

# Quadro vs GeForce GPUs for neural networks in Kaggle competitions??

Comments (9)

posted in [Kaggle Forum](#) 2 years ago

Sort By Hotness



Please [log in](#) to leave a comment



[inversion](#) • a year ago •

Options

4 |



SergiuB wrote

Hi All,

I do have a nvidia quadro k2000D and i would like your feedback if is good to keep or replace with a Titan ? (I use it only for computation primarilly).

This is a great article about GPUs for deep learning.

<http://timdettmers.com/2014/08/14/which-gpu-for-deep-learning/>

Titan X is great if you have the budget. The GTX 970 is a good sweet spot for performance vs price.



[sedielem](#) • 2 years ago •

Options

10 |



Quadro GPUs aren't for scientific computation, Tesla GPUs are. Quadro cards are designed for accelerating CAD, so they won't help you to train neural nets. They can probably be used for that purpose just fine, but it's a waste of money.

Tesla cards are for scientific computation, but they tend to be pretty expensive. The good news is that many of the features offered by Tesla cards over GeForce cards are not necessary to train neural networks.

For example, Tesla cards usually have ECC memory, which is nice to have but not a requirement. They also have much better support for double precision

computations, but single precision is plenty for neural network training, and they perform about the same as GeForce cards for that.

One useful feature of Tesla cards is that they tend to have a lot more RAM than comparable GeForce cards. More RAM is always welcome if you're planning to train bigger models (or use RAM-intensive computations like FFT-based convolutions).

If you're choosing between Quadro and GeForce, definitely pick GeForce. If you're choosing between Tesla and GeForce, pick GeForce, unless you have a lot of money and could really use the extra RAM.



[Triskelion](#) • 2 years ago •

Options

3 |



See also [Tim Dettmers](#) excellent [blog](#) on this [topic](#).



[zeros](#) • 2 years ago •

Options

0 |

Thank you very much, Sediem. You saved me a whole lot of angst!

Am certainly on a budget and will go with GeForce.



[zeros](#) • 2 years ago •

Options

0 |

Thank you, Triskelion. Very helpful link!

You guys are awesome.



[juandoso](#) • 2 years ago •

Options

0 |

Spot Instances on AWS could be an alternative: the G2 instance, with a really big Nvidia GPU (1,536 CUDA cores and 4GB), 8 cores CPU and 15 GB of RAM cost only about \$0.065 per hour



[Robert F...](#) • 2 years ago •

Options

0 |

A better set of questions might be..

Which types of machines learning emphasize FP64, Bandwidth/Memory Consumption.

How do you calculate your required memory footprint once you have selected an appropriate algorithm.

Will Amdahl's law be good to you/Is the algorithm highly parallelisable?

Can you represent this problem with a matrices and vectors?

Do you require out of order execution?

Do you require logic executing in your threads?

A xeon phi has 1Teraflop of fp64/ 2 TF of fp32 and 56'ish cores with 4 threads each capable of running either a set of extended vector instruction or x86 logic, runs a linux shell and can access nfs. 6gb ddr5 for the ones I have

but a

Titan X has 0.2 teraflops of fp64/ 6.2 Tf of fp32, several 1000 shader kernels that can do vector math but can't execute programming logic and 12gb of ddr?

Can you fit you program on one card? Does it span many? Latency can kill any advantages you get from multiple nodes.

There are lots of fun questions in this game.

I think most people just race what they have and focus on the learning part of Kaggle. If you are playing to win you likely already know the answers to the hardware questions and are either trying to decide whether spending X\$ on compute is worth the price of a lottery ticket or if you can get away with using your universities cluster for inappropriate gain :)



SergiuB • a year ago •

Options

0 |

Hi All,

I do have a nvidia quadro k2000D and i would like your feedback if is good to keep or replace with a Titan ? (I use it only for computation primarilly).

Thank you,



Le0nX • 16 days ago •

Options

0 |

Very helpful discussion. Thanks!