



PyProf

Release Notes

Table of Contents

Chapter 1. PyProf 21.04 Release Notes.....	1
Chapter 2. PyProf 21.03 Release Notes.....	2
Chapter 3. PyProf 21.02 Release Notes.....	3
Chapter 4. PyProf 21.01 Release Notes.....	4
Chapter 5. PyProf 20.12 Release Notes.....	5
Chapter 6. PyProf 20.11 Release Notes.....	6
Chapter 7. PyProf 20.10 Release Notes.....	7
Chapter 8. PyProf 20.09 Release Notes.....	8
Chapter 9. PyProf 20.08 Release Notes.....	9
Chapter 10. PyProf 20.07 Release Notes.....	10
Chapter 11. PyProf 20.06 Release Notes.....	11

Chapter 1. PyProf 21.04 Release Notes

Description

PyProf release for 21.04 is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#), and as a Python Wheel on the NVIDIA PY Index.

New Features

The key features of PyProf v3.10.0 / r21.04 are:

- ▶ Latest PyProf version supports [PyTorch 1.8.0](#) and [Nsight Systems 2020.4.3](#).
- ▶ Latest PyProf version is compatible with [DLProf v1.1.0 / r21.04](#).

Known Issues

- ▶ This software only supports PyTorch 1.8.

Resolved Issues

- ▶ None

Chapter 2. PyProf 21.03 Release Notes

Description

PyProf release for 21.03 is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#), and as a Python Wheel on the NVIDIA PY Index.

New Features

The key features of PyProf v3.9.0 / r21.03 are:

- ▶ Latest PyProf version supports [PyTorch 1.8.0](#) and [Nsight Systems 2020.4.3](#).
- ▶ Latest PyProf version is compatible with [DLProf v1.0.0 / r21.03](#).

Known Issues

- ▶ This software only supports PyTorch 1.8.

Resolved Issues

- ▶ None

Chapter 3. PyProf 21.02 Release Notes

Description

PyProf release for 21.02 is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#), and as a Python Wheel on the NVIDIA PY Index.

New Features

The key features of PyProf v3.8.0 / r21.02 are:

- ▶ Latest PyProf version supports [PyTorch 1.8.0](#) and [Nsight Systems 2020.3.4](#).
- ▶ Latest PyProf version is compatible with [DLProf v0.19.0 / r21.02](#).
- ▶ Torch.ones and Torch.zeros input can now be a list, tuple, or a sequence of integers.

Known Issues

- ▶ This software only supports PyTorch 1.8.

Resolved Issues

- ▶ Tensor.log_softmax will no longer trigger an assertion.

Chapter 4. PyProf 21.01 Release Notes

Description

The NVIDIA release for PyProf 21.01 has been canceled. The next release will be the 21.02 release which is expected to be released at the end of February.

Chapter 5. PyProf 20.12 Release Notes

Description

PyProf release for 20.12, is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#), and as a Python Wheel on the NVIDIA PY Index.

New Features

The key features of PyProf v3.7.0 / r20.12 are:

- ▶ Latest PyProf version supports [PyTorch 1.8.0](#) and [Nsight Systems 2020.3.4](#).
- ▶ Latest PyProf version is compatible with [DLProf v0.18.0 / r20.12](#).
- ▶ Monkey patching support for APEX libraries.

Known Issues

- ▶ This software only supports PyTorch 1.5.

Resolved Issues

- ▶ None

Chapter 6. PyProf 20.11 Release Notes

Description

PyProf release for 20.11, is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#), and as a Python Wheel on the NVIDIA PY Index.

New Features

The key features of PyProf v3.6.0 / r20.11 are:

- ▶ Latest PyProf version supports [PyTorch 1.5.0](#) (PyTorch 1.6.0 with DLProf) and [Nsight Systems 2020.3.2](#).
- ▶ Latest PyProf version is compatible with [DLProf v0.17.0 / r20.11](#).
- ▶ PyProf overhead was reduced to improve runtime performance:
 - ▶ Improved database query from Nsight Systems
 - ▶ Refactored nvmarker.py

Known Issues

- ▶ This software only supports PyTorch 1.5.
- ▶ Forward-Backward kernel correlation heuristics do not work correctly with PyTorch 1.6.

Recommended work arounds include:

- ▶ Use with PyTorch 1.5.
- ▶ Use the 20.11-py3 PyTorch NGC container:

```
docker pull nvcr.io/nvidia/pytorch:20.11-py3
```
- ▶ Use DLProf in the [20.11 NGC PyTorch container](#).

Resolved Issues

- ▶ None

Chapter 7. PyProf 20.10 Release Notes

Description

PyProf release for 20.10, is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#).

New Features

The key features of PyProf v3.5.0 / r20.10 are:

- ▶ Latest PyProf version supports [PyTorch 1.5.0](#) (PyTorch 1.6.0 with DLProf) and [Nsight Systems 2020.3.2](#).
- ▶ Latest PyProf version is compatible with [DLProf v0.16.0 / r20.10](#).
- ▶ Nsight System database lookup improved to speed up the runtime profile analysis time by 50x.
- ▶ Node names will now include class info and can be linked back to the original Python source.

Known Issues

- ▶ This software only supports PyTorch 1.5.
- ▶ Forward-Backward kernel correlation heuristics do not work correctly with PyTorch 1.6.

Recommended work arounds include:

- ▶ Use with PyTorch 1.5.
- ▶ Use the 20.09-py3 PyTorch NGC container:

```
docker pull nvcr.io/nvidia/pytorch:20.10-py3
```
- ▶ Use DLProf in the [20.10 NGC PyTorch container](#) or install using Python wheel.

Resolved Issues

- ▶ None

Chapter 8. PyProf 20.09 Release Notes

Description

PyProf release for 20.09, is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#).

New Features

The key features of PyProf v3.4.0 / r20.09 are:

- ▶ Latest PyProf version on GitHub supports [PyTorch 1.5.0](#) (PyTorch 1.6.0 with DLProf) and [Nsight Systems 2020.3.2](#).
- ▶ Latest PyProf version compatible with [DLProf v0.15.0 / r20.09](#).
- ▶ README and User Guide documentation has been updated with more installation options and pointers.

Known Issues

- ▶ This software only supports PyTorch 1.5 (PyTorch 1.6.0 when used as part of DLProf).
- ▶ Forward-Backward kernel correlation heuristics do not work correctly with PyTorch 1.6.

Recommended work arounds include:

- ▶ Use with PyTorch 1.5.
- ▶ Use the 20.09-py3 PyTorch NGC container:

```
docker pull nvcr.io/nvidia/pytorch:20.09-py3
```
- ▶ Use DLProf in the [20.09 NGC PyTorch container](#).

Resolved Issues

- ▶ None

Chapter 9. PyProf 20.08 Release Notes

Description

PyProf release for 20.08, is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#).

New Features

The key features of PyProf v3.3.0 / r20.08 are:

- ▶ Latest PyProf version supports [PyTorch 1.5.0](#) (PyTorch 1.6.0 with DLProf) and [Nsight Systems 2020.3.2](#).
- ▶ Latest PyProf version compatible with [DLProf v0.14.0 / r20.08](#).
- ▶ Capture PyTorch API information and data loading configuration.
- ▶ Added CUTLASS to the list of GEMM kernels.
- ▶ Added optional function stack tracing to NVTX markers. Enable with `pyprof.init(enable_function_stack=True)`.

Known Issues

- ▶ This software only supports PyTorch 1.5.
- ▶ Forward-Backward kernel correlation heuristics do not work correctly with PyTorch 1.6.

Recommended work arounds include:

- ▶ Use with PyTorch 1.5.
- ▶ Use the 20.03-py3 PyTorch NGC container:
`docker pull nvcr.io/nvidia/pytorch:20.03-py3`
- ▶ Use DLProf in the [20.08 NGC PyTorch container](#).

Resolved Issues

- ▶ None

Chapter 10. PyProf 20.07 Release Notes

Description

PyProf release for 20.07, is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#).

New Features

The key features of PyProf v3.2.0 / r20.07 are:

- ▶ Latest PyProf version supports [PyTorch 1.6.0](#) and [Nsight Systems 2020.3.2](#)
- ▶ Latest PyProf version compatible with [DLProf v0.13.0 / r20.07](#)
- ▶ Monkey patch support for APEX, fused Adam, and Layer Norm functions.
- ▶ PyYAML requirement has been removed.
- ▶ Error handling for non-existent parse file arguments has been added.

Known Issues

- ▶ This software only supports PyTorch 1.6.
- ▶ Forward-Backward kernel correlation heuristics do not work correctly with PyTorch 1.6.

Recommended work arounds include:

- ▶ Use with PyTorch 1.5.
- ▶ Use the 20.03-py3 PyTorch NGC container:

```
docker pull nvcr.io/nvidia/pytorch:20.03-py3
```
- ▶ Use DLProf in the [20.08 NGC PyTorch container](#).

Resolved Issues

- ▶ None

Chapter 11. PyProf 20.06 Release Notes

Description

The PyProf release for 20.06 is available in the NVIDIA [PyTorch](#) NGC containers and directly from the [PyProf GitHub page](#).

New Features

- ▶ Latest PyProf version supports [PyTorch 1.6.0](#) and [Nsight Systems 2020.2.5](#)
- ▶ Latest PyProf version compatible with [DLProf v0.12.0 / r20.06](#)
- ▶ Improved NVTX functionality and coverage.
- ▶ Add support required by [NVIDIA Deep Learning Profiler](#).

Known Issues

- ▶ This software only supports PyTorch 1.6.

Resolved Issues

- ▶ None

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, and cuBLAS, CUDA, cuDNN, DALI, DIGITS, DGX, DGX-1, DGX-2, DGX Station, DLProf, Jetson, Kepler, Maxwell, NCCL, Nsight Compute, Nsight Systems, NvCaffe, NVIDIA Ampere GPU Architecture, PerfWorks, Pascal, SDK Manager, Tegra, TensorRT, Triton Inference Server, Tesla, TF-TRT, and Volta are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2020-2021 NVIDIA Corporation. All rights reserved.

