


Deep Learning HDL Toolbox™ Release Notes



MATLAB®



How to Contact MathWorks



Latest news: www.mathworks.com
Sales and services: www.mathworks.com/sales_and_services
User community: www.mathworks.com/matlabcentral
Technical support: www.mathworks.com/support/contact_us



Phone: 508-647-7000



The MathWorks, Inc.
1 Apple Hill Drive
Natick, MA 01760-2098

Deep Learning HDL Toolbox™ Release Notes

© COPYRIGHT 2020—2023 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

FEDERAL ACQUISITION: This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

R2023a

Support for gated recurrent unit layer networks	1-2
Profiler and quantization support for multiple input and output networks	1-2
Updated support for layers and networks	1-2
Updates to supported software	1-3
Functionality being removed or changed	1-3
Ethernet support for Intel Arria 10 SoC board will be removed	1-3

R2022b

Custom long short-term memory network support	2-2
Support for dlnetwork objects	2-2
Support for multiple input and multiple output networks	2-2
Updates to estimatePerformance method	2-2
Updates to optimizeConfigurationForNetwork method	2-3
Updated supported layers and networks	2-3

R2022a

Generated deep learning processor IP core integration with other IP cores	3-2
Generated deep learning processor deployment without a MATLAB connection	3-2

Network custom layer creation, registration, validation, and deployment	3-2
.....	
Updates to deep learning network layer activations results	3-2
Updated supported layers and networks	3-2
Functionality being removed or changed	3-3
adder module of dldhdl.ProcessorConfig object has been renamed	3-3

R2021b

Trimmed deep learning processor configuration	4-2
Generic deep learning processor generation	4-2
Custom reference design functionality for custom boards for deep learning processor IP core integration	4-2
Deep learning processor streaming handshake modes	4-2
Updates to estimatePerformance	4-3
Updates to estimateResources	4-3
Enhancements for quantization of directed acyclic graph (DAG) networks	4-3
.....	
Network prototyping and validation without hardware	4-4
Updated supported layers	4-4
Functionality being removed or changed	4-4
KernelDataType property of dldhdl.ProcessorConfig object has been removed	4-4
.....	
LUT property of estimateResources function has been removed	4-4

R2021a

Custom directed acyclic graph (DAG) network support	5-2
Performance estimation and profiling	5-2
Resource estimation	5-2
Updated supported layers	5-2

MATLAB Emulation for validate method of dlquantizer object	5-3
dlhdl.Workflow name-value argument pair update	5-3
Updates to supported software	5-3
Functionality being removed or changed	5-3
estimate function for dlhdl.Workflow object has been removed	5-3
'ProcessorConfig' option in dlhdl.Workflow has been removed	5-3

R2020b

Introducing Deep Learning HDL Toolbox: Prototype and implement deep learning networks on FPGAs and SoCs	6-2
Prototype on FPGAs	6-2
Custom series network support	6-2
Portable Verilog and VHDL code	6-2
Tune user-configurable parameters	6-2
Custom board support	6-2
Performance estimation and profiling	6-2
Hardware Support	6-2
Support Package for Intel FPGA and SoCs	6-3
Support Package for Xilinx FPGA and SoCs	6-3

R2023a

Version: 1.5

New Features

Bug Fixes

Compatibility Considerations

Support for gated recurrent unit layer networks

In R2023a, Deep Learning HDL Toolbox supports gated recurrent unit (GRU) layer networks. Use Deep Learning HDL Toolbox to:

- Compile your GRU layer networks.
- Deploy your custom GRU layer networks by using the shipping long short-term memory (LSTM) bitstreams or generating a custom bitstream. Retrieve predictions from the deployed network by using MATLAB®.

See “How Deep Learning HDL Toolbox Compiles the GRU Layer”.

Profiler and quantization support for multiple input and output networks

In R2023a, Deep Learning HDL Toolbox supports `int8` quantization for multiple output networks in single and multiple frame modes. Deep Learning HDL Toolbox does not support `int8` quantization of multiple output `dlnetwork` objects.

You can now profile the layer level performance of multiple output networks for single and `int8` data types and multiple input networks for single data types. See “Deploy YAMNet Networks to FPGAs With and Without Cross-Layer Equalization”.

Updated support for layers and networks

In R2023a, Deep Learning HDL Toolbox supports these layers:

- Gated recurrent unit (GRU) layers with these limitations:
 - Inputs must be of single data type.
 - You must set the GRU layer `OutputMode` to `sequence`.
- Dilated convolution layer with these limitations:
 - The dilation factor must have a maximum value of 16 and be symmetric. For example, 16-by-16.
 - When the dilation factor is a multiple of three the calculated dilated filter size must have a maximum value of the existing convolution filter size limit. In all other cases, the filter size can be as large as the maximum value of the existing convolution filter size.

For more information, see “Supported Layers”.

Deep Learning HDL Toolbox now supports these networks:

- YAMNet network. See “Deploy YAMNet Networks to FPGAs With and Without Cross-Layer Equalization”.
- Sequence-to-sequence classification networks that use the GRU layer. See “Run Sequence Forecasting Using a GRU Layer on an FPGA”.
- Semantic Segmentation networks that use dilated convolution layers. See “Deploy Semantic Segmentation Network Using Dilated Convolutions on FPGA”.
- Pruned and quantized image recognition network. See “Deploy Image Recognition Network on FPGA With and Without Pruning”.

-
- Very-deep super-resolution networks. See “Increase Image Resolution Using VDSR Network Running on FPGA”.

Updates to supported software

Deep Learning HDL Toolbox has been tested with:

- Xilinx® Vivado® Design Suite 2022.1
- Intel® Quartus® Prime Standard 21.1

For more information, see “Supported Networks, Layers, Boards, and Tools”.

Functionality being removed or changed

Ethernet support for Intel Arria 10 SoC board will be removed

Ethernet support for the Intel Arria® 10 SoC board will be deprecated in a future release. When you target this board, do not set the `dlhdl.Target object Interface` argument to Ethernet. Ethernet support for this board will be deprecated in a future release. Set `Interface` to JTAG instead.

R2022b

Version: 1.4

New Features

Bug Fixes

Custom long short-term memory network support

In R2022b, Deep Learning HDL Toolbox supports custom long short-term memory (LSTM) networks. Use Deep Learning HDL Toolbox to:

- Compile your LSTM networks.
- Deploy your custom LSTM networks by using the custom LSTM shipping bitstreams or generating a custom bitstream. Retrieve predictions from the deployed network by using MATLAB.
- Retrieve layer-level activations by using the `activations` method.
- Reset layer states by using the `resetState` method of the `dlhdl.Workflow` object.
- Predict responses by using a trained and deployed recurrent neural network and update the deployed network state by using the `predictAndUpdateState` method of the `dlhdl.Workflow` object.

For more information, see Time Series and Sequence Data Networks.

Support for `dlnetwork` objects

In R2022b, Deep Learning HDL Toolbox supports `dlnetwork` objects. Use Deep Learning HDL Toolbox to:

- Deploy the `dlnetwork` object to a target FPGA board and retrieve the prediction results by using MATLAB.
- Visualize the layer level activations for a `dlnetwork` object.
- Estimate the performance of a `dlnetwork` object.
- Optimize the processor configuration for a `dlnetwork` object.

Deep Learning HDL Toolbox supports only initialized `dlnetwork` objects and `dlnetwork` objects that have an image input layer as the network input layer.

When you use a `dlnetwork` object, the inputs must be a `dlarray` of type numeric array. Deep Learning HDL Toolbox does not support `dlnetwork` objects as an input to `dlhdl.Simulator` or for `int8` quantization. See Detect Objects Using YOLO v3 Network Deployed to FPGA.

Support for multiple input and multiple output networks

In R2022b, Deep Learning HDL Toolbox supports a maximum of seven inputs and outputs. Deep Learning HDL Toolbox does not support multiple input and output networks as an input to `dlhdl.Simulator` or for `int8` quantization.

Updates to `estimatePerformance` method

In R2022b, you can use the `estimatePerformance` method to estimate the performance of any Xilinx and Intel devices that have a custom board registration file and custom reference design. Use the

- `buildCalibrationBitstream` method to generate the board-specific calibration bitstream. See `buildCalibrationBitstream`.
- `deployCalibrationBitstream` method to deploy the calibration bitstream and retrieve the board calibration data. See `deployCalibrationBitstream`.

Use the retrieved calibration data to accurately estimate the performance of a network for the target hardware board. See [Deep Learning Processor IP Core Generation for Custom Board](#).

Updates to optimizeConfigurationForNetwork method

In R2022b, you can retrieve the optimized deep learning processor configuration for a network by using the `optimizeConfigurationForNetwork` method and specifying the target frames per second (FPS) value as a name-value pair argument. Generate a custom optimized bitstream by using the optimized deep learning processor configuration as an input to the `dlhdl.buildProcessor` function. Deploy the custom network to your target board by using the custom optimized bitstream. See [Optimize Deep Learning Processor Configuration for Network Performance](#).

Updated supported layers and networks

In R2022b, Deep Learning HDL Toolbox supports these layers:

- Long short-term memory (LSTM) layer
- `resize2DLayer`
- Hyperbolic tangent layer
- `featureInputLayer`
- `sequenceInputLayer`

In R2022b, Deep Learning HDL Toolbox supports these networks:

- Sequence-to-sequence classification networks. See [Run Sequence-to-Sequence Classification on FPGAs by Using Deep Learning HDL Toolbox](#).
- Word-by-word text generation. See [Generate Word-By-Word Text on FPGAs by Using Deep Learning HDL Toolbox](#).
- Time series forecasting networks. See [Run Sequence Forecasting on FPGA by Using Deep Learning HDL Toolbox™](#).

R2022a

Version: 1.3

New Features

Bug Fixes

Compatibility Considerations

Generated deep learning processor IP core integration with other IP cores

In R2022a, interface the generated deep learning processor IP core with other IP cores by specifying the base input and output addresses for the generated deep learning processor IP core. You can specify the addresses through network registers or direct port connections. Before changing the `InputStart` signal, specify the base input and output address. See [Interface with the Deep Learning Processor IP Core](#).

Generated deep learning processor deployment without a MATLAB connection

In R2022a, you can use the Deep Learning HDL Toolbox to deploy your network-associated weights, biases, and instructions to a custom binary file. To parse the custom binary file and initialize the target memory, use a utility to program the deep learning processor IP core without a MATLAB connection. See [Initialize Deployed Deep Learning Processor Without Using a MATLAB Connection](#).

Network custom layer creation, registration, validation, and deployment

In R2022a, you can use the Deep Learning HDL Toolbox to create, register, validate, and deploy your network that has custom layers to your target hardware device. See [Create Deep Learning Processor Configuration for Custom Layers](#).

Deep Learning HDL Toolbox does not support:

- Resource estimation for the deep learning processor configuration that has registered custom layers.
- Performance estimation of networks that have custom layers.

Updates to deep learning network layer activations results

The activation results have been updated in R2022a. You can retrieve deep learning network intermediate layer results for:

- The `Image Input` layer.
- Two output `Maxpool` layers. When you retrieve activation results for the outputs of a `Maxpool` layer that has the `HasUnpoolingIndices` argument set to `true`, the only supported output is `out`.
- Two output `Max Unpooling` layers.

See [activations](#).

Updated supported layers and networks

In R2022a, Deep Learning HDL Toolbox supports these layers:

- Sigmoid
- Transposed convolution 2D

-
- Max unpooling

For `int8` data type quantization, Deep Learning HDL Toolbox supports these layers:

- Sigmoid
- Transposed convolution 2D

In R2022a, Deep Learning HDL Toolbox supports:

- A max pooling 2D layer that has the `HasUnpoolingOutputs` set to `true`.
- 1-by-N and N-by-1 size filters.
- Nonsquare size filters.

Deep Learning HDL Toolbox optimizes nonsymmetric stride sizes by converting them to symmetric stride sizes that produce an equivalent result.

In R2022a, Deep Learning HDL Toolbox implements the normalization parameter of the image input layer on hardware. Deep Learning HDL Toolbox supports only these normalizations for hardware implementation:

- `zerocenter`
- `zscore`

See [Image Input Layer Normalization Hardware Implementation](#).

In R2022a, Deep Learning HDL Toolbox supports these networks:

- U-Net
- Reduced U-Net
- PoseNet. See [Human Pose Estimation by Using Segmentation DAG Network Deployed to FPGA](#).
- SegNet
- Speech Command Recognition. See [Speech Command Recognition by Using FPGA](#).
- Modulation Classification. See [Modulation Classification by Using FPGA](#).

For `int8` data type quantization, Deep Learning HDL Toolbox supports these networks:

- U-Net
- Reduce U-Net
- PoseNet

Functionality being removed or changed

adder module `ofdIhdl.ProcessorConfig` object has been renamed

Behavior change

This property has been renamed to `custom`.

R2021b

Version: 1.2

New Features

Bug Fixes

Compatibility Considerations

Trimmed deep learning processor configuration

Generate a resource-optimized deep learning processor IP core and bitstream that suit your custom convolution module layers only or fully connected module layers only networks. Generate the optimized deep learning processor IP core and bitstream by removing the convolution, fully connected, or adder modules from the deep learning processor configuration. To remove the modules:

- Set the `ModuleGeneration` property to `off`.
- Use the `optimizeConfigurationForNetwork` function.

To further optimize your processor configuration:

- Remove the Local Response Normalization (LRN) block from the processor configuration by setting the `LRNBlockGeneration` property to `off`.
- Remove the Softmax block from the processor configuration by setting the `SoftmaxBlockGeneration` property to `off`. When you set this property to `off`, the Softmax layer is still implemented in software.

See [Generate Custom Bitstream to Meet Custom Deep Learning Network Requirements](#).

Generic deep learning processor generation

Generate a custom generic deep learning processor IP core by specifying `Generic Deep Learning Processor` for the `TargetPlatform` property of the `dlhdl.ProcessorConfig` object. Integrate the generated IP core with your larger FPGA design. You can:

- Specify a name for your project folder by using the `ProjectFolder` name-value argument.
- Name your deep learning processor IP core by using the `ProcessorName` name-value argument.
- Specify HDL code generation options, such as target language for the generated HDL code, by using the `HDLCoderConfig` name-value argument.

Custom reference design functionality for custom boards for deep learning processor IP core integration

Use custom reference design functionality for custom boards and designs for deep learning processor IP core integration. You can:

- Register a custom board to target for deep learning.
- Register a custom reference design to integrate the deep learning processor IP core.
- Specify your board and reference design by using the `TargetPlatform` and `ReferenceDesign` properties of the `dlhdl.ProcessorConfig` object.

See [registerDeepLearningMemoryAddressSpace](#), [registerDeepLearningTargetInterface](#), and [validateReferenceDesignForDeepLearning](#).

Deep learning processor streaming handshake modes

In R2021b, the generated deep learning processor IP core supports streaming handshaking modes. You can send multiple data frames to and receive multiple data frames from the deep learning

processor IP core by using buffer mode or streaming mode. See [Interface with the Deep Learning Processor IP Core](#).

Updates to estimatePerformance

Prior to R2021b, you could estimate performance of a network for only these bitstreams:

- `zcu102_single`
- `zcu102_int8`
- `zc706_single`
- `zc706_int8`
- `arria10soc_single`
- `arria10soc_int8`

In R2021b, you can estimate performance for your custom bitstream by using the `estimatePerformance` function. Create a processor configuration by using `dlhdl.ProcessorConfig`. Estimate performance by using the created processor configuration.

Estimate the performance of your network for multiple frames and for a bitstream by using the `FrameCount` name-value argument of the `estimatePerformance` function.

For more information, see `estimatePerformance`.

Updates to estimateResources

In R2021b, you can use the `estimateResources` function to:

- Estimate the resource usage for any Xilinx and Intel devices that have been registered by using a device registration function.
- Display the resource estimates as a percentage of the total resources for Xilinx devices.
- Retrieve the lookup table (LUT) utilization estimates for these devices:
 - Xilinx Zynq®-7000 ZC706
 - Intel Arria 10 SoC
 - Xilinx Zynq UltraScale+™ MPSoC ZCU102

Enhancements for quantization of directed acyclic graph (DAG) networks

In R2021b, Deep Learning HDL Toolbox supports quantization of:

- Addition layers that have more than two inputs.
- Addition layer followed by ReLU, Leaky ReLU, and Clipped ReLU layers.

Deep Learning HDL Toolbox supports channel-wise quantization for depth-wise separable convolution layers for improved accuracy of quantized network predictions.

Deep Learning HDL Toolbox supports quantization of these DAG networks:

- GoogLeNet
- MobileNet
- SqueezeNet

Network prototyping and validation without hardware

In R2021b, you can prototype, verify prediction accuracy, and retrieve intermediate layer-level results for your custom deep learning networks without the need for hardware. Create a simulation object by using the `dlhdl.Simulation` class. Verify network prediction accuracy by using the `prediction` function of the simulation object. Retrieve intermediate layer-level performance by using the `activations` function of the simulation object. The `dlhdl.Simulation` object accepts `single` data type networks, `int8` data type quantized networks, and `dlhdl.ProcessorConfig` objects as inputs. See `dlhdl.Simulator`.

Updated supported layers

Deep Learning HDL Toolbox now provides support for these layers:

- Softmax layer hardware implementation

For `int8` data type quantization, Deep Learning HDL Toolbox now provides support for these layers:

- Depth concatenation layer
- Softmax layer
- Addition layer followed by ReLU, leaky ReLU, or clipped ReLU layers

Functionality being removed or changed

KernelDataType property of `dlhdl.ProcessorConfig` object has been removed

Errors

This property has been removed. Use the `ProcessorDataType` property of the `dlhdl.ProcessorConfig` object instead.

LUT property of `estimateResources` function has been removed

Errors

This property has been removed as the `estimateResources` method reports LUT utilization by default.

R2021a

Version: 1.1

New Features

Compatibility Considerations

Custom directed acyclic graph (DAG) network support

Compile and deploy your custom DAG networks. Retrieve predictions from the deployed network by using MATLAB. For a list of supported networks, see Supported Networks, Layers, Boards, and Tools. The deep learning compiler analyzes the DAG network graph and generates the instructions, address mapping, and schedule to run the DAG network on the new deep learning processor. Deploy larger DAG networks onto FPGA boards with smaller resources by quantizing your DAG networks to use `int8` data types. See Quantization of Deep Neural Networks.

Performance estimation and profiling

Estimate performance by using the `estimatePerformance` function on the `dlhdl.ProcessorConfig` object before building your custom deep learning processor. For more information, see `estimatePerformance`. Retrieve the processor configuration of the shipping (reference) bitstream, by using the `dlhdl.ProcessorConfig` object. See `dlhdl.ProcessorConfig`. Perform design space exploration to find the deep learning processor configuration that fits your performance requirements by comparing the performance of your custom deep learning processor configuration to the performance of the shipping (reference) bitstream processor configuration.

You cannot estimate performance by using the `estimate` method for the `dlhdl.Workflow` object. For more information, see “Functionality being removed or changed” on page 5-3.

Resource estimation

Estimate resource utilization by using the `estimateResources` function on the `dlhdl.ProcessorConfig` object before building your custom deep learning processor. For more information, see `estimateResources`. Retrieve resource utilization of shipping (reference) bitstreams by using the `getBuildInfo` function on the `dlhdl.Workflow` object. See `getBuildInfo`. Perform design space exploration to find the deep learning processor configuration that fits your FPGA resource budget by comparing the resource utilization of your custom deep learning processor configuration to the resource utilization of the shipping (reference) bitstream.

Updated supported layers

Deep Learning HDL Toolbox now provides support for these layers:

- Addition layer
- Depth-wise separable convolution layer
- Depth concatenation layer

For `int8` data type quantization, Deep Learning HDL Toolbox now provides support for these layers:

- Average pooling layer
- Global average pooling layer
- Addition layer
- Clipped ReLU layer
- Leaky ReLU layer
- Depth-wise separable convolution layer

See Supported Networks, Layers, Boards, and Tools.

MATLAB Emulation for validate method of dlquantizer object

Validate the performance of your quantized network by comparing the prediction accuracy of your quantized network to that of your nonquantized network, without the need for hardware by using MATLAB emulation. See `validate`.

dlhdl.Workflow name-value argument pair update

For the list of name-value pair arguments that have been removed from `dlhdl.Workflow`, see “Functionality being removed or changed” on page 5-3.

Updates to supported software

Deep Learning HDL Toolbox has been tested with:

- Xilinx Vivado Design Suite 2020.1
- Intel Quartus Pro 18.1

Functionality being removed or changed

estimate function for dlhdl.Workflow object has been removed

Errors

This function has been removed.

'ProcessorConfig' option in dlhdl.Workflow has been removed

Errors

The 'ProcessorConfig' name-value pair for `dlhdl.Workflow` has been removed.

R2020b

Version: 1.0

New Features

Introducing Deep Learning HDL Toolbox: Prototype and implement deep learning networks on FPGAs and SoCs

With Deep Learning HDL Toolbox, you can prototype and implement deep learning networks on FPGAs and SoCs. Deploy and run deep learning networks on supported Xilinx and Intel FPGA and SoC devices. Improve deep learning network design, performance, and resource utilization by using profiling and estimating tools to explore tradeoffs and customize the network. Using HDL Coder™, you can generate HDL and an IP core to target FPGAs or SoCs.

Prototype on FPGAs

Use MATLAB and fixed bitstreams to compile, deploy, and run inference for pretrained series networks on target Intel and Xilinx FPGA and SoC boards. For more information, see Prototype Deep Learning Networks on FPGA.

Custom series network support

Compile and deploy your custom series networks using the same fixed-bitstreams as the pre-trained networks. For more information, see Prototype Deep Learning Networks on FPGA and SoCs Workflow.

Portable Verilog and VHDL code

Generate portable Verilog® and VHDL® code from your series deep learning network.

Tune user-configurable parameters

Customize your deep learning network implementation by tuning user-configurable parameters such as Thread Number, Input, and Output Memory Size. For more information, see Custom Processor Configuration Workflow.

Custom board support

Integrate the code generated from your customized design into your reference design for deploying to your custom board. For more information, see Generate Custom Processor IP.

Performance estimation and profiling

Gather layer-level latency and throughput estimates for your series networks. For more information, see estimate.

Hardware Support

Prototype and deploy deep learning networks to Intel and Xilinx FPGA boards. Use Ethernet based LIBIIO to rapidly deploy your series deep learning networks to your target Intel and Xilinx FPGA and SoC boards. For more information, see LIBIIO/Ethernet Connection Based Deployment.

Support Package for Intel FPGA and SoCs

You can use the Deep Learning HDL Toolbox Support Package for Intel FPGA and SoC Devices to communicate with, deploy series networks, and retrieve inference results from target Intel FPGA and SoC platforms. To download the support package, use the Add-on Explorer. For more information, see Deep Learning HDL Toolbox Support Package for Intel FPGA and SoC Devices.

Support Package for Xilinx FPGA and SoCs

You can use the Deep Learning HDL Toolbox Support Package for Xilinx FPGA and SoC Devices to communicate with, deploy series networks, and retrieve inference results from target Xilinx FPGA and SoC platforms. To download the support package, use the Add-on Explorer. For more information, see Deep Learning HDL Toolbox Support Package for Xilinx FPGA and SoC Devices.

