

Is there any noticeable difference in TensorFlow performance if using Quadro GPUs vs GeForce GPUs?

7

**e.g. does it use double precision operations or something else that would cause a drop in GeForce cards?**

I am about to buy a GPU for TensorFlow, and wanted to know if a GeForce would be ok. Thanks and appreciate your help

1

[machine-learning](#) [gpu](#) [gpgpu](#) [tensorflow](#)

share improve this question

edited Nov 30 '16 at 14:20

asked Jan 11 '16 at 5:57



[gordatron](#)

2,133 2 13 34



[user2771184](#)

69 1 4

3 This is 100% not opinion based. – [Goddard Jul 30 '16 at 16:05](#)

1 This question is not opinion based, tensorflow is a specific application and there are specific hardware differences between these cards. The question even points to one of the main differences and asks whether the programming library uses that technology. – [gordatron Nov 30 '16 at 14:18](#)

[add a comment](#)

## 1 Answer

[active](#)

[oldest](#)

[votes](#)

I think GeForece TITAN is great and widely used in Machine Learning (ML) ares. In ML, single precision is enough in most of cases.

14

Detail performance of GTX in wiki [here](#).

A comment in [here](#).

For comparison, an “entry-level” \$700 Quadro 4000 is significantly slower than a \$530 high-end GeForce GTX 680, at least according to my measurements using several Vrui applications, and the closest performance-equivalent to a GeForce GTX 680 I could find was a Quadro 6000 for a whopping \$3660.

And specific for ML, such as deep learning, a comment as below in [kaggle](#), which compared Quadro, GeForce and Tesla:

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.

**NOTE:** Be careful what platform you are working on and what's the default precision in it. For example in [here](#), the developer works on R and want to accelerate its applications by NVIDIA Geforce Series GPU. However, the default precision in R is double precision so worse performance on new GPU. Actually, Tesla GPU is better for double precision.

Double precision performance of Titan X is pretty low.

share improve this answer

edited Sep 20 '16 at 6:59

answered Jan 11 '16 at 6:28



Patric

1,077 7 11

---

2 "Quadro GPUs aren't for scientific computation, Tesla GPUs are" +1 – [Guy Coder](#) Jan 11 '16 at 12:26

---

yes, definitely the Titan X or the new 1080 tli – [fabrizioM](#) Mar 4 at 6:28

---