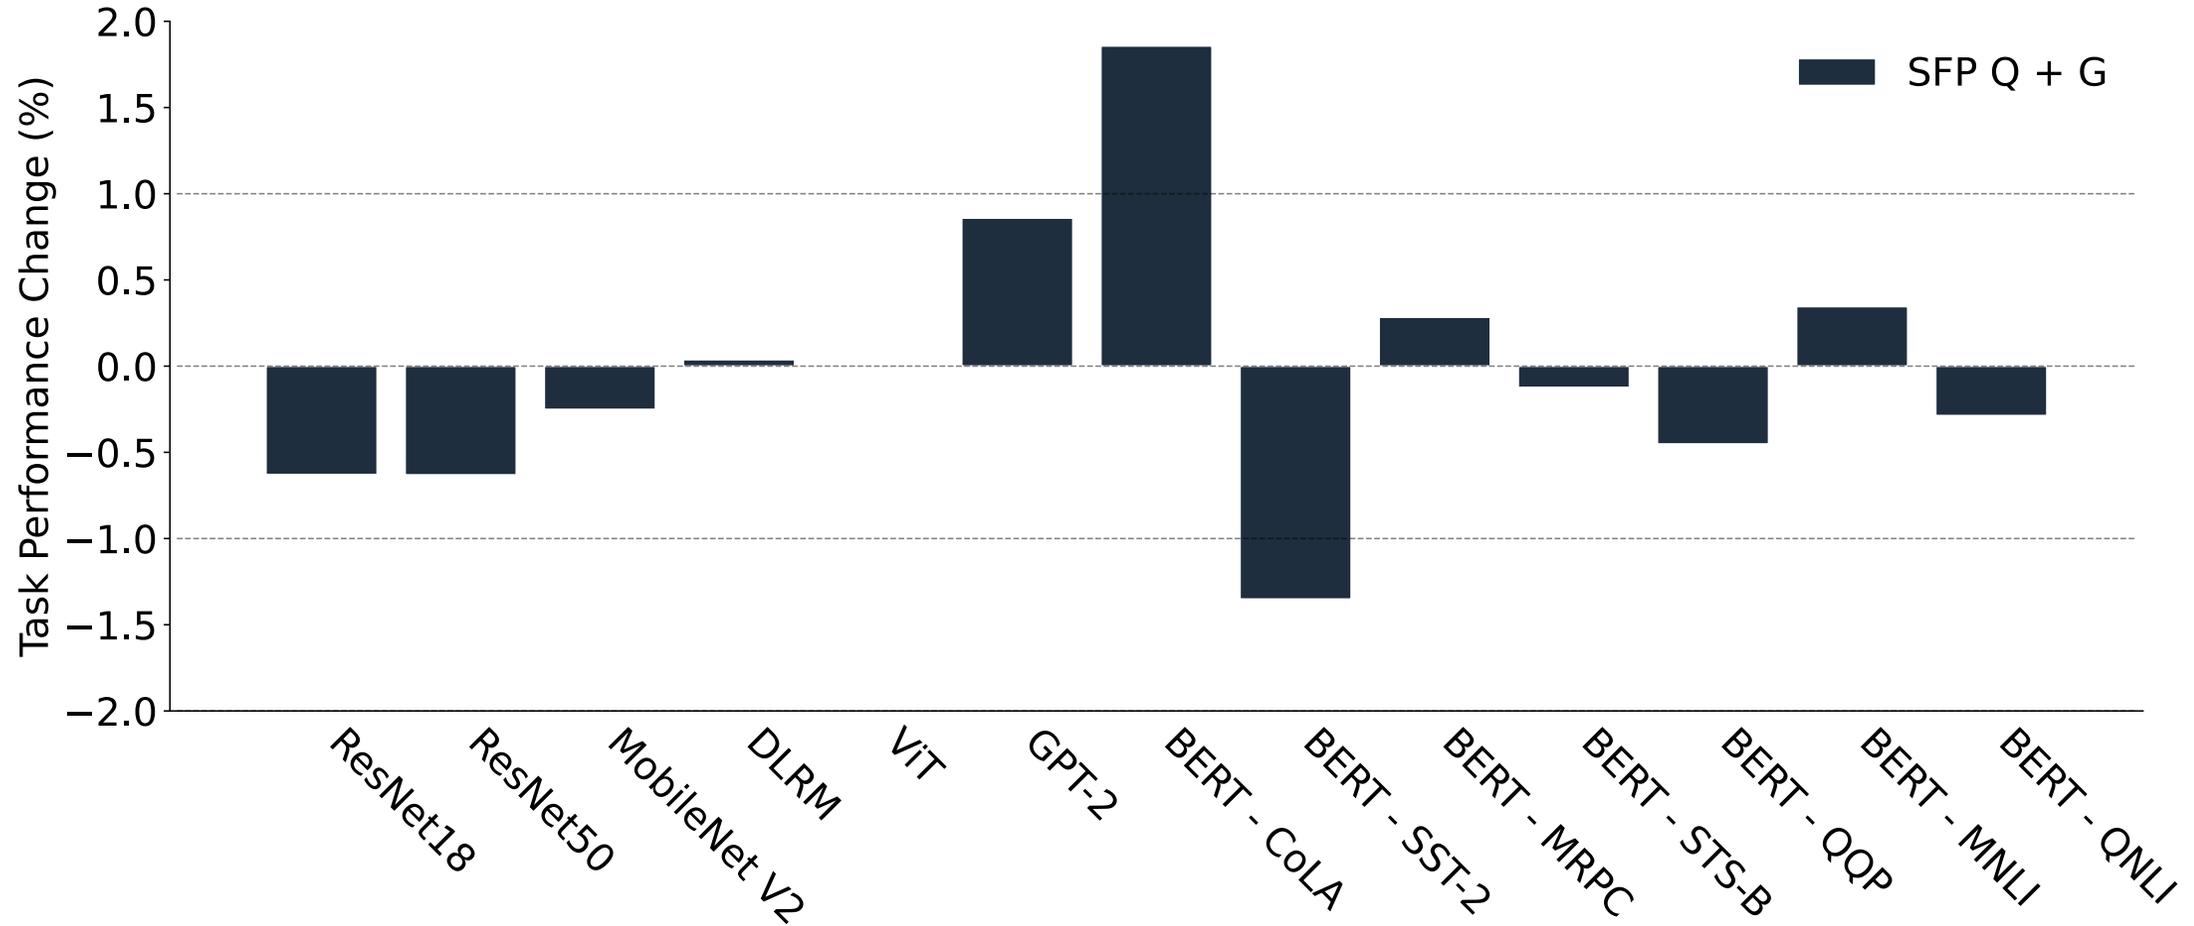
Massive footprint reduction on ResNet18!



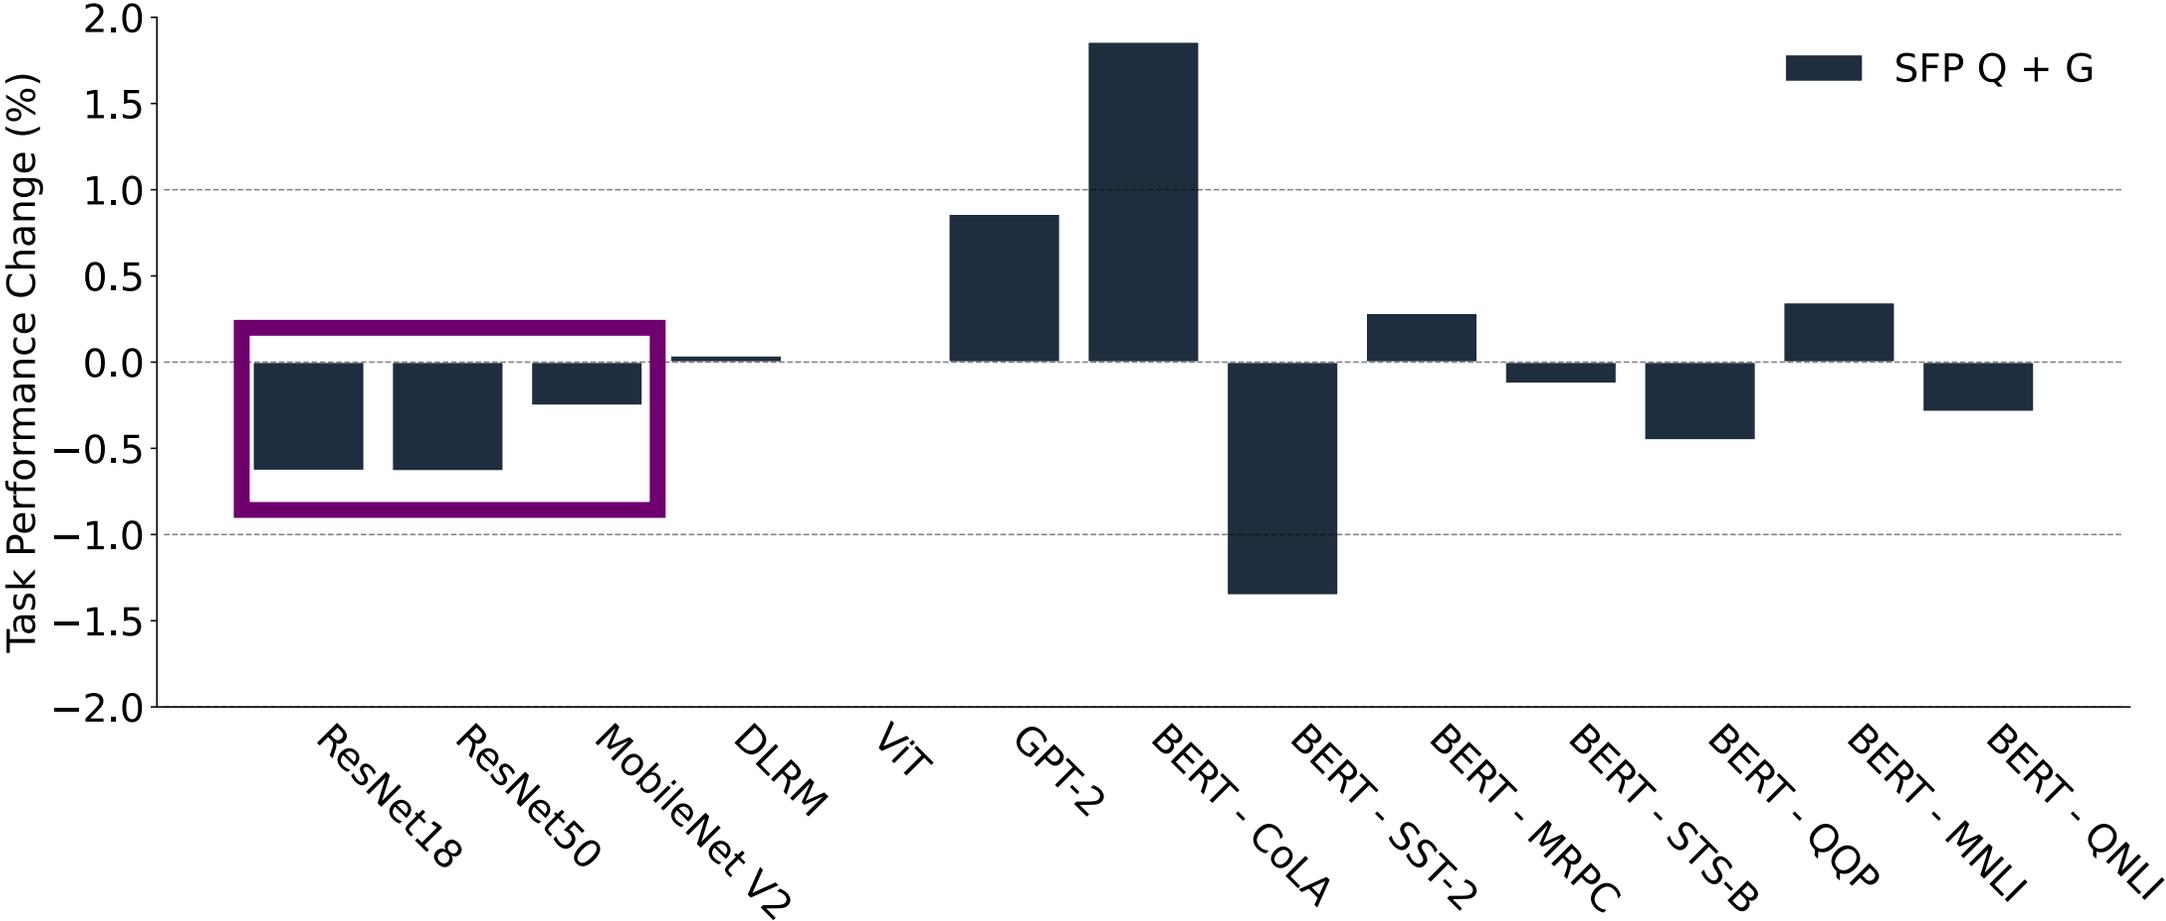
Massive footprint reduction on ResNet18!



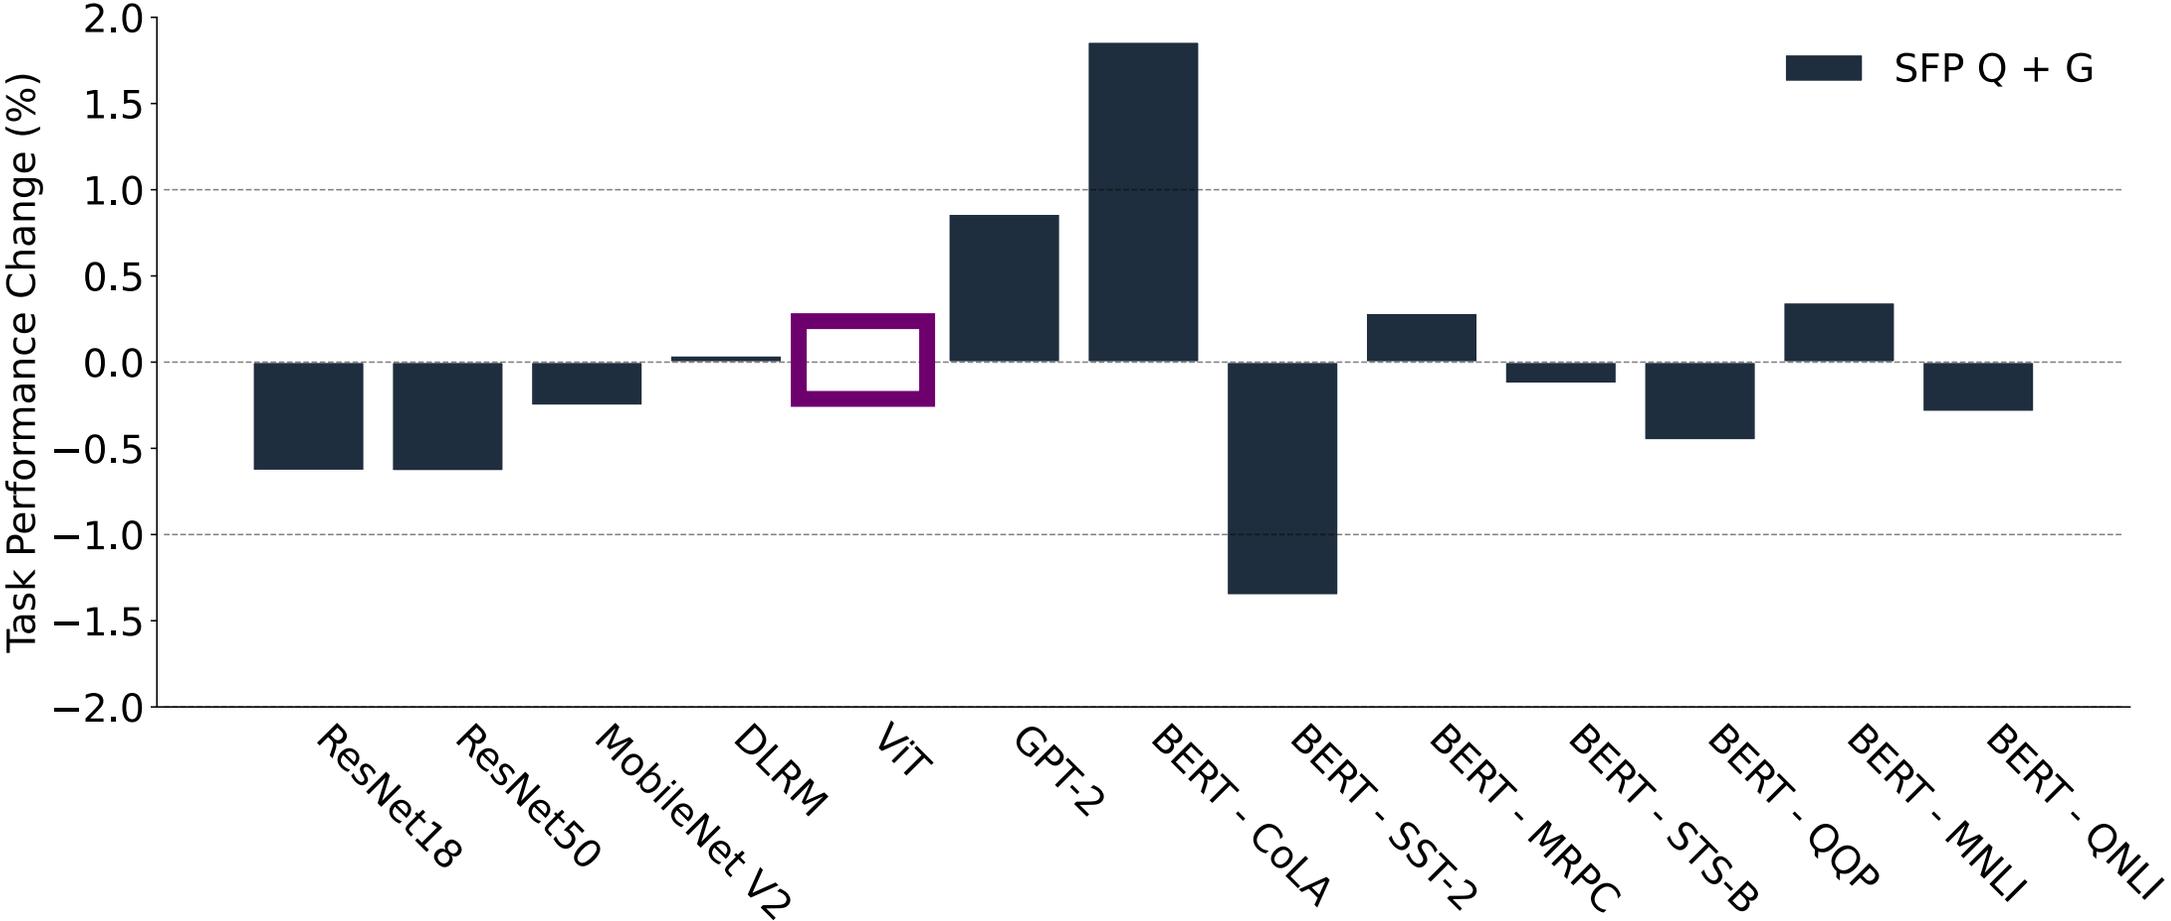
Similar Task Performance!



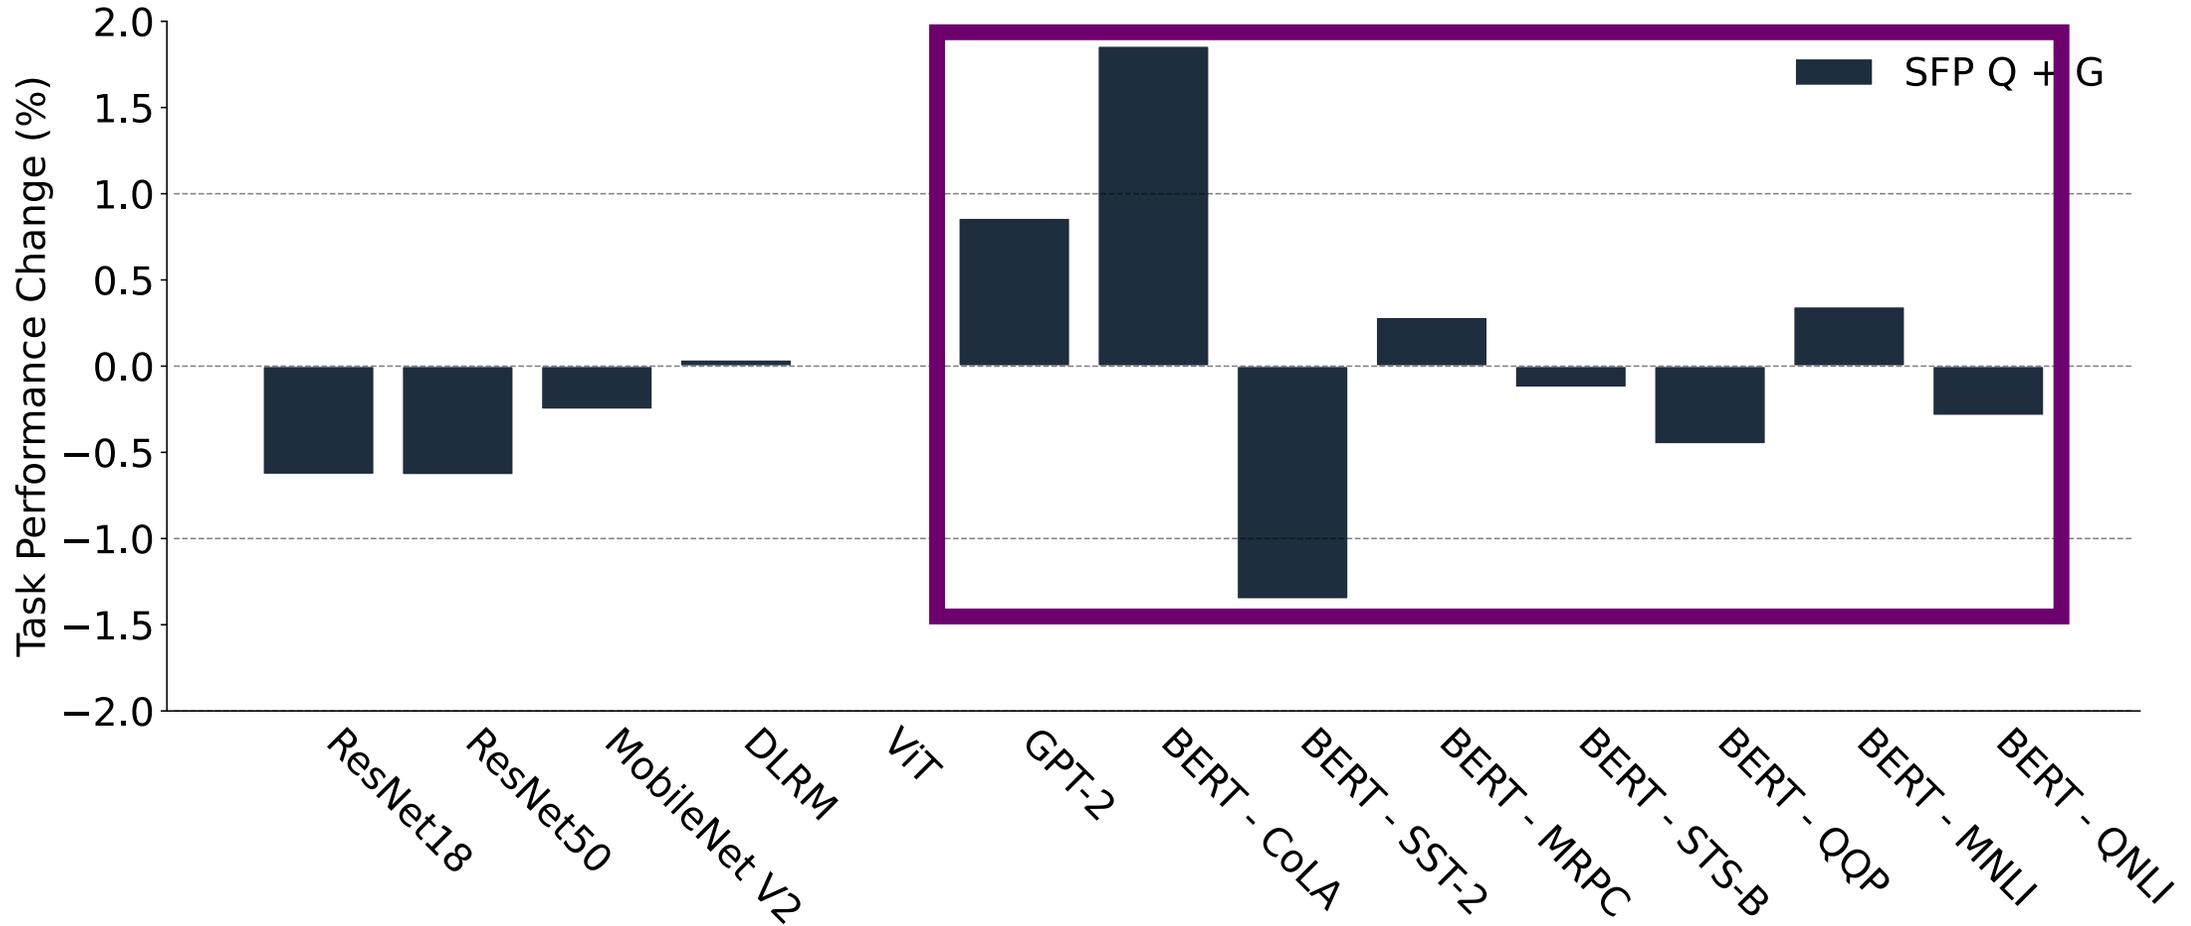
Similar Task Performance!



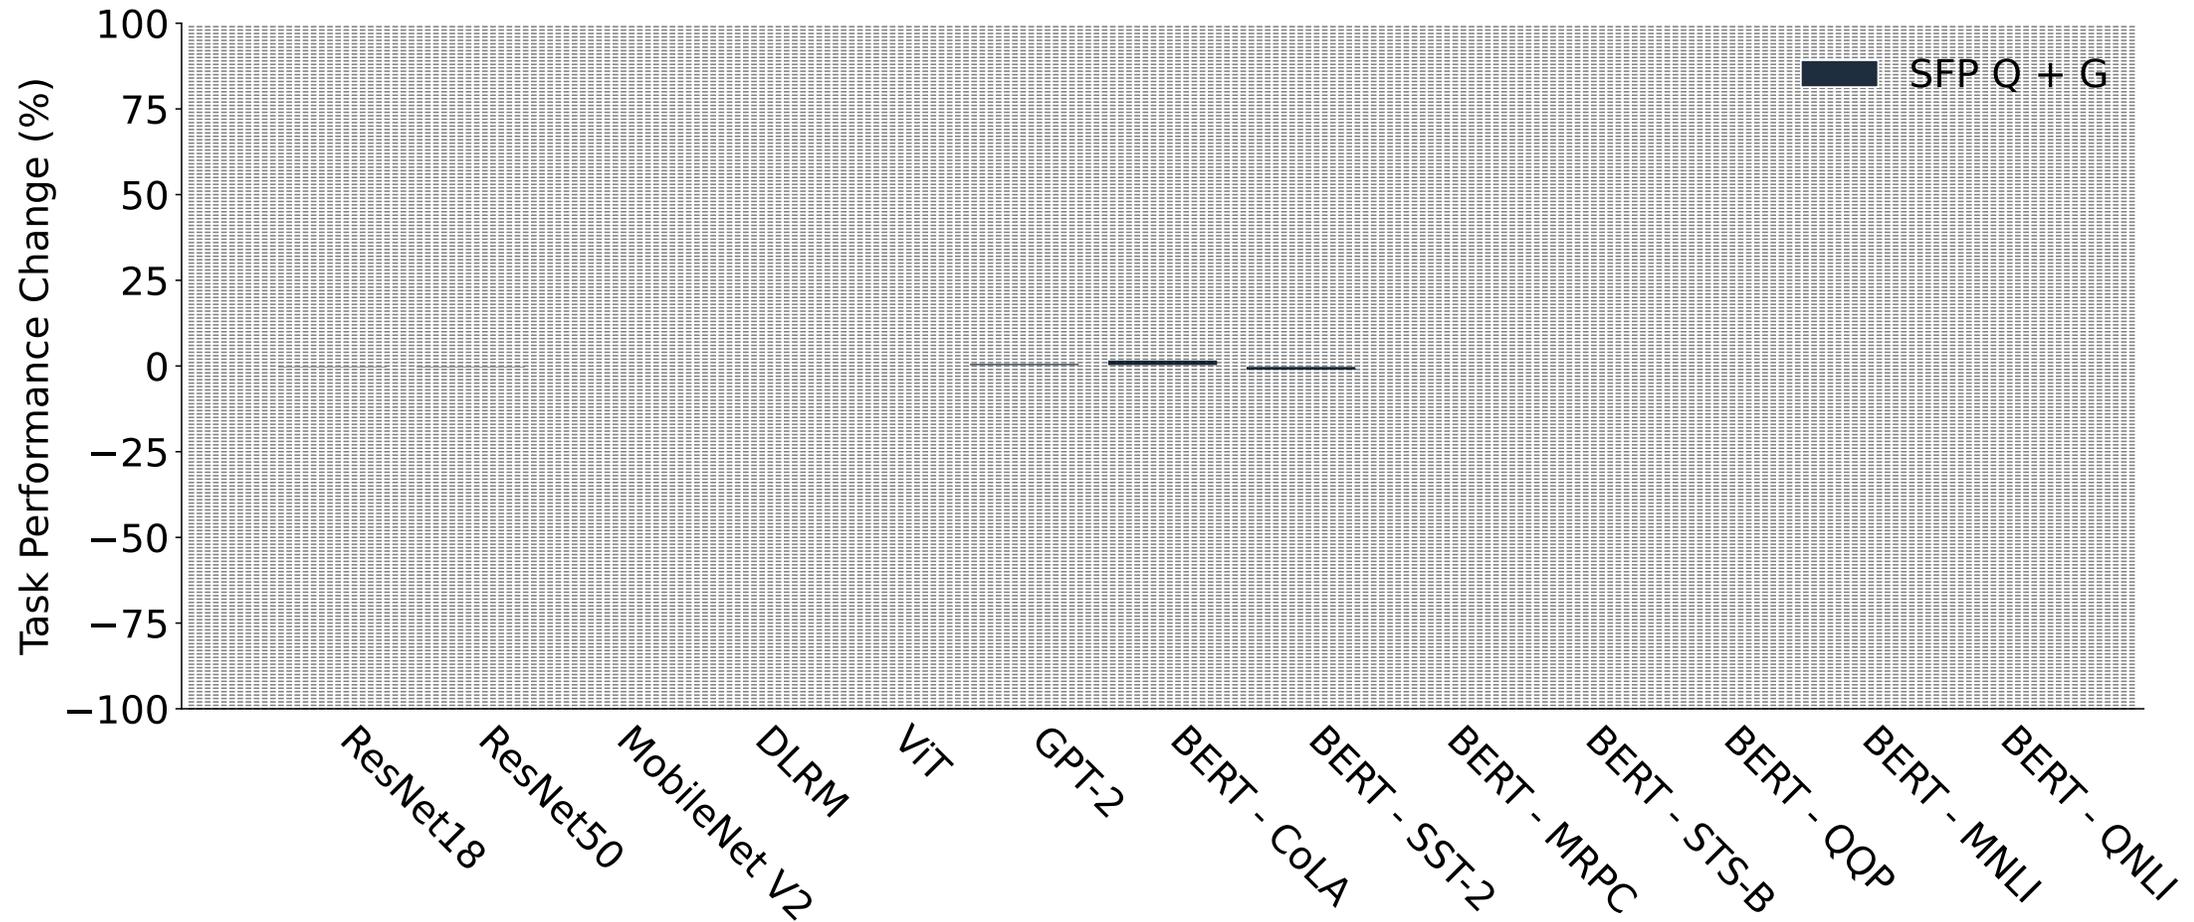
Similar Task Performance!



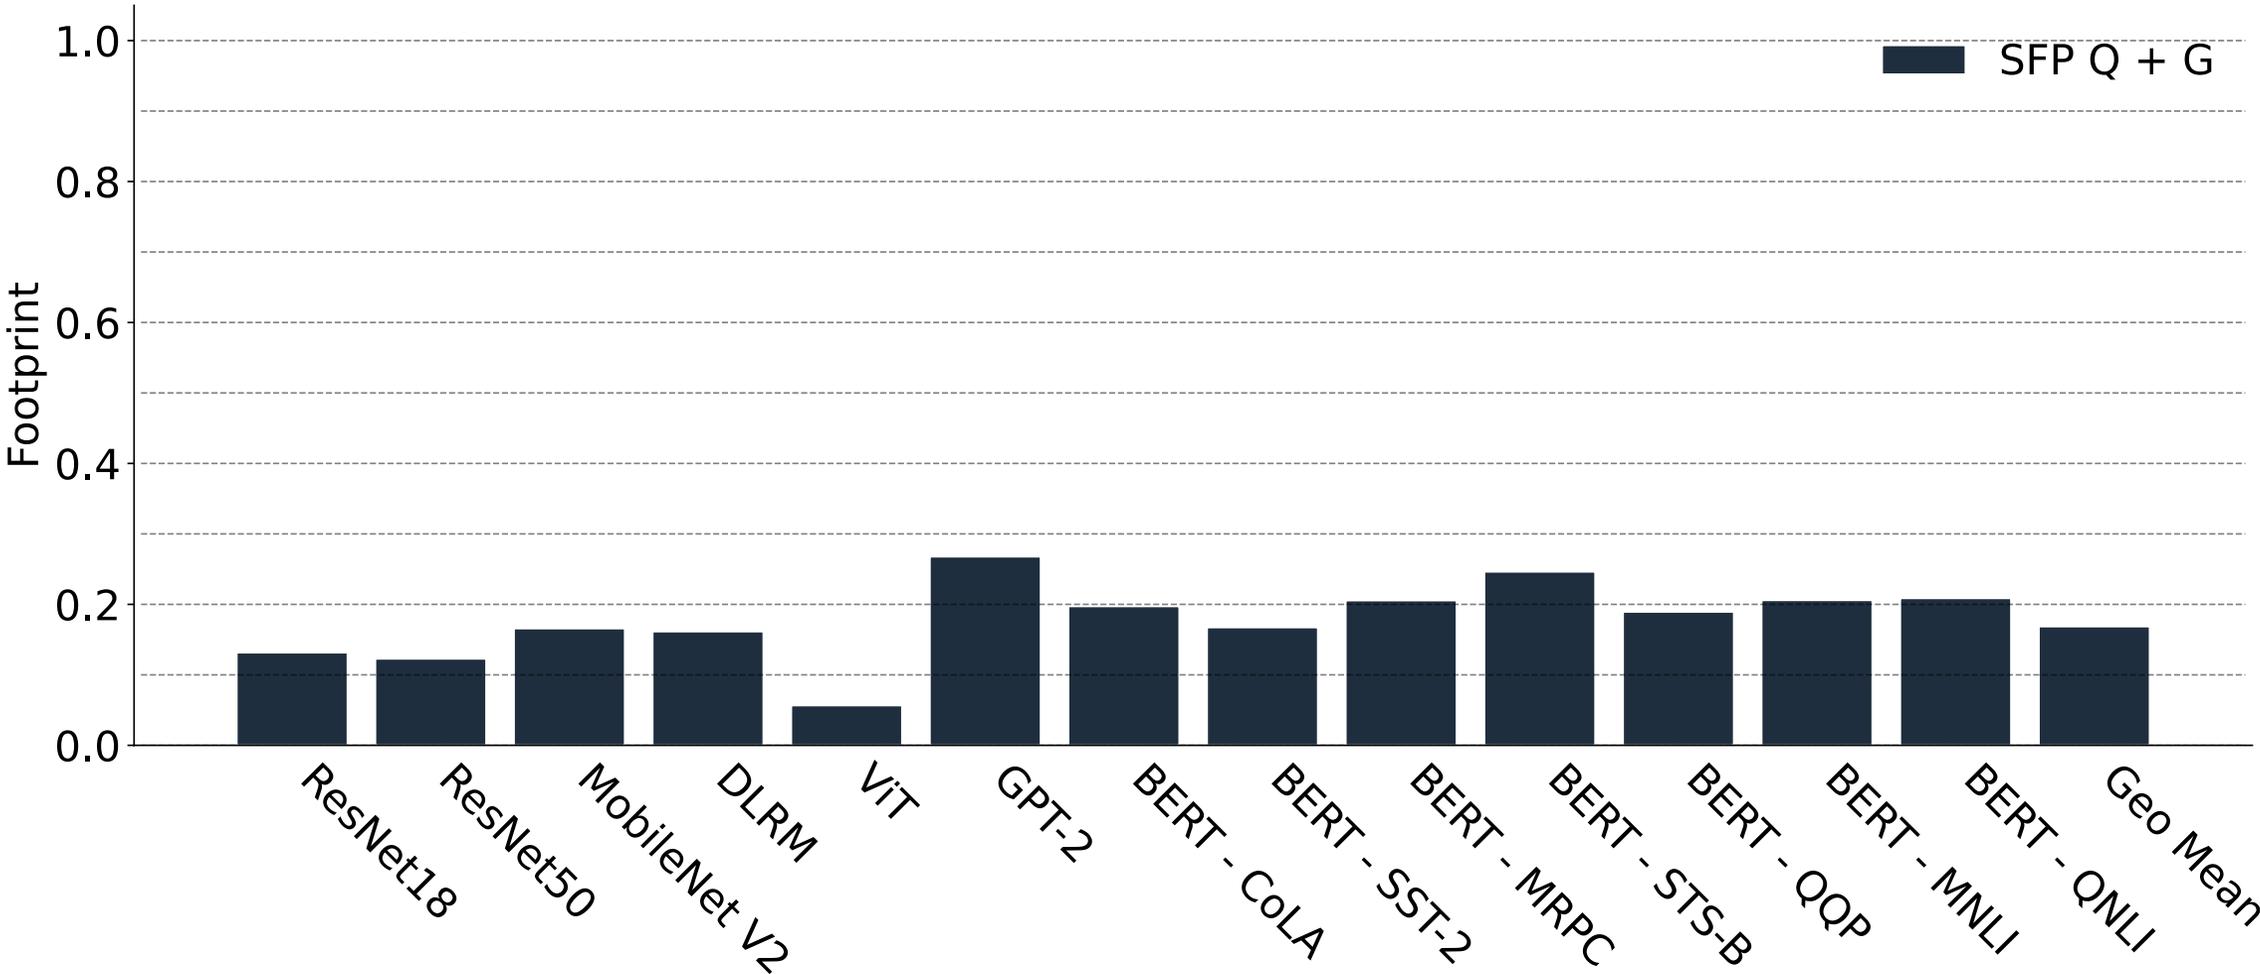
Similar Task Performance!



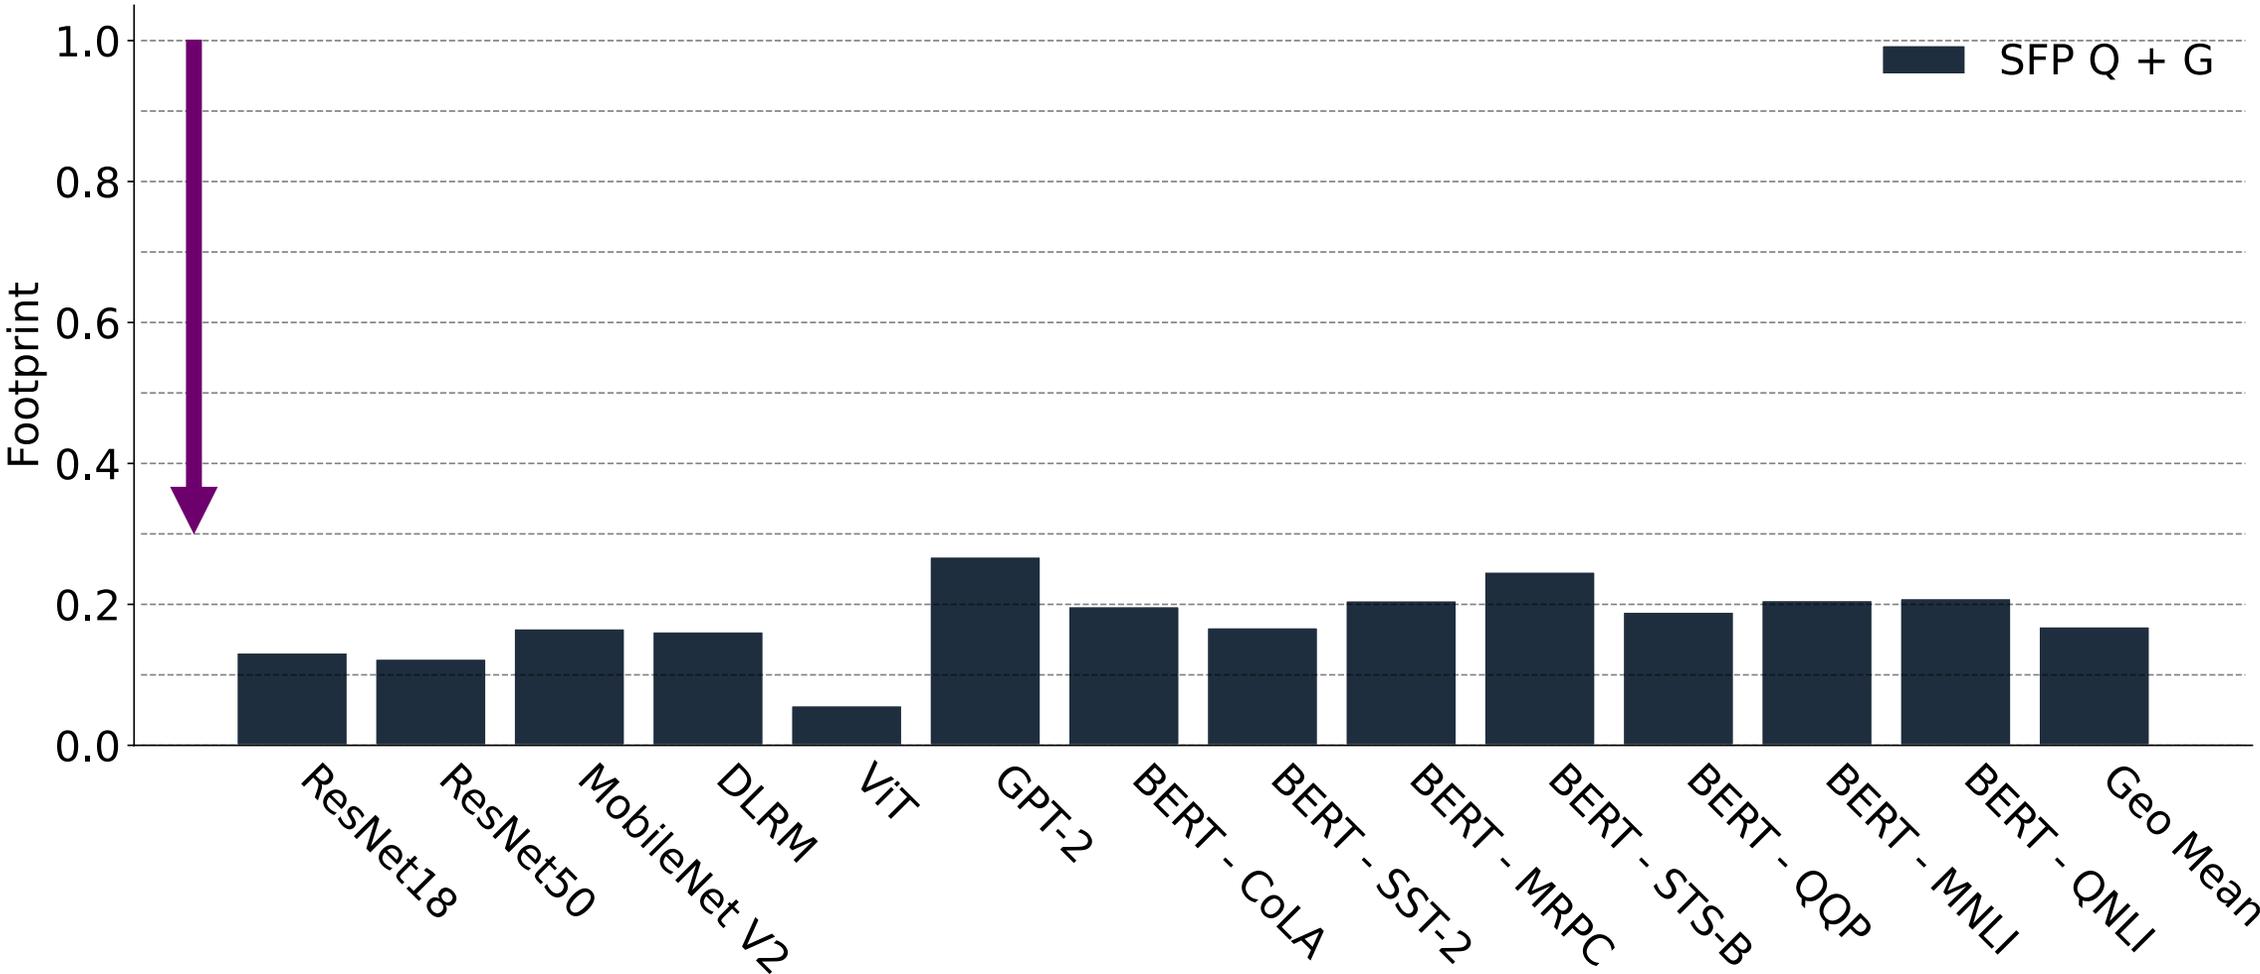
Similar Task Performance!



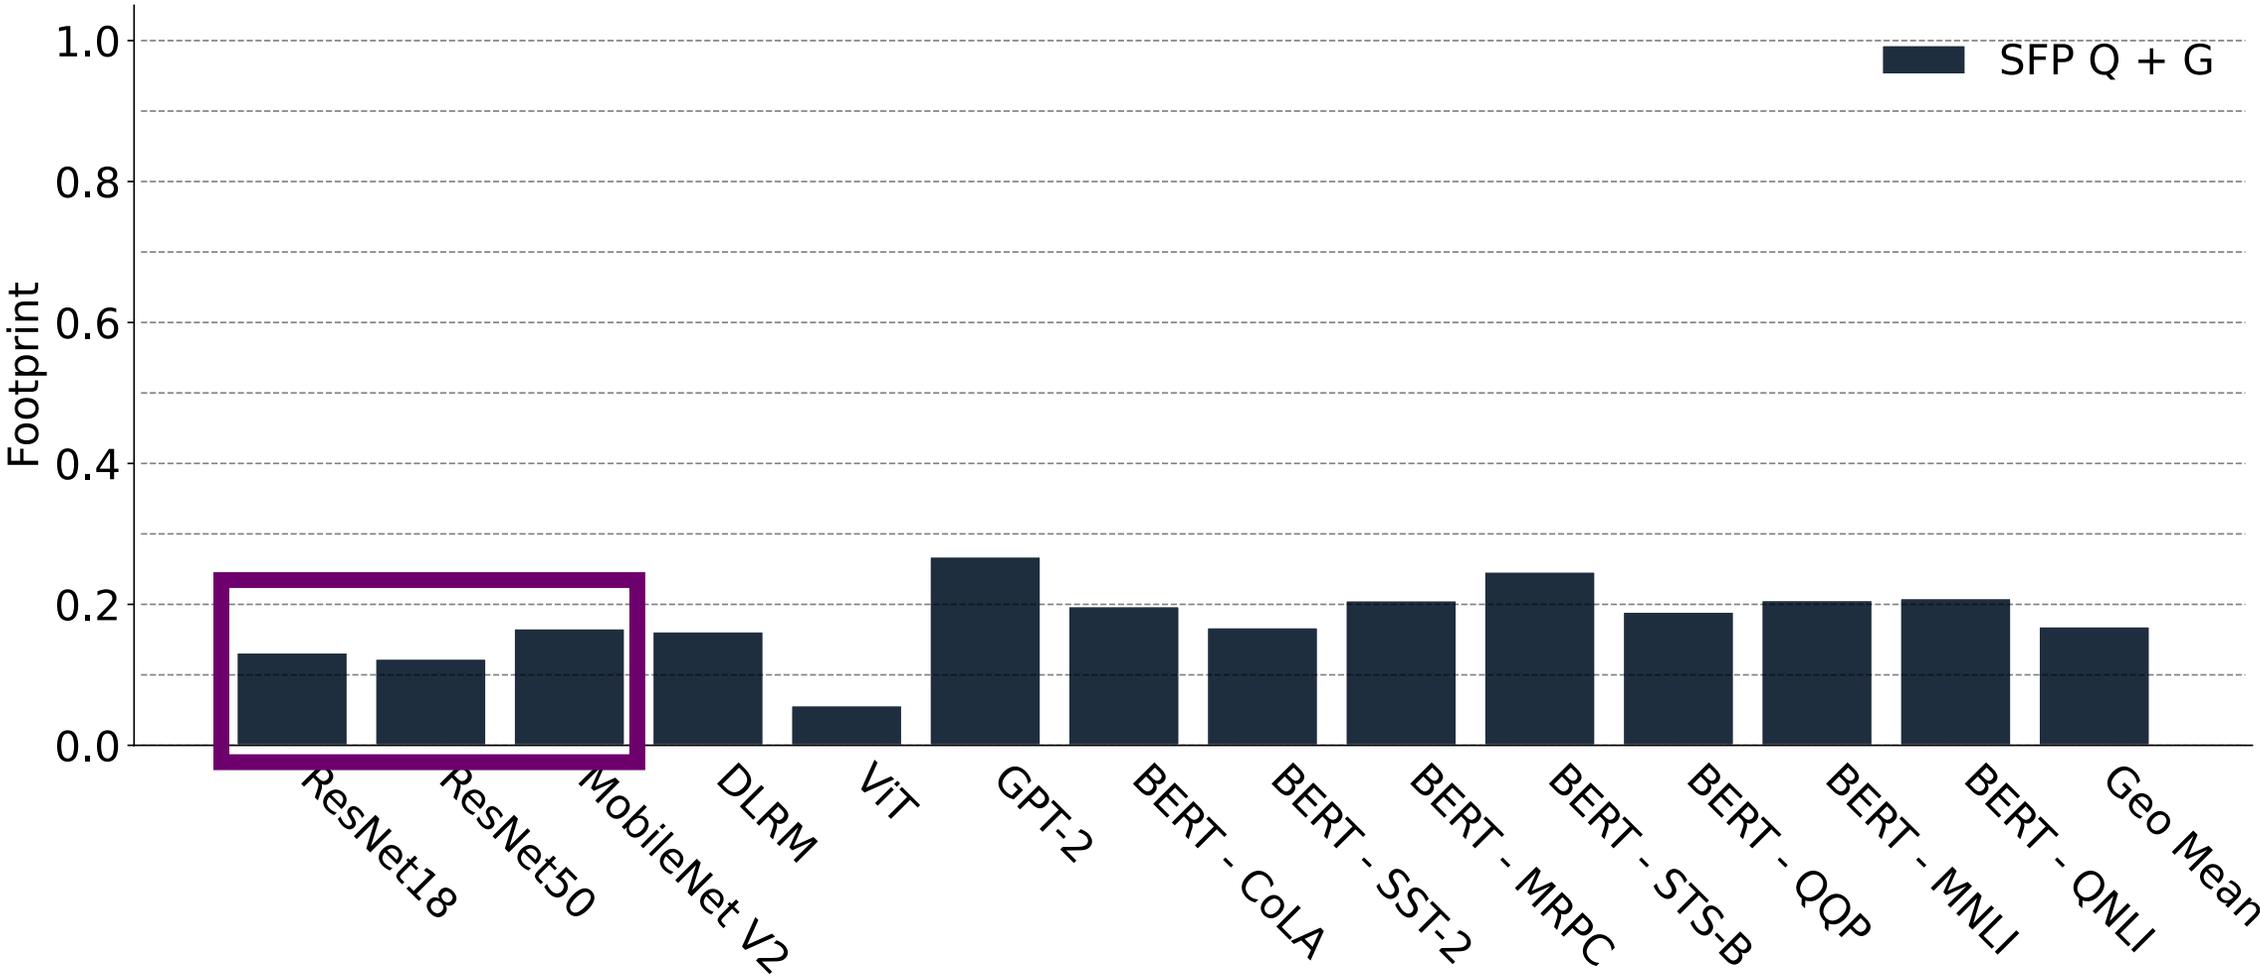
Massive footprint reduction across the board!



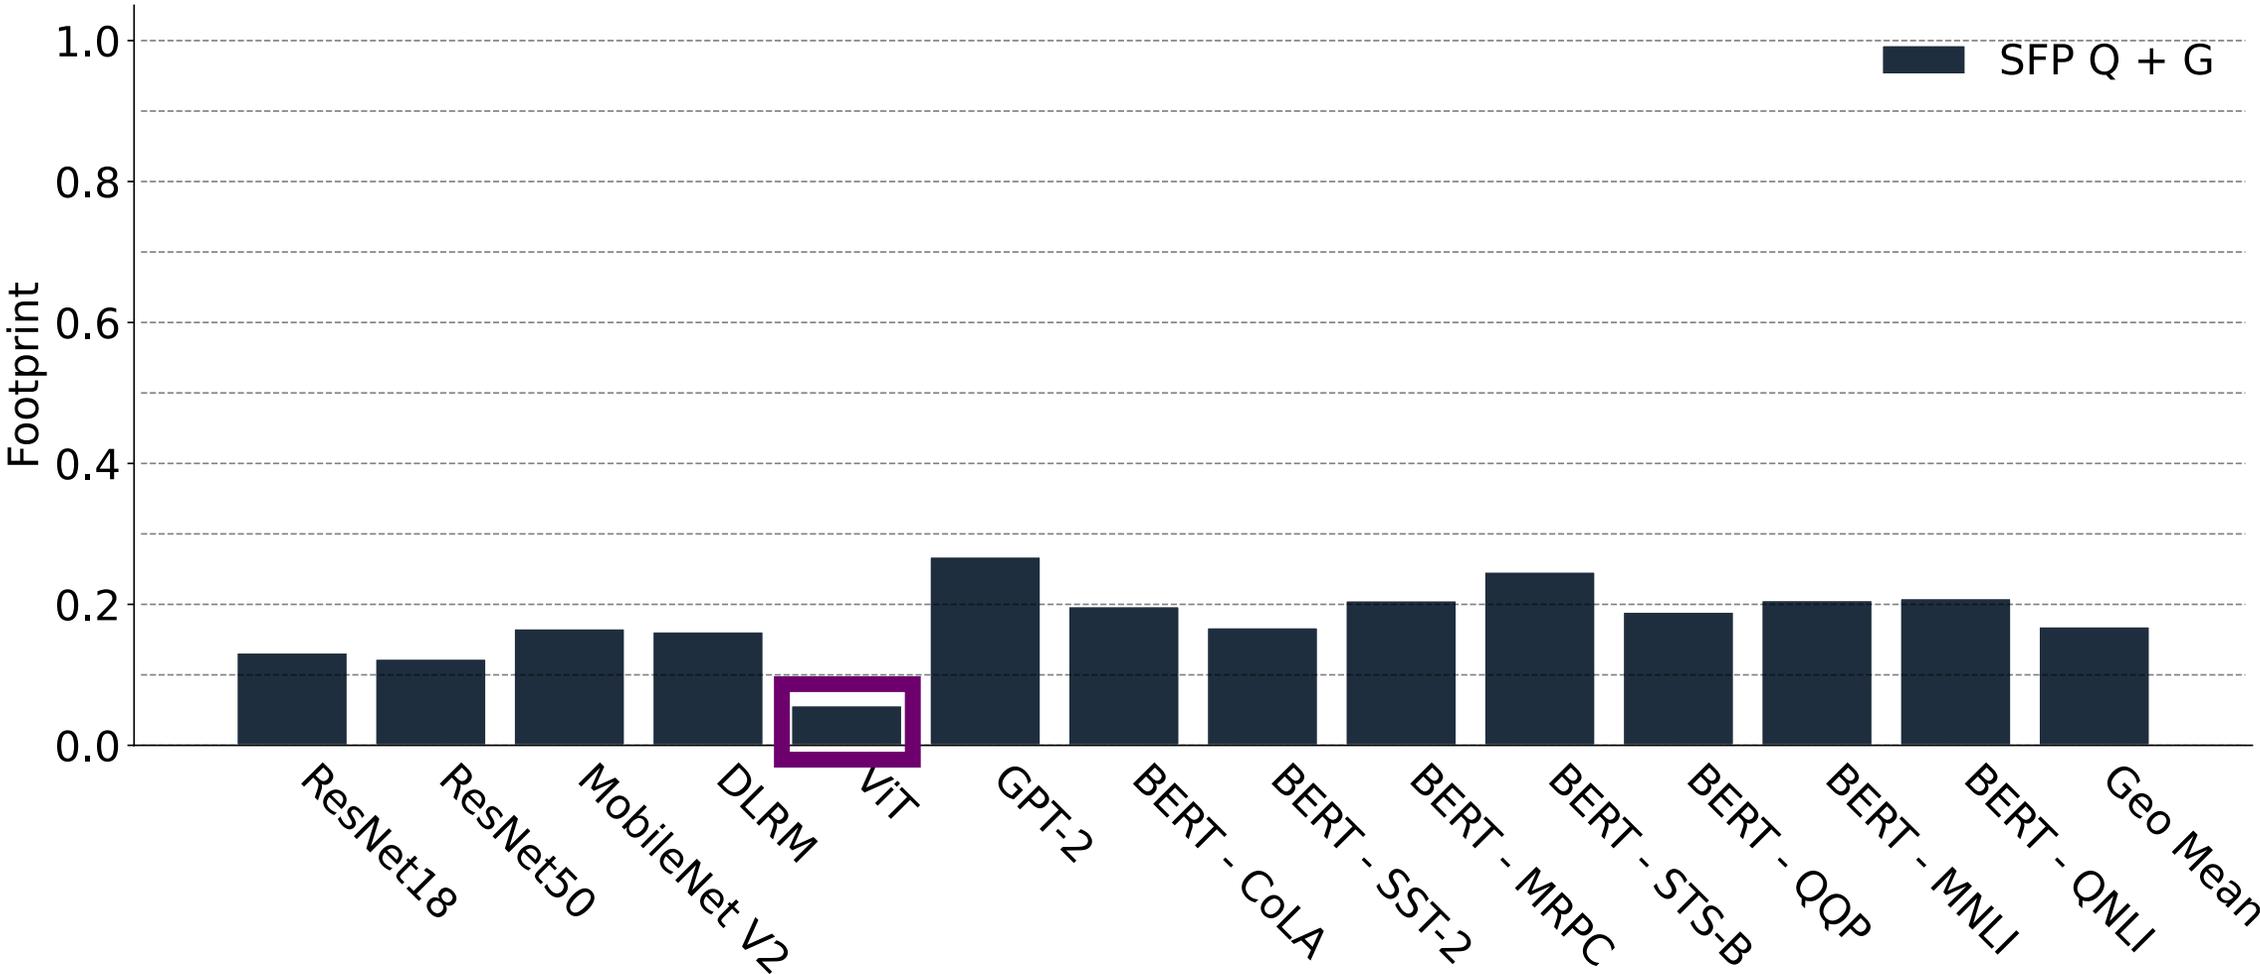
Massive footprint reduction across the board!



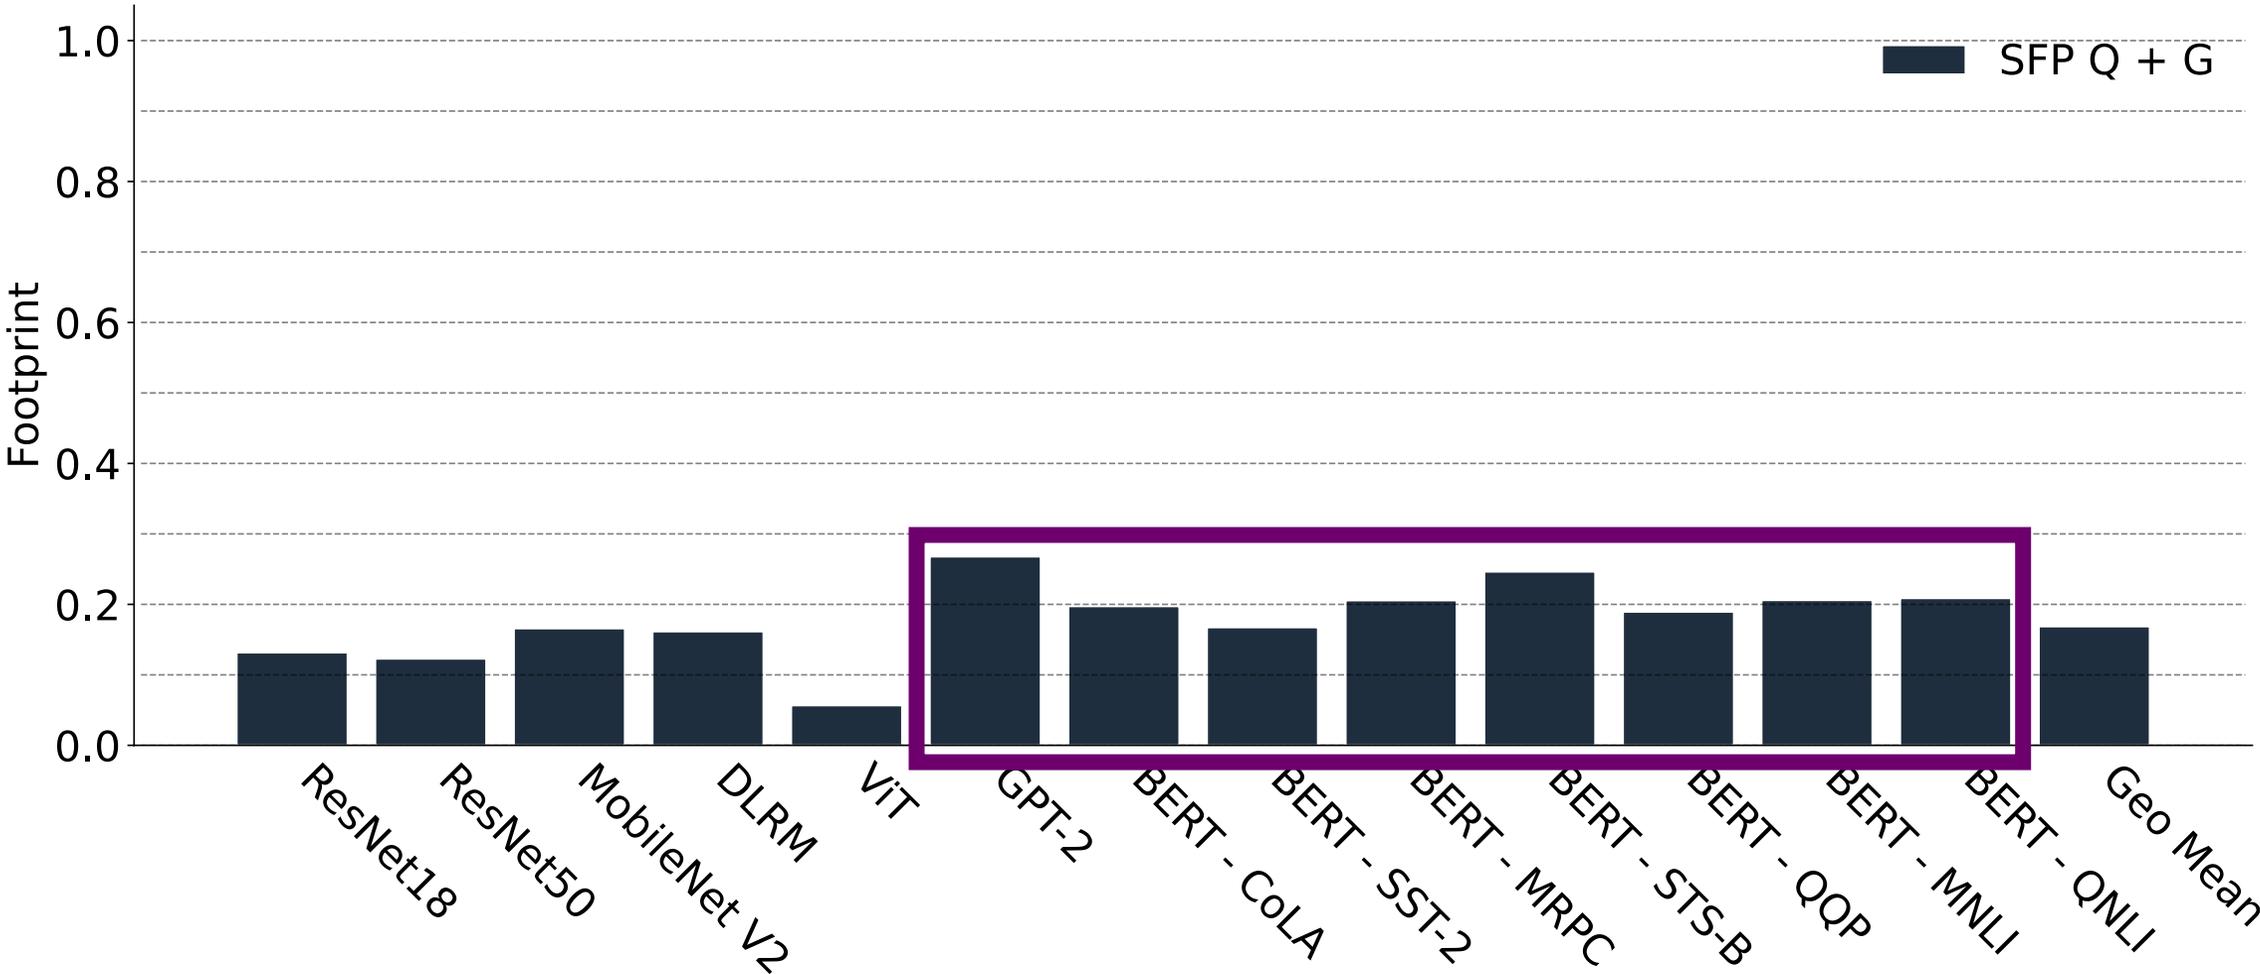
Massive footprint reduction across the board!



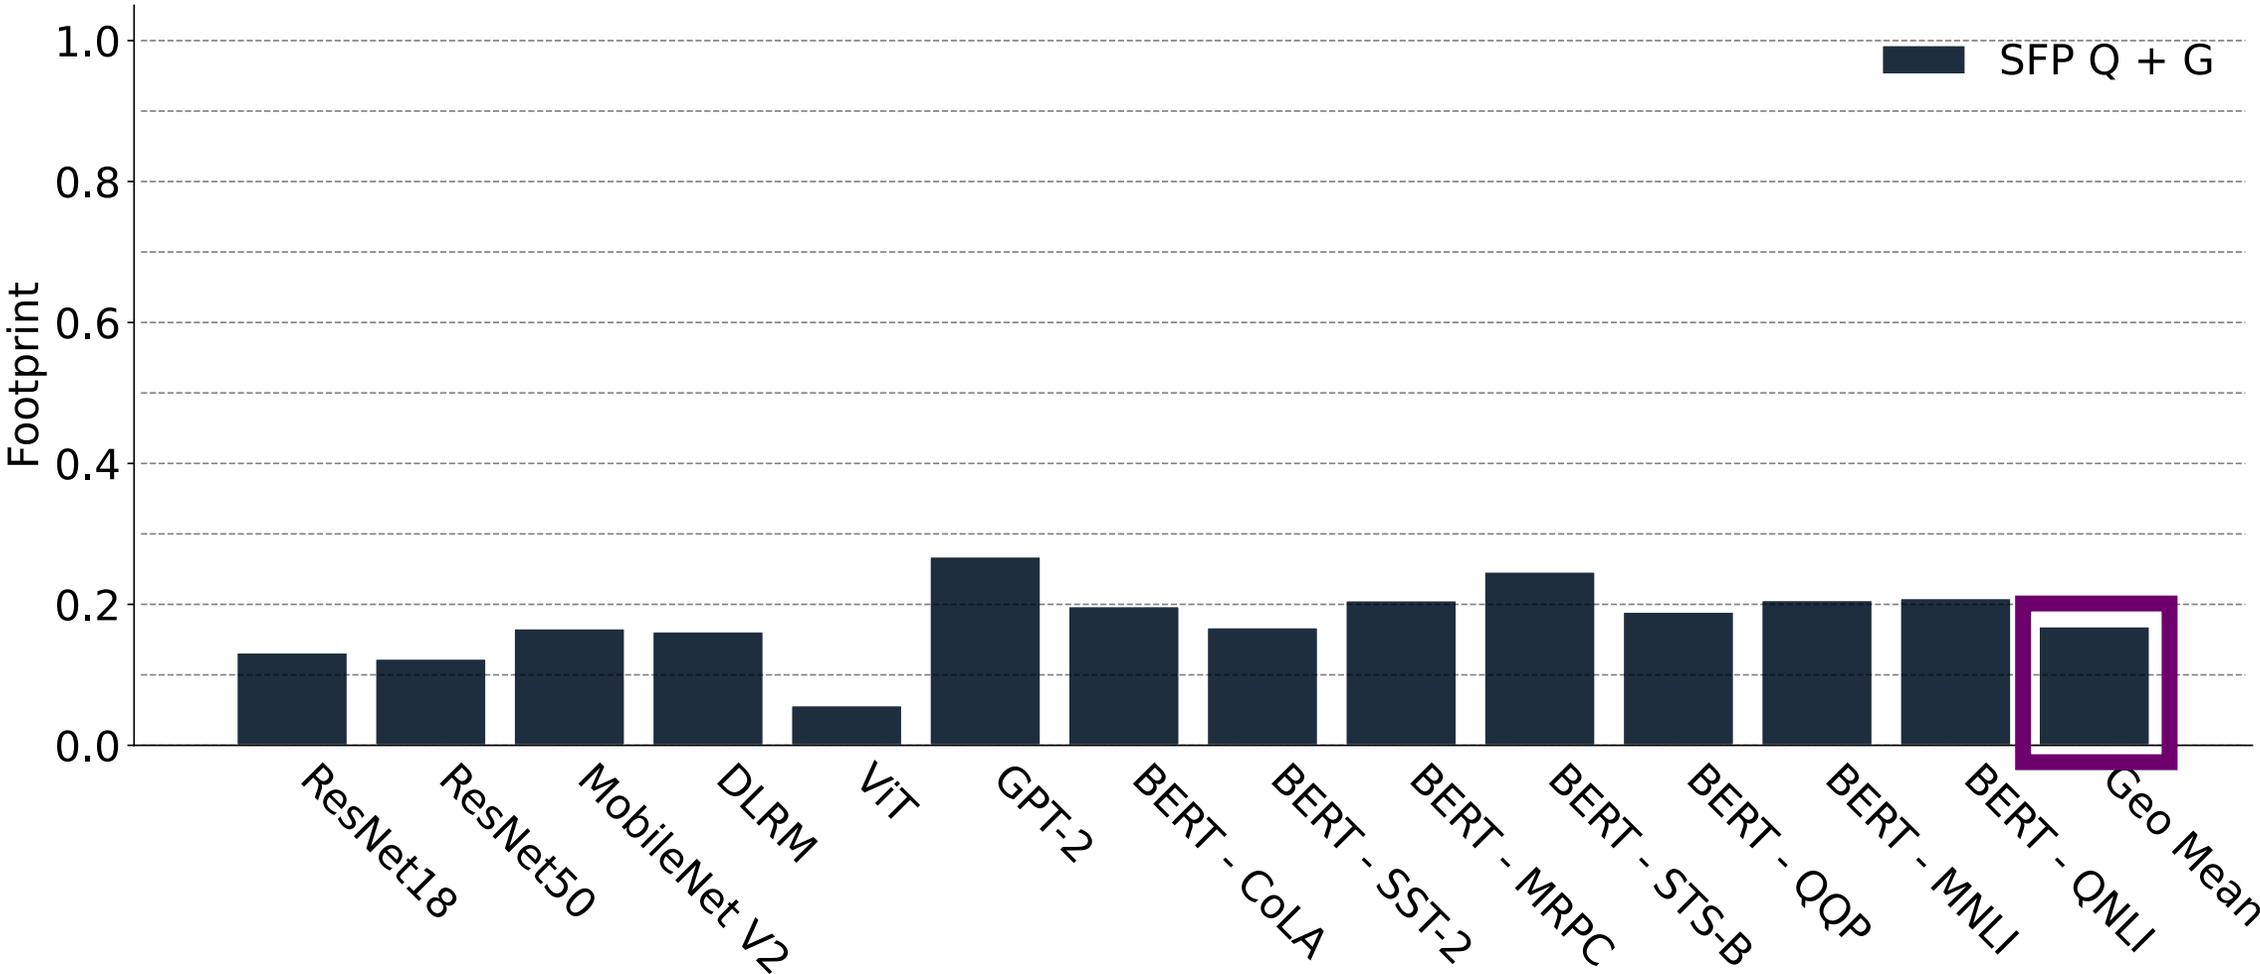
Massive footprint reduction across the board!



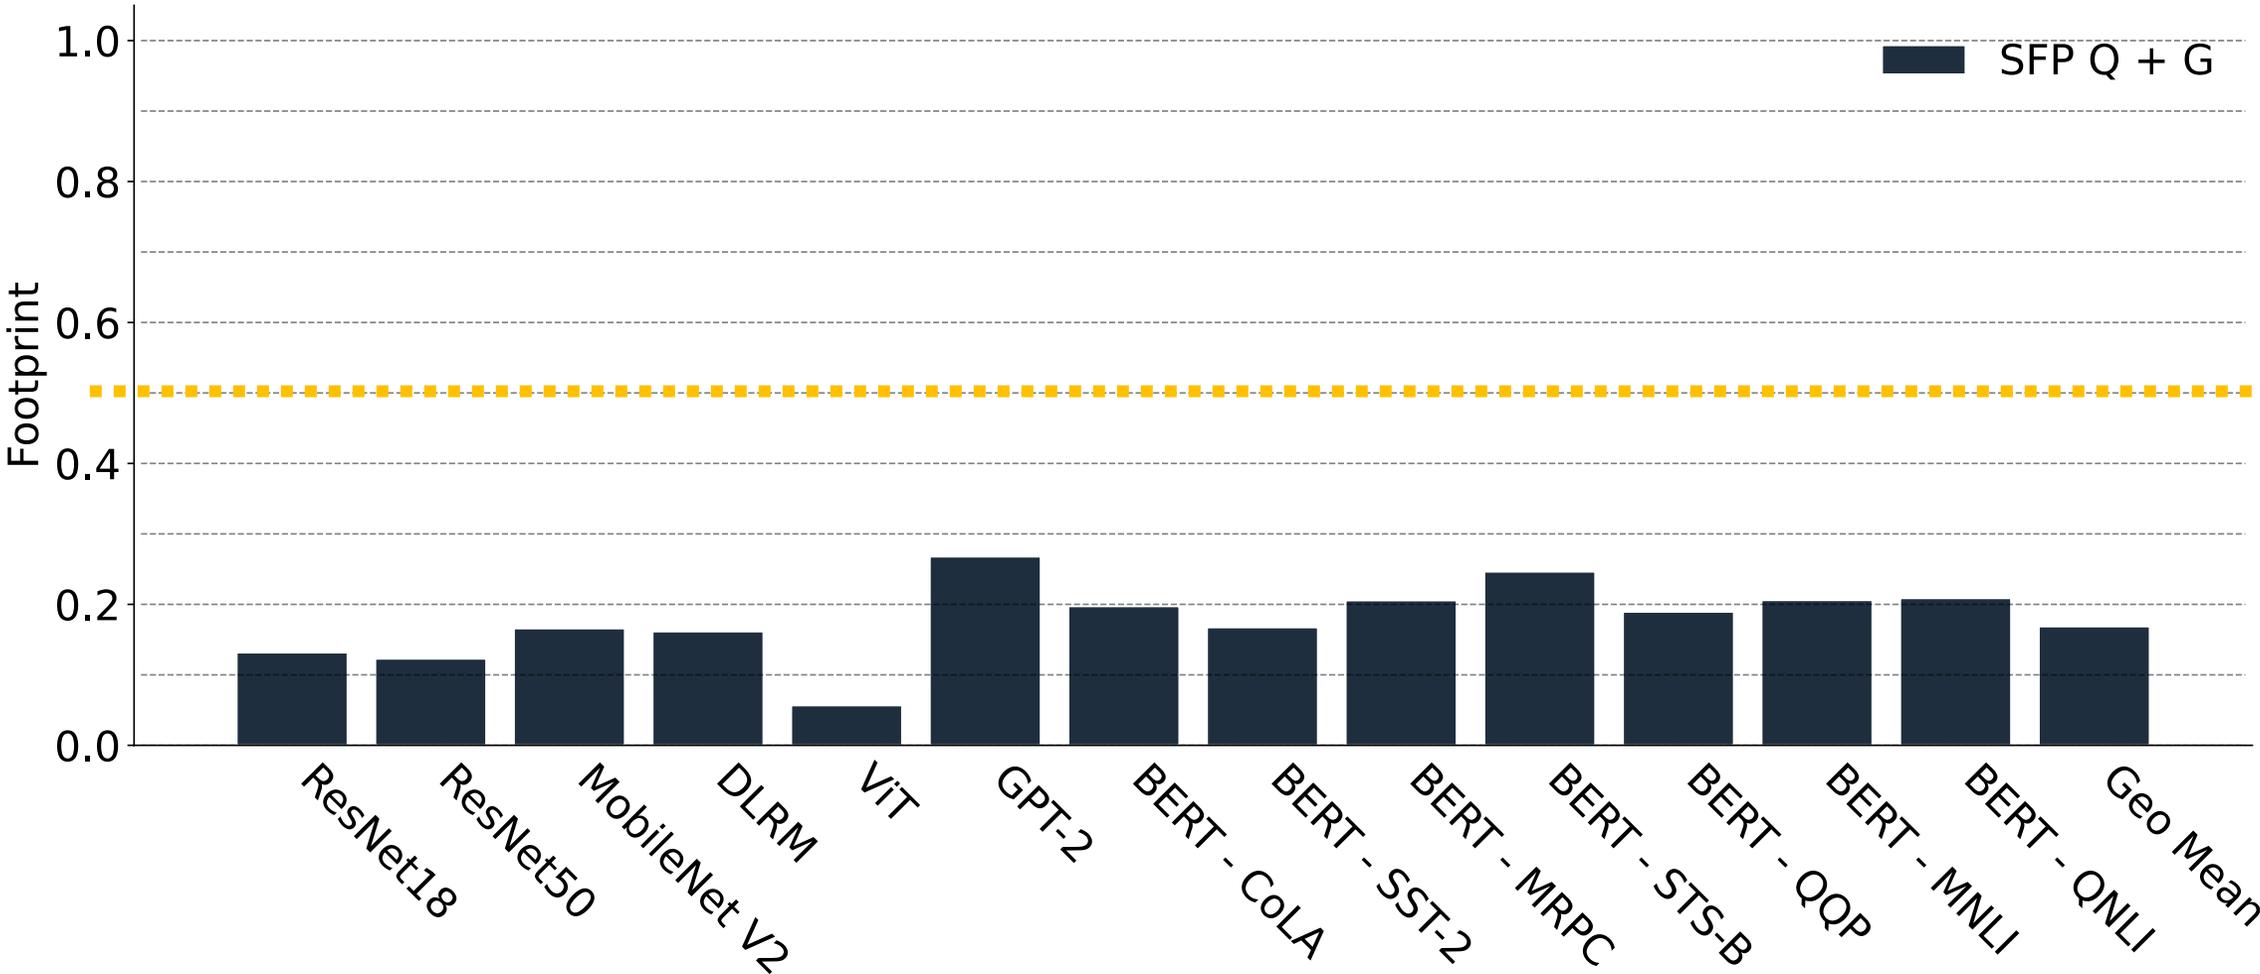
Massive footprint reduction across the board!



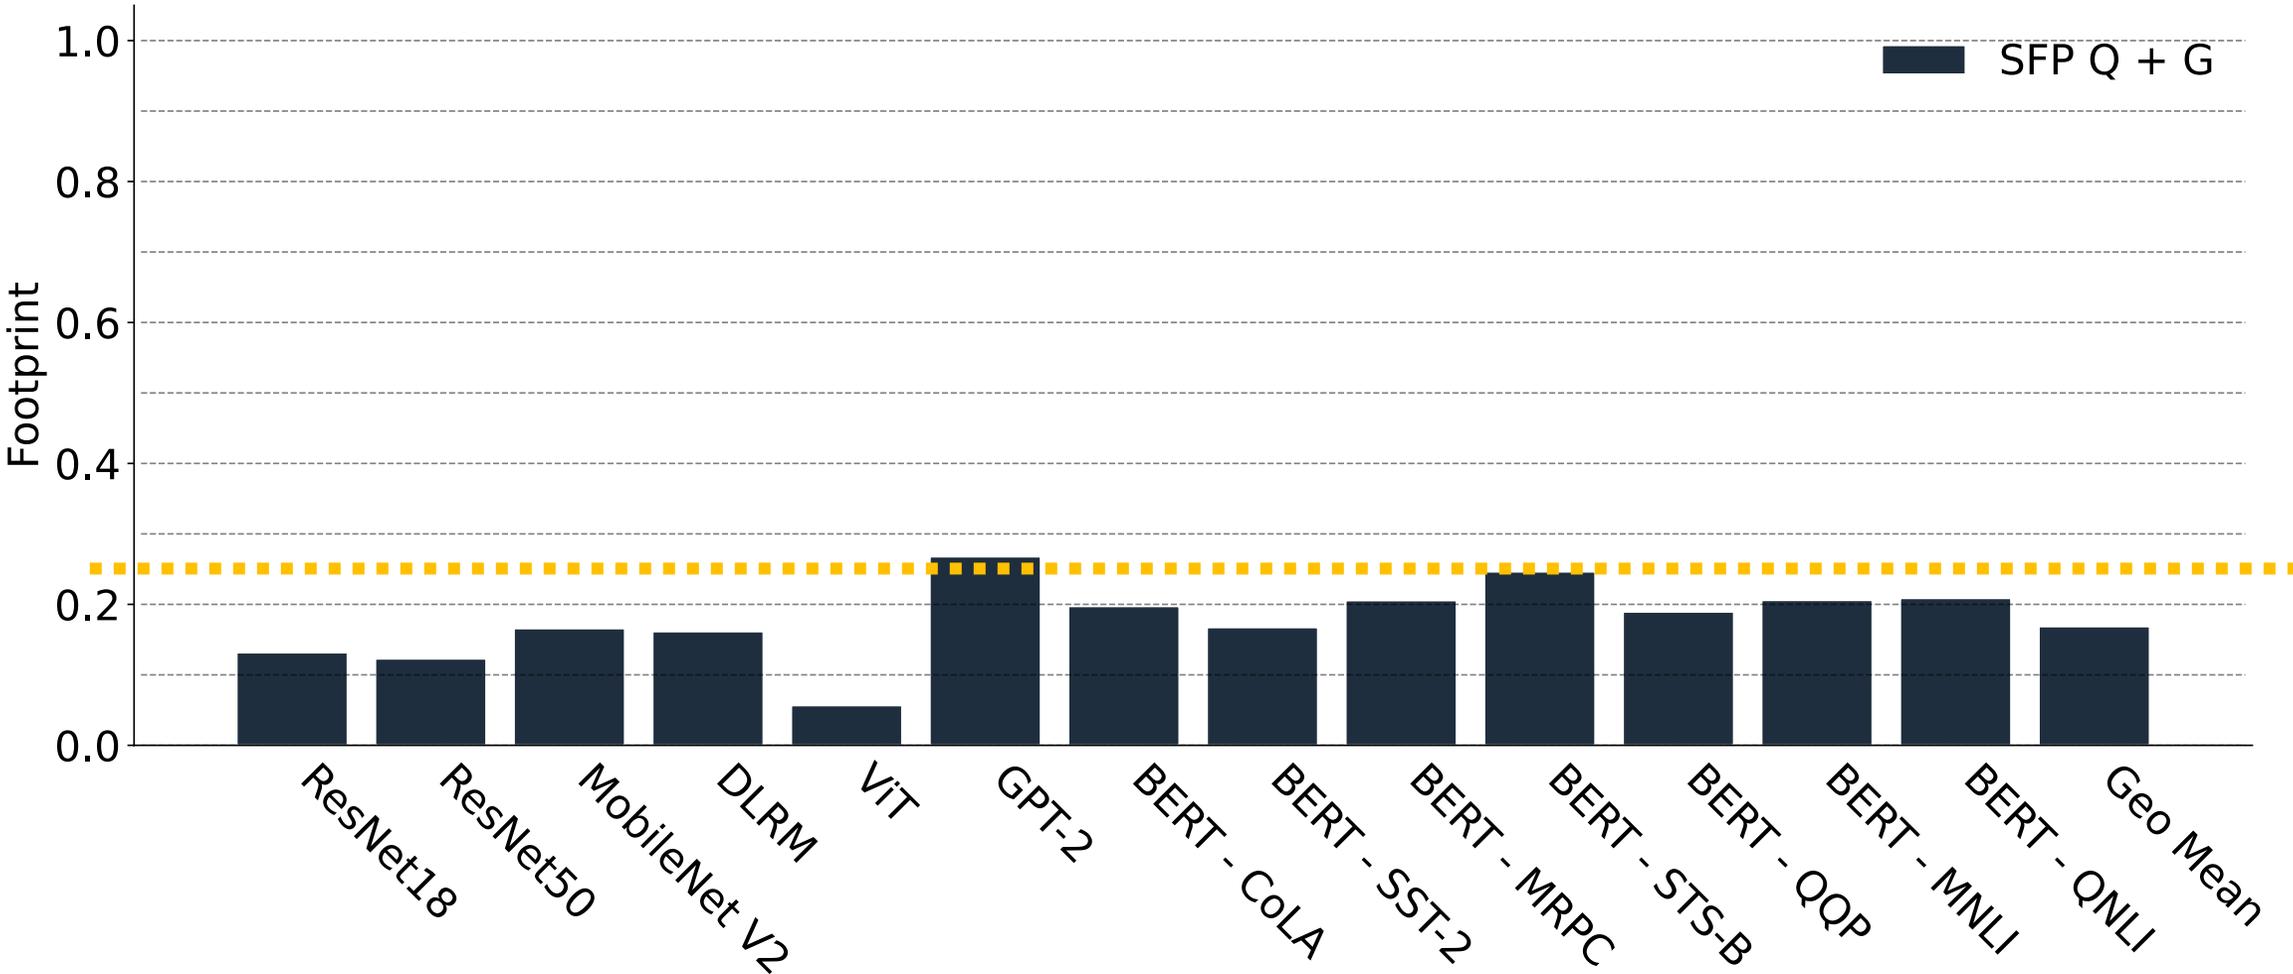
Massive footprint reduction across the board!



Massive footprint reduction across the board!

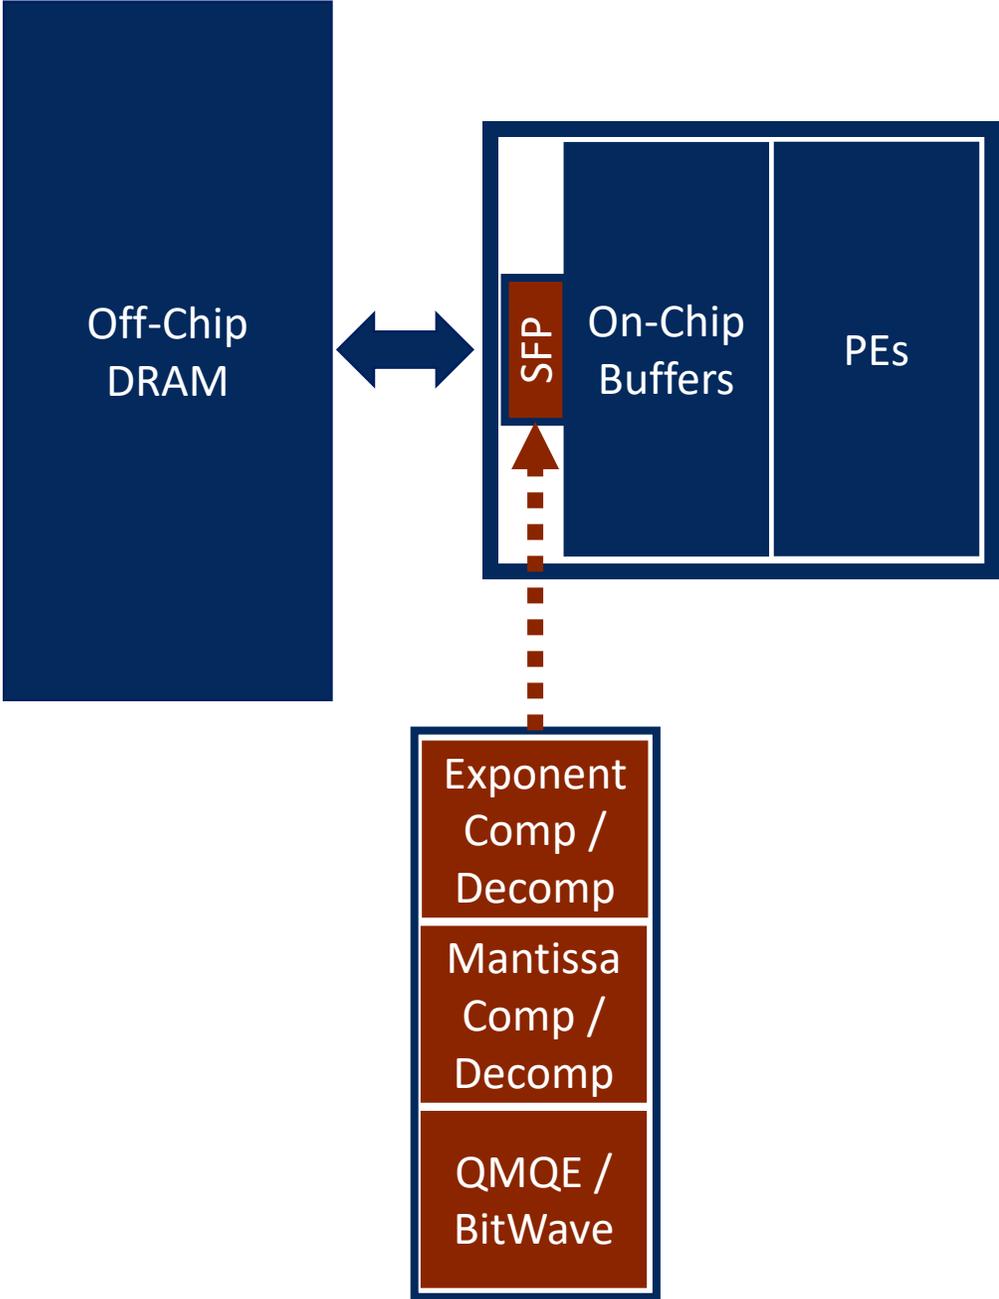


Massive footprint reduction across the board!



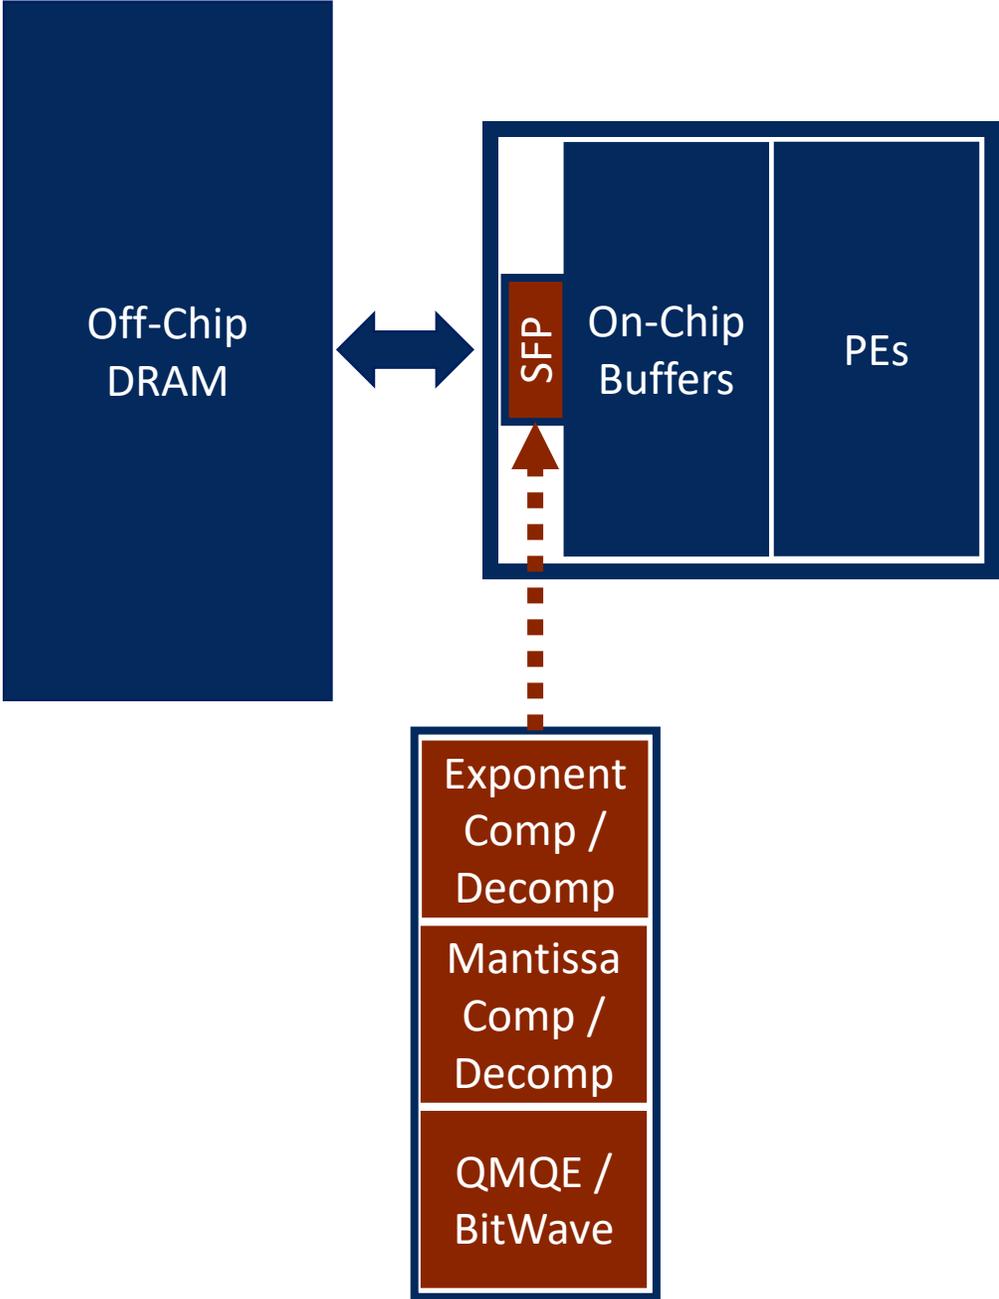
How to exploit?

- Between on- and off-chip memory



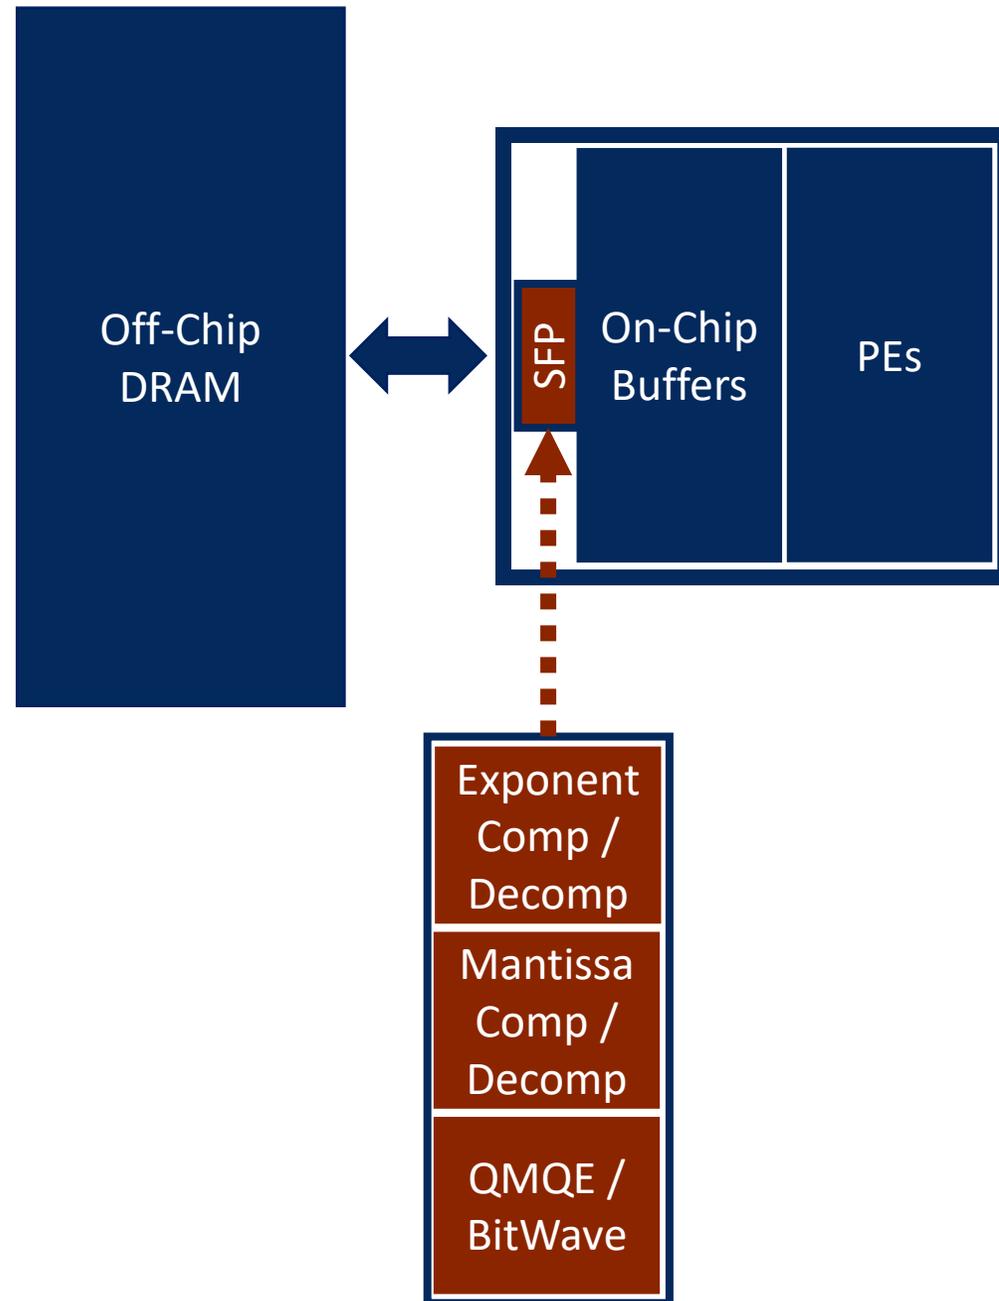
How to exploit?

- Between on- and off-chip memory
- Seamless



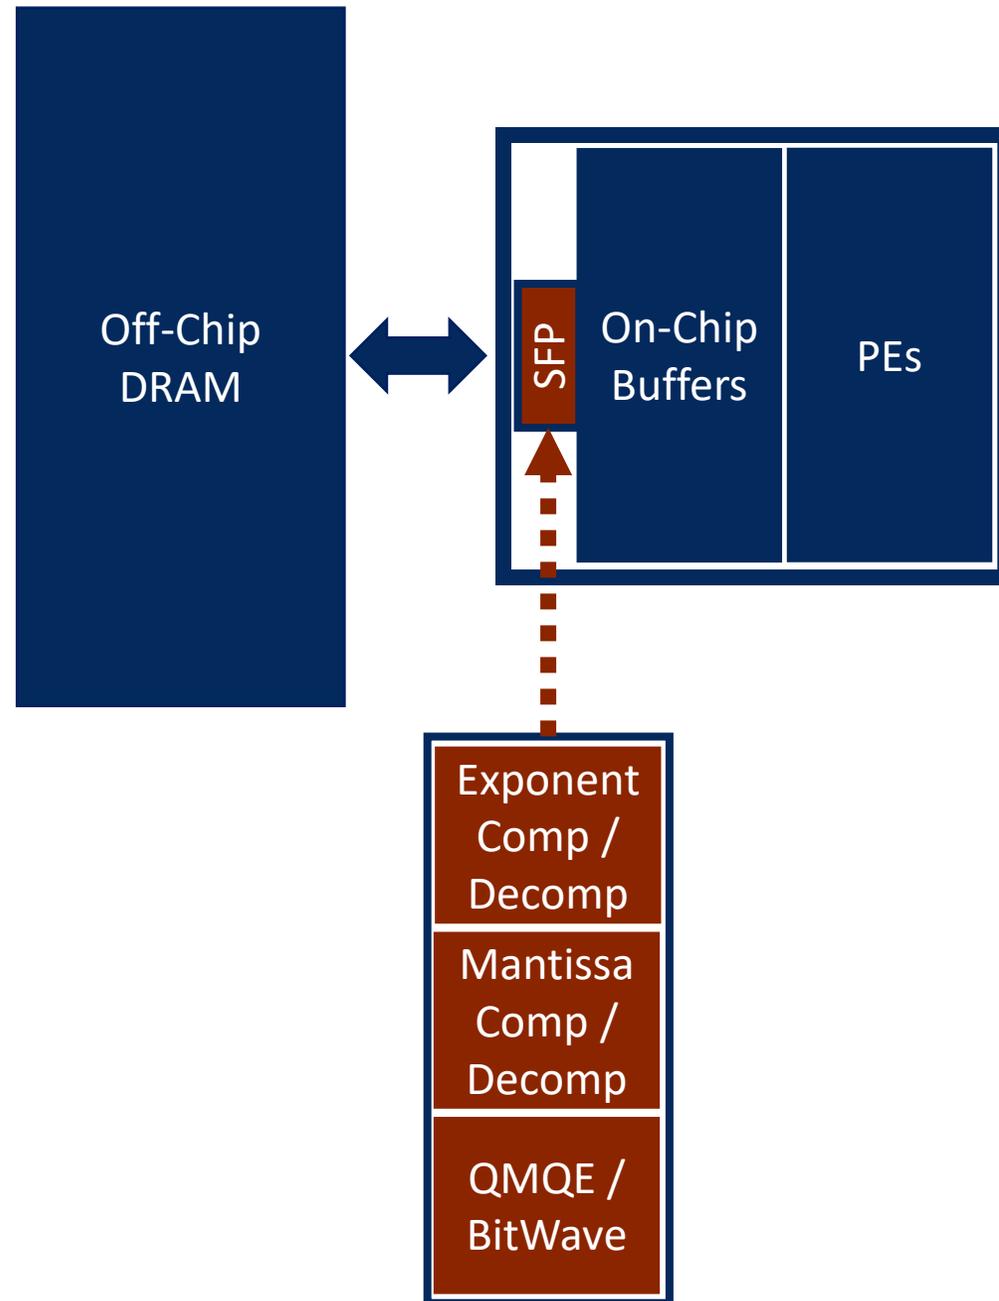
How to exploit?

- Between on- and off-chip memory
- Seamless
- Evaluate in hardware
 - Can be done in software



How to exploit?

- Between on- and off-chip memory
- Seamless
- Evaluate in hardware
 - Can be done in software
- vs. FP8 baseline
 - **2.6x** performance
 - **2.3x** energy efficiency



More in the paper!

- *BitWave* description & results
- Behavior through time
- Results for every network
- Hardware IP blocks
- Simulation results
- Simulation details

Conclusion

- **Schrödinger's FP**: Automatic datatype selection and compression methods



Conclusion

- **Schrödinger's FP**: Automatic datatype selection and compression methods
 - Machine Learning
 - **4.7x** & **5.6x** (+G)



Conclusion

- **Schrödinger's FP**: Automatic datatype selection and compression methods
 - Machine Learning
 - **4.7x** & **5.6x** (+G)
 - Sampling
 - **3.2x** & **4.6x** (+G)



Conclusion

- **Schrödinger's FP**: Automatic datatype selection and compression methods
 - Machine Learning
 - **4.7x** & **5.6x** (+G)
 - Sampling
 - **3.2x** & **4.6x** (+G)
 - Optionally +Gecko



Conclusion

- **Schrödinger's FP**: Automatic datatype selection and compression methods
 - Machine Learning
 - **4.7x** & **5.6x** (+G)
 - Sampling
 - **3.2x** & **4.6x** (+G)
 - Optionally +Gecko



Questions?

Miloš Nikolić

[<milos.nikolic@mail.utoronto.ca>](mailto:milos.nikolic@mail.utoronto.ca)