



CUDNN LIBRARY

DU-06702-001_v4.0 | February 2016

User Guide



Chapter 1.

INTRODUCTION

NVIDIA® cuDNN is a GPU-accelerated library of primitives for deep neural networks. It provides highly tuned implementations of routines arising frequently in DNN applications:

- ▶ Convolution forward and backward, including cross-correlation
- ▶ Pooling forward and backward
- ▶ Softmax forward and backward
- ▶ Neuron activations forward and backward:
 - ▶ Rectified linear (ReLU)
 - ▶ Sigmoid
 - ▶ Hyperbolic tangent (TANH)
- ▶ Tensor transformation functions
- ▶ LRN, LCN and batch normalization forward and backward

cuDNN's convolution routines aim for performance competitive with the fastest GEMM (matrix multiply) based implementations of such routines while using significantly less memory.

cuDNN features customizable data layouts, supporting flexible dimension ordering, striding, and subregions for the 4D tensors used as inputs and outputs to all of its routines. This flexibility allows easy integration into any neural network implementation and avoids the input/output transposition steps sometimes necessary with GEMM-based convolutions.

cuDNN offers a context-based API that allows for easy multithreading and (optional) interoperability with CUDA streams.

Chapter 2.

GENERAL DESCRIPTION

2.1. Programming Model

The cuDNN Library exposes a Host API but assumes that for operations using the GPU, the necessary data is directly accessible from the device.

An application using cuDNN must initialize a handle to the library context by calling `cudaDnnCreate()`. This handle is explicitly passed to every subsequent library function that operates on GPU data. Once the application finishes using cuDNN, it can release the resources associated with the library handle using `cudaDnnDestroy()`. This approach allows the user to explicitly control the library's functioning when using multiple host threads, GPUs and CUDA Streams. For example, an application can use `cudaSetDevice()` to associate different devices with different host threads and in each of those host threads, use a unique cuDNN handle which directs library calls to the device associated with it. cuDNN library calls made with different handles will thus automatically run on different devices. The device associated with a particular cuDNN context is assumed to remain unchanged between the corresponding `cudaDnnCreate()` and `cudaDnnDestroy()` calls. In order for the cuDNN library to use a different device within the same host thread, the application must set the new device to be used by calling `cudaSetDevice()` and then create another cuDNN context, which will be associated with the new device, by calling `cudaDnnCreate()`.

2.2. Notation

As of CUDNN v4 we have adopted a mathematically-inspired notation for layer inputs and outputs using \mathbf{x} , \mathbf{y} , \mathbf{dx} , \mathbf{dy} , \mathbf{b} , \mathbf{w} for common layer parameters. This was done to improve readability and ease of understanding of parameters meaning. All layers now follow a uniform convention that during inference

$\mathbf{x} = \text{layerFunction}(\mathbf{y}, \text{otherParams})$.

And during backpropagation

$(\mathbf{dx}, \text{dOtherParams}) = \text{layerFunctionGradient}(\mathbf{x}, \mathbf{y}, \mathbf{dy}, \text{otherParams})$

For convolution the notation is

$$\mathbf{y} = \mathbf{x} * \mathbf{w} + \mathbf{b}$$

where \mathbf{w} is the matrix of filter weights, \mathbf{x} is the previous layer's data (during inference), \mathbf{y} is the next layer's data, \mathbf{b} is the bias and $*$ is the convolution operator. In backpropagation routines the parameters keep their meanings. $d\mathbf{x}$, $d\mathbf{y}$, $d\mathbf{w}$, $d\mathbf{b}$ always refer to the gradient of the final network error function with respect to a given parameter. So $d\mathbf{y}$ in all backpropagation routines always refers to error gradient backpropagated through the network computation graph so far. Similarly other parameters in more specialized layers, such as, for instance, $d\mathbf{Means}$ or $d\mathbf{BnBias}$ refer to gradients of the loss function wrt those parameters.



w is used in the API for both the width of the \mathbf{x} tensor and convolution filter matrix. To resolve this ambiguity we use w and `filter` notation interchangeably for convolution filter weight matrix. The meaning is clear from the context since the layer width is always referenced near its height.

2.3. Tensor Descriptor

The cuDNN Library describes data holding images, videos and any other data with contents with a generic n-D tensor defined with the following parameters :

- ▶ a dimension `dim` from 3 to 8
- ▶ a data type (32-bit floating point, 64 bit-floating point, 16 bit floating point...)
- ▶ `dim` integers defining the size of each dimension
- ▶ `dim` integers defining the stride of each dimension (e.g the number of elements to add to reach the next element from the same dimension)

The first two dimensions define respectively the batch number `n` and the number of features maps `c`. This tensor definition allows for example to have some dimensions overlapping each others within the same tensor by having the stride of one dimension smaller than the product of the dimension and the stride of the next dimension. In cuDNN, unless specified otherwise, all routines will support tensors with overlapping dimensions for forward pass input tensors, however, dimensions of the output tensors cannot overlap. Even though this tensor format supports negative strides (which can be useful for data mirroring), cuDNN routines do not support tensors with negative strides unless specified otherwise.

2.3.1. WXYZ Tensor Descriptor

Tensor descriptor formats are identified using acronyms, with each letter referencing a corresponding dimension. In this document, the usage of this terminology implies :

- ▶ all the strides are strictly positive
- ▶ the dimensions referenced by the letters are sorted in decreasing order of their respective strides

2.3.2. 4-D Tensor Descriptor

A 4-D Tensor descriptor is used to define the format for batches of 2D images with 4 letters : N,C,H,W for respectively the batch number, the number of feature maps, the height and the width. The letters are sorted in decreasing order of the strides. The commonly used 4-D tensor formats are :

- ▶ NCHW
- ▶ NHWC
- ▶ CHWN

2.3.3. 5-D Tensor Description

A 5-D Tensor descriptor is used to define the format of batch of 3D images with 5 letters : N,C,D,H,W for respectively the batch number, the number of feature maps, the depth, the height and the width. The letters are sorted in decreasing order of the strides. The commonly used 5-D tensor formats are called :

- ▶ NCDHW
- ▶ NDHWC
- ▶ CDHWN

2.3.4. Fully-packed tensors

A tensor is defined as **XYZ-fully-packed** if and only if :

- ▶ the number of tensor dimensions is equal to the number of letters preceding the **fully-packed** suffix.
- ▶ the stride of the i -th dimension is equal to the product of the $(i+1)$ -th dimension by the $(i+1)$ -th stride.
- ▶ the stride of the last dimension is 1.

2.3.5. Partially-packed tensors

The partially 'XYZ-packed' terminology only applies in a context of a tensor format described with a superset of the letters used to define a partially-packed tensor. A WXYZ tensor is defined as **XYZ-packed** if and only if :

- ▶ the strides of all dimensions NOT referenced in the -packed suffix are greater or equal to the product of the next dimension by the next stride.
- ▶ the stride of each dimension referenced in the -packed suffix in position i is equal to the product of the $(i+1)$ -st dimension by the $(i+1)$ -st stride.
- ▶ if last tensor's dimension is present in the -packed suffix, it's stride is 1.

For example a NHWC tensor WC-packed means that the `c_stride` is equal to 1 and `w_stride` is equal to `c_dim` x `c_stride`. In practice, the -packed suffix is usually with slowest changing dimensions of a tensor but it is also possible to refer to a NCHW tensor that is only N-packed.

2.3.6. Spatially packed tensors

Spatially-packed tensors are defined as partially-packed in spatial dimensions.

For example a spatially-packed 4D tensor would mean that the tensor is either NCHW HW-packed or CNHW HW-packed.

2.3.7. Overlapping tensors

A tensor is defined to be overlapping if a iterating over a full range of dimensions produces the same address more than once.

In practice an overlapped tensor will have $\text{stride}[i-1] < \text{stride}[i] * \text{dim}[i]$ for some of the i from $[1, \text{nbDims}]$ interval.

2.4. Thread Safety

The library is thread safe and its functions can be called from multiple host threads, even with the same handle. When sharing a handle across host threads, extreme care needs to be taken to ensure that any changes to the handle configuration in one thread do not adversely affect cuDNN function calls in others. This is especially true for the destruction of the handle. It is not recommended that multiple threads share the same cuDNN handle.

2.5. Reproducibility (determinism)

By design, most of cuDNN's routines from a given version generate the same bit-wise results at every run when executed on GPUs with the same architecture and the same number of SMs. However, bit-wise reproducibility is not guaranteed across versions, as the implementation of a given routine may change. With the current release, the following routines do not guarantee reproducibility because they use atomic add operations:

- ▶ `cudaDnnConvolutionBackwardFilter` when `CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0` or `CUDNN_CONVOLUTION_BWD_FILTER_ALGO_3` is used
- ▶ `cudaDnnConvolutionBackwardData` when `CUDNN_CONVOLUTION_BWD_DATA_ALGO_0` is used
- ▶ `cudaDnnPoolingBackward` when `CUDNN_POOLING_MAX` is used

2.6. Scaling parameters `alpha` and `beta`

Many cuDNN routines like `cudaDnnConvolutionForward` take pointers to scaling factors (in host memory) used to blend the source value with prior value in the destination tensor as follows: $\text{dstValue} = \text{alpha}[0] * \text{srcValue} + \text{beta}[0] * \text{priorDstValue}$. When `beta[0]` is zero, the output is not read and can contain any uninitialized data (including NaN). The

storage data type for alpha[0], beta[0] is float for HALF and SINGLE tensors, and double for DOUBLE tensors. These parameters are passed using a host memory pointer.



For improved performance it is advised to use beta[0] = 0.0. Use a non-zero value for beta[0] only when blending with prior values stored in the output tensor is needed.

2.7. GPU and driver requirements

cuDNN v4.0 supports NVIDIA GPUs of compute capability 3.0 and higher and requires an NVIDIA Driver compatible with CUDA Toolkit 7.0.

2.8. Backward compatibility and deprecation policy

When we want to change the API of an already implemented functionality, we will create a new routine with the same name appended with a suffix `_v<n>` where `n` represents the cuDNN version where that new API was introduced. This way, we can guarantee backward compatibility with the previous version of cuDNN Library. However, in order to facilitate the maintenance and keep the API as clean as possible, we want to be able to deprecate and then eliminate the previous API entries and eventually remove also the new suffixed ones to only keep the unsuffixed entries mapped to the new API. To achieve this, we introduce the following deprecation policy:

- ▶ the former API entries will be renamed with a suffix `_v<f>` where `f` represents a cuDNN version anterior to the new API.
- ▶ the unsuffixed API entry will be mapped to the deprecated entry until the former API entry with an anterior suffix `_v<f>` appeared in the API header file.
- ▶ the deprecated former API entry with an anterior suffix `_v<f>` and new API entry with suffix `_v<n>` will be present for at least one release
- ▶ In one of the subsequent cuDNN release, the suffixed entries of a given entry will be removed simultaneously.

To sum up, when a routine appears in two forms, one with a suffix and one with no suffix, it means that the non-suffixed entry is deprecated. In this case, it is strongly advised to users to use the new suffixed API entry to guarantee backwards compatibility with the next cuDNN release. When a routine appears with multiples suffixes, the unsuffixed API entry will be mapped to the higher number. In that case it is strongly advised to use the non-suffixed API entry to guarantee backward compatibility with the next cuDNN release.

Chapter 3.

CUDNN DATATYPES REFERENCE

This chapter describes all the types and enums of the cuDNN library API.

3.1. cudnnHandle_t

cudnnHandle_t is a pointer to an opaque structure holding the cuDNN library context. The cuDNN library context must be created using **cudnnCreate()** and the returned handle must be passed to all subsequent library function calls. The context should be destroyed at the end using **cudnnDestroy()**. The context is associated with only one GPU device, the current device at the time of the call to **cudnnCreate()**. However multiple contexts can be created on the same GPU device.

3.2. cudnnStatus_t

cudnnStatus_t is an enumerated type used for function status returns. All cuDNN library functions return their status, which can be one of the following values:

Value	Meaning
CUDNN_STATUS_SUCCESS	The operation completed successfully.
CUDNN_STATUS_NOT_INITIALIZED	The cuDNN library was not initialized properly. This error is usually returned when a call to cudnnCreate() fails or when cudnnCreate() has not been called prior to calling another cuDNN routine. In the former case, it is usually due to an error in the CUDA Runtime API called by cudnnCreate() or by an error in the hardware setup.
CUDNN_STATUS_ALLOC_FAILED	Resource allocation failed inside the cuDNN library. This is usually caused by an internal cudaMalloc() failure. To correct: prior to the function call, deallocate previously allocated memory as much as possible.

Value	Meaning
<code>CUDNN_STATUS_BAD_PARAM</code>	An incorrect value or parameter was passed to the function. To correct: ensure that all the parameters being passed have valid values.
<code>CUDNN_STATUS_ARCH_MISMATCH</code>	The function requires a feature absent from the current GPU device. Note that cuDNN only supports devices with compute capabilities greater than or equal to 3.0. To correct: compile and run the application on a device with appropriate compute capability.
<code>CUDNN_STATUS_MAPPING_ERROR</code>	An access to GPU memory space failed, which is usually caused by a failure to bind a texture. To correct: prior to the function call, unbind any previously bound textures. Otherwise, this may indicate an internal error/bug in the library.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The GPU program failed to execute. This is usually caused by a failure to launch some cuDNN kernel on the GPU, which can occur for multiple reasons. To correct: check that the hardware, an appropriate version of the driver, and the cuDNN library are correctly installed. Otherwise, this may indicate a internal error/bug in the library.
<code>CUDNN_STATUS_INTERNAL_ERROR</code>	An internal cuDNN operation failed.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The functionality requested is not presently supported by cuDNN.
<code>CUDNN_STATUS_LICENSE_ERROR</code>	The functionality requested requires some license and an error was detected when trying to check the current licensing. This error can happen if the license is not present or is expired or if the environment variable <code>NVIDIA_LICENSE_FILE</code> is not set properly.

3.3. `cudaTensorDescriptor_t`

`cudaCreateTensorDescriptor_t` is a pointer to an opaque structure holding the description of a generic n-D dataset. `cudaCreateTensorDescriptor()` is used to create one instance, and one of the routines `cudaSetTensorNdDescriptor()`, `cudaSetTensor4dDescriptor()` or `cudaSetTensor4dDescriptorEx()` must be used to initialize this instance.

3.4. cudnnFilterDescriptor_t

`cudnnFilterDescriptor_t` is a pointer to an opaque structure holding the description of a filter dataset. `cudnnCreateFilterDescriptor()` is used to create one instance, and `cudnnSetFilterDescriptor()` must be used to initialize this instance.

3.5. cudnnConvolutionDescriptor_t

`cudnnConvolutionDescriptor_t` is a pointer to an opaque structure holding the description of a convolution operation. `cudnnCreateConvolutionDescriptor()` is used to create one instance, and `cudnnSetConvolutionNdDescriptor()` or `cudnnSetConvolution2dDescriptor()` must be used to initialize this instance.

3.6. cudnnNanPropagation_t

`cudnnNanPropagation_t` is an enumerated type used to indicate if some routines should propagate **Nan** numbers. This enumerated type is used as a field for the `cudnnActivationDescriptor_t` descriptor and `cudnnPoolingDescriptor_t` descriptor.

Value	Meaning
<code>CUDNN_NOT_PROPAGATE_NAN</code>	Nan numbers are not propagated
<code>CUDNN_PROPAGATE_NAN</code>	Nan numbers are propagated

3.7. cudnnPoolingDescriptor_t

`cudnnPoolingDescriptor_t` is a pointer to an opaque structure holding the description of a pooling operation. `cudnnCreatePoolingDescriptor()` is used to create one instance, and `cudnnSetPoolingNdDescriptor()` or `cudnnSetPooling2dDescriptor()` must be used to initialize this instance.

3.8. cudnnDataType_t

`cudnnDataType_t` is an enumerated type indicating the data type to which a tensor descriptor or filter descriptor refers.

Value	Meaning
<code>CUDNN_DATA_FLOAT</code>	The data is 32-bit single-precision floating point (<code>float</code>).
<code>CUDNN_DATA_DOUBLE</code>	The data is 64-bit double-precision floating point (<code>double</code>).

Value	Meaning
CUDNN_DATA_HALF	The data is 16-bit floating point.

3.9. cudnnTensorFormat_t

`cudnnTensorFormat_t` is an enumerated type used by `cudnnSetTensor4dDescriptor()` to create a tensor with a pre-defined layout.

Value	Meaning
CUDNN_TENSOR_NCHW	This tensor format specifies that the data is laid out in the following order: image, features map, rows, columns. The strides are implicitly defined in such a way that the data are contiguous in memory with no padding between images, feature maps, rows, and columns; the columns are the inner dimension and the images are the outermost dimension.
CUDNN_TENSOR_NHWC	This tensor format specifies that the data is laid out in the following order: image, rows, columns, features maps. The strides are implicitly defined in such a way that the data are contiguous in memory with no padding between images, rows, columns, and features maps; the feature maps are the inner dimension and the images are the outermost dimension.

3.10. cudnnAddMode_t

`cudnnAddMode_t` is an enumerated type used by `cudnnAddTensor()` to specify how a bias tensor is added to an input/output tensor.

Value	Meaning
CUDNN_ADD_IMAGE or CUDNN_ADD_SAME_HW	In this mode, the bias tensor is defined as one image with one feature map. This image will be added to every feature map of every image of the input/output tensor.
CUDNN_ADD_FEATURE_MAP or CUDNN_ADD_SAME_CHW	In this mode, the bias tensor is defined as one image with multiple feature maps. This image will be added to every image of the input/output tensor.
CUDNN_ADD_SAME_C	In this mode, the bias tensor is defined as one image with multiple feature maps of dimension 1x1; it can be seen as an vector of feature maps. Each feature map of the bias tensor will be added to the corresponding feature map of all height-by-width pixels of every image of the input/output tensor.

Value	Meaning
CUDNN_ADD_FULL_TENSOR	In this mode, the bias tensor has the same dimensions as the input/output tensor. It will be added point-wise to the input/output tensor.

3.11. cudnnConvolutionMode_t

`cudnnConvolutionMode_t` is an enumerated type used by `cudnnSetConvolutionDescriptor()` to configure a convolution descriptor. The filter used for the convolution can be applied in two different ways, corresponding mathematically to a convolution or to a cross-correlation. (A cross-correlation is equivalent to a convolution with its filter rotated by 180 degrees.)

Value	Meaning
CUDNN_CONVOLUTION	In this mode, a convolution operation will be done when applying the filter to the images.
CUDNN_CROSS_CORRELATION	In this mode, a cross-correlation operation will be done when applying the filter to the images.

3.12. cudnnConvolutionFwdPreference_t

`cudnnConvolutionFwdPreference_t` is an enumerated type used by `cudnnGetConvolutionForwardAlgorithm()` to help the choice of the algorithm used for the forward convolution.

Value	Meaning
CUDNN_CONVOLUTION_FWD_NO_WORKSPACE	In this configuration, the routine <code>cudnnGetConvolutionForwardAlgorithm()</code> is guaranteed to return an algorithm that does not require any extra workspace to be provided by the user.
CUDNN_CONVOLUTION_FWD_PREFER_FASTEST	In this configuration, the routine <code>cudnnGetConvolutionForwardAlgorithm()</code> will return the fastest algorithm regardless how much workspace is needed to execute it.
CUDNN_CONVOLUTION_FWD_SPECIFY_WORKSPACE_LIMIT	In this configuration, the routine <code>cudnnGetConvolutionForwardAlgorithm()</code> will return the fastest algorithm that fits within the memory limit that the user provided.

3.13. cudnnConvolutionFwdAlgo_t

`cudnnConvolutionFwdAlgo_t` is an enumerated type that exposes the different algorithms available to execute the forward convolution operation.

Value	Meaning
<code>CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_GEMM</code>	This algorithm expresses the convolution as a matrix product without actually explicitly form the matrix that holds the input tensor data.
<code>CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_PRECOMPUTED_GEMM</code>	This algorithm expresses the convolution as a matrix product without actually explicitly form the matrix that holds the input tensor data, but still needs some memory workspace to precompute some indices in order to facilitate the implicit construction of the matrix that holds the input tensor data
<code>CUDNN_CONVOLUTION_FWD_ALGO_GEMM</code>	This algorithm expresses the convolution as an explicit matrix product. A significant memory workspace is needed to store the matrix that holds the input tensor data.
<code>CUDNN_CONVOLUTION_FWD_ALGO_DIRECT</code>	This algorithm expresses the convolution as a direct convolution (e.g without implicitly or explicitly doing a matrix multiplication).
<code>CUDNN_CONVOLUTION_FWD_ALGO_FFT</code>	This algorithm uses a Fast-Fourier Transform approach to compute the convolution. A significant memory workspace is needed to store intermediate results.
<code>CUDNN_CONVOLUTION_FWD_ALGO_FFT_TILING</code>	This algorithm uses a Fast-Fourier Transform approach but splits the inputs into 32x32 tiles. A significant memory workspace is needed to store intermediate results but significantly less than <code>CUDNN_CONVOLUTION_FWD_ALGO_FFT</code> for big size images.

3.14. `cudaConvolutionFwdAlgoPerf_t`

`cudaConvolutionFwdAlgoPerf_t` is a structure containing performance results returned by `cudaFindConvolutionForwardAlgorithm()`.

Member Name	Explanation
<code>cudaConvolutionFwdAlgo_t algo</code>	The algorithm run to obtain the associated performance metrics.
<code>cudaStatus_t status</code>	<p>If any error occurs during the workspace allocation or timing of <code>cudaConvolutionForward()</code>, this status will represent that error. Otherwise, this status will be the return status of <code>cudaConvolutionForward()</code>.</p> <ul style="list-style-type: none"> ▶ <code>CUDNN_STATUS_ALLOC_FAILED</code> if any error occurred during workspace allocation. ▶ <code>CUDNN_STATUS_INTERNAL_ERROR</code> if any error occurred during timing calculations or workspace deallocation.

Member Name	Explanation
	▶ Otherwise, this will be the return status of <code>cudaConvolutionForward()</code> .
<code>float time</code>	The execution time of <code>cudaConvolutionForward()</code> (in milliseconds).
<code>size_t memory</code>	The workspace size (in bytes).

3.15. `cudaConvolutionBwdFilterPreference_t`

`cudaConvolutionBwdFilterPreference_t` is an enumerated type used by `cudaGetConvolutionBackwardFilterAlgorithm()` to help the choice of the algorithm used for the backward filter convolution.

Value	Meaning
<code>CUDA_CONVOLUTION_BWD_FILTER_NO_WORKSPACE</code>	In this configuration, the routine <code>cudaGetConvolutionBackwardFilterAlgorithm()</code> is guaranteed to return an algorithm that does not require any extra workspace to be provided by the user.
<code>CUDA_CONVOLUTION_BWD_FILTER_PREFER_FASTEST</code>	In this configuration, the routine <code>cudaGetConvolutionBackwardFilterAlgorithm()</code> will return the fastest algorithm regardless how much workspace is needed to execute it.
<code>CUDA_CONVOLUTION_BWD_FILTER_SPECIFY_WORKSPACE</code>	In this configuration, the routine <code>cudaGetConvolutionBackwardFilterAlgorithm()</code> will return the fastest algorithm that fits within the memory limit that the user provided.

3.16. `cudaConvolutionBwdFilterAlgo_t`

`cudaConvolutionBwdFilterAlgo_t` is an enumerated type that exposes the different algorithms available to execute the backward filter convolution operation.

Value	Meaning
<code>CUDA_CONVOLUTION_BWD_FILTER_ALGO_0</code>	This algorithm expresses the convolution as a sum of matrix product without actually explicitly form the matrix that holds the input tensor data. The sum is done using atomic adds operation, thus the results are non-deterministic.
<code>CUDA_CONVOLUTION_BWD_FILTER_ALGO_1</code>	This algorithm expresses the convolution as a matrix product without actually explicitly form the matrix that holds the input tensor data. The results are deterministic.
<code>CUDA_CONVOLUTION_BWD_FILTER_ALGO_FFT</code>	This algorithm uses a Fast-Fourier Transform approach to compute the convolution. A significant memory workspace is needed to

Value	Meaning
	store intermediate results. The results are deterministic.
<code>CUDNN_CONVOLUTION_BWD_FILTER_ALGO_3</code>	This algorithm is similar to <code>CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0</code> but uses some small workspace to precompute some indices. The results are also non-deterministic.

3.17. `cudaConvolutionBwdFilterAlgoPerf_t`

`cudaConvolutionBwdFilterAlgoPerf_t` is a structure containing performance results returned by `cudaFindConvolutionBackwardFilterAlgorithm()`.

Member Name	Explanation
<code>cudaConvolutionBwdFilterAlgo_t algo</code>	The algorithm run to obtain the associated performance metrics.
<code>cudaStatus_t status</code>	If any error occurs during the workspace allocation or timing of <code>cudaConvolutionBackwardFilter_v3()</code> , this status will represent that error. Otherwise, this status will be the return status of <code>cudaConvolutionBackwardFilter_v3()</code> . <ul style="list-style-type: none"> ▶ <code>CUDNN_STATUS_ALLOC_FAILED</code> if any error occurred during workspace allocation. ▶ <code>CUDNN_STATUS_INTERNAL_ERROR</code> if any error occurred during timing calculations or workspace deallocation. ▶ Otherwise, this will be the return status of <code>cudaConvolutionBackwardFilter_v3()</code>.
<code>float time</code>	The execution time of <code>cudaConvolutionBackwardFilter_v3()</code> (in milliseconds).
<code>size_t memory</code>	The workspace size (in bytes).

3.18. `cudaConvolutionBwdDataPreference_t`

`cudaConvolutionBwdDataPreference_t` is an enumerated type used by `cudaGetConvolutionBackwardDataAlgorithm()` to help the choice of the algorithm used for the backward data convolution.

Value	Meaning
<code>CUDNN_CONVOLUTION_BWD_DATA_NO_WORKSPACE</code>	In this configuration, the routine <code>cudaGetConvolutionBackwardDataAlgorithm()</code> is guaranteed to return an algorithm that does not require any extra workspace to be provided by the user.

Value	Meaning
<code>CUDNN_CONVOLUTION_BWD_DATA_PREFER_FASTEST</code>	In this configuration, the routine <code>cudaGetConvolutionBackwardDataAlgorithm()</code> will return the fastest algorithm regardless how much workspace is needed to execute it.
<code>CUDNN_CONVOLUTION_BWD_DATA_SPECIFY_WORKSPACE_LIMIT</code>	In this configuration, the routine <code>cudaGetConvolutionBackwardDataAlgorithm()</code> will return the fastest algorithm that fits within the memory limit that the user provided.

3.19. `cudaConvolutionBwdDataAlgo_t`

`cudaConvolutionBwdDataAlgo_t` is an enumerated type that exposes the different algorithms available to execute the backward data convolution operation.

Value	Meaning
<code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_0</code>	This algorithm expresses the convolution as a sum of matrix product without actually explicitly form the matrix that holds the input tensor data. The sum is done using atomic adds operation, thus the results are non-deterministic.
<code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_1</code>	This algorithm expresses the convolution as a matrix product without actually explicitly form the matrix that holds the input tensor data. The results are deterministic.
<code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_FFT</code>	This algorithm uses a Fast-Fourier Transform approach to compute the convolution. A significant memory workspace is needed to store intermediate results. The results are deterministic.
<code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_FFT_TILING</code>	This algorithm uses a Fast-Fourier Transform approach but splits the inputs into 32x32 tiles. A significant memory workspace is needed to store intermediate results but significantly less than <code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_FFT</code> for big size images. A significant memory workspace is needed to store intermediate results. The results are deterministic.

3.20. `cudaConvolutionBwdDataAlgoPerf_t`

`cudaConvolutionBwdDataAlgoPerf_t` is a structure containing performance results returned by `cudaFindConvolutionBackwardDataAlgorithm()`.

Member Name	Explanation
<code>cudaConvolutionBwdDataAlgo_t algo</code>	The algorithm run to obtain the associated performance metrics.

Member Name	Explanation
<code>cudaStatus_t status</code>	<p>If any error occurs during the workspace allocation or timing of <code>cudaConvolutionBackwardData_v3()</code>, this status will represent that error. Otherwise, this status will be the return status of <code>cudaConvolutionBackwardData_v3()</code>.</p> <ul style="list-style-type: none"> ▶ <code>CUDA_STATUS_ALLOC_FAILED</code> if any error occurred during workspace allocation. ▶ <code>CUDA_STATUS_INTERNAL_ERROR</code> if any error occurred during timing calculations or workspace deallocation. ▶ Otherwise, this will be the return status of <code>cudaConvolutionBackwardData_v3()</code>.
<code>float time</code>	The execution time of <code>cudaConvolutionBackwardData_v3()</code> (in milliseconds).
<code>size_t memory</code>	The workspace size (in bytes).

3.21. `cudaSoftmaxAlgorithm_t`

`cudaSoftmaxAlgorithm_t` is used to select an implementation of the softmax function used in `cudaSoftmaxForward()` and `cudaSoftmaxBackward()`.

Value	Meaning
<code>CUDA_SOFTMAX_FAST</code>	This implementation applies the straightforward softmax operation.
<code>CUDA_SOFTMAX_ACCURATE</code>	This implementation scales each point of the softmax input domain by its maximum value to avoid potential floating point overflows in the softmax evaluation.
<code>CUDA_SOFTMAX_LOG</code>	This entry performs the Log softmax operation, avoiding overflows by scaling each point in the input domain as in <code>CUDA_SOFTMAX_ACCURATE</code>

3.22. `cudaSoftmaxMode_t`

`cudaSoftmaxMode_t` is used to select over which data the `cudaSoftmaxForward()` and `cudaSoftmaxBackward()` are computing their results.

Value	Meaning
<code>CUDA_SOFTMAX_MODE_INSTANCE</code>	The softmax operation is computed per image (N) across the dimensions C,H,W.

Value	Meaning
CUDNN_SOFTMAX_MODE_CHANNEL	The softmax operation is computed per spatial location (H,W) per image (N) across the dimension C.

3.23. cudnnPoolingMode_t

`cudnnPoolingMode_t` is an enumerated type passed to `cudnnSetPoolingDescriptor()` to select the pooling method to be used by `cudnnPoolingForward()` and `cudnnPoolingBackward()`.

Value	Meaning
CUDNN_POOLING_MAX	The maximum value inside the pooling window will be used.
CUDNN_POOLING_AVERAGE_COUNT_INCLUDE_PADDING	The values inside the pooling window will be averaged. The number of padded values will be taken into account when computing the average value
CUDNN_POOLING_AVERAGE_COUNT_EXCLUDE_PADDING	The values inside the pooling window will be averaged. The number of padded values will not be taken into account when computing the average value

3.24. cudnnActivationMode_t

`cudnnActivationMode_t` is an enumerated type used to select the neuron activation function used in `cudnnActivationForward()` and `cudnnActivationBackward()`.

Value	Meaning
CUDNN_ACTIVATION_SIGMOID	Selects the sigmoid function.
CUDNN_ACTIVATION_RELU	Selects the rectified linear function.
CUDNN_ACTIVATION_TANH	Selects the hyperbolic tangent function.
CUDNN_ACTIVATION_CLIPPED_RELU	Selects the clipped rectified linear function

3.25. cudnnLRNMode_t

`cudnnLRNMode_t` is an enumerated type used to specify the mode of operation in `cudnnLRNCrossChannelForward()` and `cudnnLRNCrossChannelBackward()`.

Value	Meaning
CUDNN_LRN_CROSS_CHANNEL_DIM1	LRN computation is performed across tensor's dimension dimA[1].

3.26. cudnnDivNormMode_t

cudnnDivNormMode_t is an enumerated type used to specify the mode of operation in **cudnnDivisiveNormalizationForward()** and **cudnnDivisiveNormalizationBackward()**.

Value	Meaning
CUDNN_DIVNORM_PRECOMPUTED_MEANS	The means tensor data pointer is expected to contain means or other kernel convolution values precomputed by the user. The means pointer can also be NULL, in that case it's considered to be filled with zeroes. This is equivalent to spatial LRN. Note that in the backward pass the means are treated as independent inputs and the gradient over means is computed independently. In this mode to yield a net gradient over the entire LCN computational graph the destDiffMeans result should be backpropagated through the user's means layer (which can be implemented using average pooling) and added to the destDiffData tensor produced by cudnnDivisiveNormalizationBackward.

3.27. cudnnBatchNormMode_t

cudnnBatchNormMode_t is an enumerated type used to specify the mode of operation in **cudnnBatchNormalizationForwardInference()**, **cudnnBatchNormalizationForwardTraining()**, **cudnnBatchNormalizationBackward()** and **cudnnDeriveBNTensorDescriptor()** routines.

Value	Meaning
CUDNN_BATCHNORM_PER_ACTIVATION	Normalization is performed per-activation. This mode is intended to be used after non-convolutional network layers. In this mode bnBias and bnScale tensor dimensions are 1xCxHxW.
CUDNN_BATCHNORM_SPATIAL	Normalization is performed over N+spatial dimensions. This mode is intended for use after convolutional layers (where spatial invariance is desired). In this mode bnBias, bnScale tensor dimensions are 1xCx1x1.

Chapter 4.

CUDNN API REFERENCE

This chapter describes the API of all the routines of the cuDNN library.

4.1. cudnnGetVersion

```
size_t cudnnGetVersion()
```

This function returns the version number of the cuDNN Library. It returns the **CUDNN_VERSION** define present in the cudnn.h header file. Starting with release R2, the routine can be used to identify dynamically the current cuDNN Library used by the application. The define **CUDNN_VERSION** can be used to have the same application linked against different cuDNN versions using conditional compilation statements.

4.2. cudnnGetErrorString

```
const char * cudnnGetErrorString(cudnnStatus_t status)
```

This function returns a human-readable character string describing the **cudnnStatus_t** enumerate passed as input parameter.

4.3. cudnnCreate

```
cudnnStatus_t cudnnCreate(cudnnHandle_t *handle)
```

This function initializes the cuDNN library and creates a handle to an opaque structure holding the cuDNN library context. It allocates hardware resources on the host and device and must be called prior to making any other cuDNN library calls. The cuDNN library context is tied to the current CUDA device. To use the library on multiple devices, one cuDNN handle needs to be created for each device. For a given device, multiple cuDNN handles with different configurations (e.g., different current CUDA streams) may be created. Because **cudnnCreate** allocates some internal resources, the release of those resources by calling **cudnnDestroy** will implicitly call **cudaDeviceSynchronize**; therefore, the recommended best practice is to call **cudnnCreate/cudnnDestroy** outside of performance-critical code paths. For multithreaded applications that use the same device from different threads, the

recommended programming model is to create one (or a few, as is convenient) cuDNN handle(s) per thread and use that cuDNN handle for the entire life of the thread.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The initialization succeeded.
<code>CUDNN_STATUS_NOT_INITIALIZED</code>	CUDA Runtime API initialization failed.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	The resources could not be allocated.

4.4. cudnnDestroy

```
cudaError_t cudnnDestroy(cudaDnnHandle_t handle)
```

This function releases hardware resources used by the cuDNN library. This function is usually the last call with a particular handle to the cuDNN library. Because `cudnnCreate` allocates some internal resources, the release of those resources by calling `cudnnDestroy` will implicitly call `cudaDeviceSynchronize`; therefore, the recommended best practice is to call `cudnnCreate/cudnnDestroy` outside of performance-critical code paths.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The cuDNN context destruction was successful.
<code>CUDNN_STATUS_NOT_INITIALIZED</code>	The library was not initialized.

4.5. cudnnSetStream

```
cudaError_t cudnnSetStream(cudaDnnHandle_t handle, cudaStream_t streamId)
```

This function sets the cuDNN library stream, which will be used to execute all subsequent calls to the cuDNN library functions with that particular handle. If the cuDNN library stream is not set, all kernels use the default (**NULL**) stream. In particular, this routine can be used to change the stream between kernel launches and then to reset the cuDNN library stream back to **NULL**.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The stream was set successfully.

4.6. cudnnGetStream

```
cudaError_t cudnnGetStream(cudaDnnHandle_t handle, cudaStream_t *streamId)
```

This function gets the cuDNN library stream, which is being used to execute all calls to the cuDNN library functions. If the cuDNN library stream is not set, all kernels use the *default* **NULL** stream.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The stream was returned successfully.

4.7. cudnnCreateTensorDescriptor

```
cudaStatus_t cudnnCreateTensorDescriptor(cudaTensorDescriptor_t *tensorDesc)
```

This function creates a generic Tensor descriptor object by allocating the memory needed to hold its opaque structure. The data is initialized to be all zero.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was created successfully.
CUDNN_STATUS_ALLOC_FAILED	The resources could not be allocated.

4.8. cudnnSetTensor4dDescriptor

```
cudaStatus_t
cudnnSetTensor4dDescriptor(cudaTensorDescriptor_t tensorDesc,
                           cudaTensorFormat_t format,
                           cudaDataType_t dataType,
                           int n,
                           int c,
                           int h,
                           int w )
```

This function initializes a previously created generic Tensor descriptor object into a 4D tensor. The strides of the four dimensions are inferred from the format parameter and set in such a way that the data is contiguous in memory with no padding between dimensions.



The total size of a tensor including the potential padding between dimensions is limited to 2 Giga-elements of type `dataType`.

Param	In/out	Meaning
tensorDesc	input/output	Handle to a previously created tensor descriptor.
format	input	Type of format.
datatype	input	Data type.
n	input	Number of images.
c	input	Number of feature maps per image.
h	input	Height of each feature map.
w	input	Width of each feature map.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the parameters <code>n</code> , <code>c</code> , <code>h</code> , <code>w</code> was negative or <code>format</code> has an invalid enumerant value or <code>dataType</code> has an invalid enumerant value.
CUDNN_STATUS_NOT_SUPPORTED	The total size of the tensor descriptor exceeds the maximim limit of 2 Giga-elements.

4.9. cudnnSetTensor4dDescriptorEx

```

cudnnStatus_t
cudnnSetTensor4dDescriptorEx( cudnnTensorDescriptor_t tensorDesc,
                             cudnnDataType_t dataType,
                             int n,
                             int c,
                             int h,
                             int w,
                             int nStride,
                             int cStride,
                             int hStride,
                             int wStride );

```

This function initializes a previously created generic Tensor descriptor object into a 4D tensor, similarly to `cudnnSetTensor4dDescriptor` but with the strides explicitly passed as parameters. This can be used to lay out the 4D tensor in any order or simply to define gaps between dimensions.



At present, some cuDNN routines have limited support for strides; Those routines will return CUDNN_STATUS_NOT_SUPPORTED if a Tensor4D object with an unsupported stride is used. `cudnnTransformTensor` can be used to convert the data to a supported layout.



The total size of a tensor including the potential padding between dimensions is limited to 2 Giga-elements of type `dataType`.

Param	In/out	Meaning
tensorDesc	input/ output	Handle to a previously created tensor descriptor.
datatype	input	Data type.
n	input	Number of images.
c	input	Number of feature maps per image.
h	input	Height of each feature map.
w	input	Width of each feature map.
nStride	input	Stride between two consecutive images.

Param	In/out	Meaning
cStride	input	Stride between two consecutive feature maps.
hStride	input	Stride between two consecutive rows.
wStride	input	Stride between two consecutive columns.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the parameters <code>n</code> , <code>c</code> , <code>h</code> , <code>w</code> or <code>nStride</code> , <code>cStride</code> , <code>hStride</code> , <code>wStride</code> is negative or <code>dataType</code> has an invalid enumerant value.
CUDNN_STATUS_NOT_SUPPORTED	The total size of the tensor descriptor exceeds the maximim limit of 2 Giga-elements.

4.10. cudnnGetTensor4dDescriptor

```

cudnnStatus_t
cudnnGetTensor4dDescriptor( cudnnTensorDescriptor_t tensorDesc,
                           cudnnDataType_t *dataType,
                           int *n,
                           int *c,
                           int *h,
                           int *w,
                           int *nStride,
                           int *cStride,
                           int *hStride,
                           int *wStride )

```

This function queries the parameters of the previously initialized Tensor4D descriptor object.

Param	In/out	Meaning
tensorDesc	input	Handle to a previously initalized tensor descriptor.
datatype	output	Data type.
n	output	Number of images.
c	output	Number of feature maps per image.
h	output	Height of each feature map.
w	output	Width of each feature map.
nStride	output	Stride between two consecutive images.
cStride	output	Stride between two consecutive feature maps.
hStride	output	Stride between two consecutive rows.
wStride	output	Stride between two consecutive columns.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The operation succeeded.

4.11. cudnnSetTensorNdDescriptor

```

cudnnStatus_t
cudnnSetTensorNdDescriptor( cudnnTensorDescriptor_t  tensorDesc,

                             cudnnDataType_t dataType,
                             int nbDims,
                             int dimA[],
                             int strideA[])

```

This function initializes a previously created generic Tensor descriptor object.



The total size of a tensor including the potential padding between dimensions is limited to 2 Giga-elements of type `dataType`.

Param	In/out	Meaning
<code>tensorDesc</code>	input/ output	Handle to a previously created tensor descriptor.
<code>dataType</code>	input	Data type.
<code>nbDims</code>	input	Dimension of the tensor.
<code>dimA</code>	input	Array of dimension <code>nbDims</code> that contain the size of the tensor for every dimension.
<code>strideA</code>	input	Array of dimension <code>nbDims</code> that contain the stride of the tensor for every dimension.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the elements of the array <code>dimA</code> was negative or zero, or <code>dataType</code> has an invalid enumerant value.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	the parameter <code>nbDims</code> exceeds <code>CUDNN_DIM_MAX</code> or the total size of the tensor descriptor exceeds the maximim limit of 2 Giga-elements.

4.12. cudnnGetTensorNdDescriptor

```

cudnnStatus_t
cudnnGetTensorNdDescriptor( const cudnnTensorDescriptor_t  tensorDesc,

                           int  nbDimsRequested,
                           cudnnDataType_t *dataType,
                           int  *nbDims,
                           int  dimA[],
                           int  strideA[])

```

This function retrieves values stored in a previously initialized Tensor descriptor object.

Param	In/out	Meaning
tensorDesc	input	Handle to a previously initialized tensor descriptor.
nbDimsRequested	input	Number of dimensions to extract from a given tensor descriptor. It is also the minimum size of the arrays <code>dimA</code> and <code>strideA</code> . If this number is greater than the resulting <code>nbDims[0]</code> , only <code>nbDims[0]</code> dimensions will be returned.
datatype	output	Data type.
nbDims	output	Actual number of dimensions of the tensor will be returned in <code>nbDims[0]</code> .
dimA	output	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the dimensions from the provided tensor descriptor.
strideA	input	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the strides from the provided tensor descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The results were returned successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	Either <code>tensorDesc</code> or <code>nbDims</code> pointer is NULL.

4.13. cudnnDestroyTensorDescriptor

```

cudnnStatus_t cudnnDestroyTensorDescriptor(cudnnTensorDescriptor_t tensorDesc)

```

This function destroys a previously created Tensor descriptor object.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was destroyed successfully.

4.14. cudnnTransformTensor

```

cudnnStatus_t
cudnnTransformTensor( cudnnHandle_t      handle,
                     const void         *alpha,
                     const cudnnTensorDescriptor_t xDesc,
                     const void         *x,
                     const void         *beta,
                     const cudnnTensorDescriptor_t yDesc,
                     void                *y )

```

This function copies the scaled data from one tensor to another tensor with a different layout. Those descriptors need to have the same dimensions but not necessarily the same strides. The input and output tensors must not overlap in any way (i.e., tensors cannot be transformed in place). This function can be used to convert a tensor with an unsupported format to a supported one.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the source value with prior value in the destination tensor as follows: $\text{dstValue} = \text{alpha}[0] * \text{srcValue} + \text{beta}[0] * \text{priorDstValue}$. Please refer to this section for additional details.
xDesc	input	Handle to a previously initialized tensor descriptor.
x	input	Pointer to data of the tensor described by the <code>xDesc</code> descriptor.
yDesc	input	Handle to a previously initialized tensor descriptor.
y	output	Pointer to data of the tensor described by the <code>yDesc</code> descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	The dimensions <code>n, c, h, w</code> or the <code>dataType</code> of the two tensor descriptors are different.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.15. cudnnAddTensor_v2

```

cudnnStatus_t
cudnnAddTensor_v2( cudnnHandle_t      handle,
                  cudnnAddMode_t      mode,
                  const void         *alpha,
                  const cudnnTensorDescriptor_t bDesc,
                  const void         *b,
                  const void         *beta,
                  const cudnnTensorDescriptor_t yDesc,
                  void                *y )

```

This function adds the scaled values of one tensor to another tensor. The `mode` parameter can be used to select different ways of performing the scaled addition. The amount of data described by the `biasDesc` descriptor must match exactly the amount of data needed to perform the addition. Therefore, the following conditions must be met:

- ▶ Except for the `CUDNN_ADD_SAME_C` mode, the dimensions `h`, `w` of the two tensors must match.
- ▶ In the case of `CUDNN_ADD_IMAGE` mode, the dimensions `n`, `c` of the bias tensor must be 1.
- ▶ In the case of `CUDNN_ADD_FEATURE_MAP` mode, the dimension `n` of the bias tensor must be 1 and the dimension `c` of the two tensors must match.
- ▶ In the case of `CUDNN_ADD_FULL_TENSOR` mode, the dimensions `n`, `c` of the two tensors must match.
- ▶ In the case of `CUDNN_ADD_SAME_C` mode, the dimensions `n`, `w`, `h` of the bias tensor must be 1 and the dimension `c` of the two tensors must match.



Up to dimension 4, all tensor formats are supported. Beyond those dimensions, this routine is not supported.



This routine is deprecated. `cudaAddTensor` should be used instead.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
mode	input	Addition mode that describe how the addition is performed.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the source value with prior value in the destination tensor as follows: $dstValue = alpha[0]*srcValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
bDesc	input	Handle to a previously initialized tensor descriptor.
b	input	Pointer to data of the tensor described by the <code>bDesc</code> descriptor.
yDesc	input/ output	Handle to a previously initialized tensor descriptor.
y	input/ output	Pointer to data of the tensor described by the <code>yDesc</code> descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function executed successfully.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The dimensions of the bias tensor and the output tensor dimensions are above 4.
<code>CUDNN_STATUS_BAD_PARAM</code>	The dimensions <code>n</code> , <code>c</code> , <code>h</code> , <code>w</code> of the bias tensor refer to an amount of data that is incompatible with the

Return Value	Meaning
	<code>mode</code> parameter and the output tensor dimensions or the <code>dataType</code> of the two tensor descriptors are different.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.16. cudnnAddTensor

```

cudnnStatus_t
cudnnAddTensor_(
    cudnnHandle_t          handle,
    const void             *alpha,
    const cudnnTensorDescriptor_t bDesc,
    const void             *b,
    const void             *beta,
    const cudnnTensorDescriptor_t yDesc,
    void                  *y )

```

This function adds the scaled values of a bias tensor to another tensor. Each dimension of the bias tensor `b` must match the corresponding dimension of the destination tensor `y` or must be equal to 1. In the latter case, the same value from the bias tensor for those dimensions will be used to blend into the `y` tensor.



Up to dimension 5, all tensor formats are supported. Beyond those dimensions, this routine is not supported



This routine replaces deprecated routine `cudnnAddTensor_v2`

Param	In/out	Meaning
<code>handle</code>	input	Handle to a previously created cuDNN context.
<code>alpha, beta</code>	input	Pointers to scaling factors (in host memory) used to blend the source value with prior value in the destination tensor as follows: $dstValue = alpha[0]*srcValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
<code>bDesc</code>	input	Handle to a previously initialized tensor descriptor.
<code>b</code>	input	Pointer to data of the tensor described by the <code>bDesc</code> descriptor.
<code>yDesc</code>	input/ output	Handle to a previously initialized tensor descriptor.
<code>y</code>	input/ output	Pointer to data of the tensor described by the <code>yDesc</code> descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function executed successfully.

Return Value	Meaning
CUDNN_STATUS_NOT_SUPPORTED	The dimensions of the bias tensor and the output tensor dimensions are above 5.
CUDNN_STATUS_BAD_PARAM	The dimensions of the bias tensor refer to an amount of data that is incompatible the output tensor dimensions or the <code>dataType</code> of the two tensor descriptors are different.
CUDNN_STATUS_EXECUTION_FAILED	The function failed to launch on the GPU.

4.17. cudnnAddTensor_v3

```

cudnnStatus_t
cudnnAddTensor_v3(  cudnnHandle_t          handle,
                    const void            *alpha,
                    const cudnnTensorDescriptor_t bDesc,
                    const void            *b,
                    const void            *beta,
                    const cudnnTensorDescriptor_t yDesc,
                    void                  *y )

```

This function is now equivalent to `cudnnAddTensor`

4.18. cudnnSetTensor

```

cudnnStatus_t cudnnSetTensor(
    cudnnHandle_t          handle,
    const cudnnTensorDescriptor_t yDesc,
    void                  *y,
    const void            *valuePtr );

```

This function sets all the elements of a tensor to a given value.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
yDesc	input	Handle to a previously initialized tensor descriptor.
y	input/output	Pointer to data of the tensor described by the <code>yDesc</code> descriptor.
valuePtr	input	Pointer in Host memory to a single value. All elements of the <code>y</code> tensor will be set to <code>value[0]</code> . The data type of the element in <code>value[0]</code> has to match the data type of tensor <code>y</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The function launched successfully.
CUDNN_STATUS_BAD_PARAM	one of the provided pointers is nil
CUDNN_STATUS_EXECUTION_FAILED	The function failed to launch on the GPU.

4.19. cudnnScaleTensor

```

cudnnStatus_t cudnnScaleTensor( cudnnHandle_t      handle,
                                const cudnnTensorDescriptor_t yDesc,
                                void              *y,
                                const void       *alpha);

```

This function scale all the elements of a tensor by a given factor.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
yDesc	input	Handle to a previously initialized tensor descriptor.
y	input/ output	Pointer to data of the tensor described by the <code>yDesc</code> descriptor.
alpha	input	Pointer in Host memory to a single value that all elements of the tensor will be scaled with. Please refer to this section for additional details.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	one of the provided pointers is nil
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.20. cudnnCreateFilterDescriptor

```

cudnnStatus_t cudnnCreateFilterDescriptor( cudnnFilterDescriptor_t *filterDesc)

```

This function creates a filter descriptor object by allocating the memory needed to hold its opaque structure,

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was created successfully.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	The resources could not be allocated.

4.21. cudnnSetFilter4dDescriptor

```

cudnnStatus_t
cudnnSetFilter4dDescriptor( cudnnFilterDescriptor_t filterDesc,
                            cudnnDataType_t dataType,
                            int k,
                            int c,
                            int h,
                            int w )

```


This function initializes a previously created filter descriptor object into a 4D filter. Filters layout must be contiguous in memory. When using this routine to set up a filter descriptor, the filter format is set to CUDNN_TENSOR_NCHW.

Param	In/out	Meaning
filterDesc	input/output	Handle to a previously created filter descriptor.
datatype	input	Data type.
k	input	Number of output feature maps.
c	input	Number of input feature maps.
h	input	Height of each filter.
w	input	Width of each filter.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the parameters <code>k</code> , <code>c</code> , <code>h</code> , <code>w</code> is negative or <code>dataType</code> has an invalid enumerant value.

4.22. cudnnGetFilter4dDescriptor

```

cudnnStatus_t
cudnnGetFilter4dDescriptor( cudnnFilterDescriptor_t filterDesc,
                           cudnnDataType_t *dataType,
                           int *k,
                           int *c,
                           int *h,
                           int *w )

```

This function queries the parameters of the previously initialized filter descriptor object.

Param	In/out	Meaning
filterDesc	input	Handle to a previously created filter descriptor.
datatype	output	Data type.
k	output	Number of output feature maps.
c	output	Number of input feature maps.
h	output	Height of each filter.
w	output	Width of each filter.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.

4.23. cudnnSetFilter4dDescriptor_v4

```

cudnnStatus_t
cudnnSetFilter4dDescriptor_v4( cudnnFilterDescriptor_t filterDesc,
                               cudnnDataType_t dataType,
                               cudnnTensorFormat_t format,
                               int k,
                               int c,
                               int h,
                               int w )
    
```

This function initializes a previously created filter descriptor object into a 4D filter. Filters layout must be contiguous in memory.

Tensor format CUDNN_TENSOR_NHWC has limited support in **cudnnConvolutionForward**, **cudnnConvolutionBackwardData** and **cudnnConvolutionBackwardFilter**; please refer to each function's documentation for more information.

Param	In/out	Meaning
filterDesc	input/output	Handle to a previously created filter descriptor.
datatype	input	Data type.
format	input	Type of format.
k	input	Number of output feature maps.
c	input	Number of input feature maps.
h	input	Height of each filter.
w	input	Width of each filter.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the parameters k , c , h , w is negative or dataType or format has an invalid enumerant value.

4.24. cudnnGetFilter4dDescriptor_v4

```

cudnnStatus_t
cudnnGetFilter4dDescriptor_v4( cudnnFilterDescriptor_t filterDesc,
                               cudnnDataType_t *dataType,
                               cudnnTensorFormat_t *format,
                               int *k,
                               int *c,
                               int *h,
                               int *w )

```

This function queries the parameters of the previously initialized filter descriptor object.

Param	In/out	Meaning
filterDesc	input	Handle to a previously created filter descriptor.
datatype	output	Data type.
format	output	Type of format.
k	output	Number of output feature maps.
c	output	Number of input feature maps.
h	output	Height of each filter.
w	output	Width of each filter.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.

4.25. cudnnSetFilterNdDescriptor

```

cudnnStatus_t
cudnnSetFilterNdDescriptor( cudnnFilterDescriptor_t filterDesc,
                            cudnnDataType_t dataType,
                            int nbDims,
                            int filterDimA[])

```

This function initializes a previously created filter descriptor object. Filters layout must be contiguous in memory. When using this routine to set up a filter descriptor, the filter format is set to CUDNN_TENSOR_NCHW.

Param	In/out	Meaning
filterDesc	input/ output	Handle to a previously created filter descriptor.
datatype	input	Data type.
nbDims	input	Dimension of the filter.

Param	In/out	Meaning
filterDimA	input	Array of dimension <code>nbDims</code> containing the size of the filter for each dimension.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the elements of the array <code>filterDimA</code> is negative or <code>dataType</code> has an invalid enumerant value.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	the parameter <code>nbDims</code> exceeds <code>CUDNN_DIM_MAX</code> .

4.26. cudnnGetFilterNdDescriptor

```

cudnnStatus_t
cudnnGetFilterNdDescriptor( const cudnnFilterDescriptor_t wDesc,
                           int nbDimsRequested,
                           cudnnDataType_t *dataType,
                           int *nbDims,
                           int filterDimA[])

```

This function queries a previously initialized filter descriptor object.

Param	In/out	Meaning
wDesc	input	Handle to a previously initialized filter descriptor.
nbDimsRequested	input	Dimension of the expected filter descriptor. It is also the minimum size of the arrays <code>filterDimA</code> in order to be able to hold the results
dataType	input	Data type.
nbDims	input	Actual dimension of the filter.
filterDimA	input	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the filter parameters from the provided filter descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	The parameter <code>nbDimsRequested</code> is negative.

4.27. cudnnSetFilterNdDescriptor_v4

```

cudnnStatus_t
cudnnSetFilterNdDescriptor_v4( cudnnFilterDescriptor_t filterDesc,
                               cudnnDataType_t   dataType,
                               cudnnTensorFormat_t format,
                               int nbDims,
                               int filterDimA[])

```

This function initializes a previously created filter descriptor object. Filters layout must be contiguous in memory.

Tensor format CUDNN_TENSOR_NHWC has limited support in **cudnnConvolutionForward**, **cudnnConvolutionBackwardData** and **cudnnConvolutionBackwardFilter**; please refer to each function's documentation for more information.

Param	In/out	Meaning
filterDesc	input/ output	Handle to a previously created filter descriptor.
datatype	input	Data type.
format	input	Type of format.
nbDims	input	Dimension of the filter.
filterDimA	input	Array of dimension <code>nbDims</code> containing the size of the filter for each dimension.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the elements of the array <code>filterDimA</code> is negative or <code>dataType</code> or <code>format</code> has an invalid enumerant value.
CUDNN_STATUS_NOT_SUPPORTED	the parameter <code>nbDims</code> exceeds CUDNN_DIM_MAX.

4.28. cudnnGetFilterNdDescriptor_v4

```

cudnnStatus_t
cudnnGetFilterNdDescriptor_v4( const cudnnFilterDescriptor_t wDesc,
                               int nbDimsRequested,
                               cudnnDataType_t *dataType,
                               cudnnTensorFormat_t *format,
                               int *nbDims,
                               int filterDimA[])

```

This function queries a previously initialized filter descriptor object.

Param	In/out	Meaning
wDesc	input	Handle to a previously initialized filter descriptor.
nbDimsRequested	input	Dimension of the expected filter descriptor. It is also the minimum size of the arrays <code>filterDimA</code> in order to be able to hold the results
datatype	input	Data type.
format	output	Type of format.
nbDims	input	Actual dimension of the filter.
filterDimA	input	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the filter parameters from the provided filter descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	The parameter <code>nbDimsRequested</code> is negative.

4.29. cudnnDestroyFilterDescriptor

```
cudaStatus_t cudnnDestroyFilterDescriptor(cudaFilterDescriptor_t filterDesc)
```

This function destroys a previously created Tensor4D descriptor object.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was destroyed successfully.

4.30. cudnnCreateConvolutionDescriptor

```
cudaStatus_t cudnnCreateConvolutionDescriptor(cudaConvolutionDescriptor_t *convDesc)
```

This function creates a convolution descriptor object by allocating the memory needed to hold its opaque structure,

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was created successfully.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	The resources could not be allocated.

4.31. cudnnSetConvolution2dDescriptor

```

cudnnStatus_t
cudnnSetConvolution2dDescriptor( cudnnConvolutionDescriptor_t convDesc,
                                int pad_h,
                                int pad_w,
                                int u,
                                int v,
                                int upscalex,
                                int upscaley,
                                cudnnConvolutionMode_t mode )

```

This function initializes a previously created convolution descriptor object into a 2D correlation. This function assumes that the tensor and filter descriptors corresponds to the forward convolution path and checks if their settings are valid. That same convolution descriptor can be reused in the backward path provided it corresponds to the same layer.

Param	In/out	Meaning
convDesc	input/ output	Handle to a previously created convolution descriptor.
pad_h	input	zero-padding height: number of rows of zeros implicitly concatenated onto the top and onto the bottom of input images.
pad_w	input	zero-padding width: number of columns of zeros implicitly concatenated onto the left and onto the right of input images.
u	input	Vertical filter stride.
v	input	Horizontal filter stride.
upscalex	input	Upscale the input in x-direction.
upscaley	input	Upscale the input in y-direction.
mode	input	Selects between <code>CUDNN_CONVOLUTION</code> and <code>CUDNN_CROSS_CORRELATION</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The descriptor <code>convDesc</code> is nil. ▶ One of the parameters <code>pad_h</code>, <code>pad_w</code> is strictly negative. ▶ One of the parameters <code>u</code>, <code>v</code> is negative. ▶ The parameter <code>mode</code> has an invalid enumerant value.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The parameter <code>upscalex</code> or <code>upscaley</code> is not 1.

4.32. cudnnGetConvolution2dDescriptor

```

cudnnStatus_t
cudnnGetConvolution2dDescriptor( const cudnnConvolutionDescriptor_t convDesc,
                                int* pad_h,
                                int* pad_w,
                                int* u,
                                int* v,
                                int* upscalex,
                                int* upscaley,
                                cudnnConvolutionMode_t *mode )

```

This function queries a previously initialized 2D convolution descriptor object.

Param	In/out	Meaning
convDesc	input/ output	Handle to a previously created convolution descriptor.
pad_h	output	zero-padding height: number of rows of zeros implicitly concatenated onto the top and onto the bottom of input images.
pad_w	output	zero-padding width: number of columns of zeros implicitly concatenated onto the left and onto the right of input images.
u	output	Vertical filter stride.
v	output	Horizontal filter stride.
upscalex	output	Upscale the input in x-direction.
upscaley	output	Upscale the input in y-direction.
mode	output	convolution mode.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The operation was successful.
CUDNN_STATUS_BAD_PARAM	The parameter <code>convDesc</code> is nil.

4.33. cudnnGetConvolution2dForwardOutputDim

```

cudnnStatus_t
cudnnGetConvolution2dForwardOutputDim( const cudnnConvolutionDescriptor_t
convDesc,
                                        const cudnnTensorDescriptor_t
inputTensorDesc,
                                        const cudnnFilterDescriptor_t filterDesc,
                                        int *n,
                                        int *c,
                                        int *h,
                                        int *w )

```

This function returns the dimensions of the resulting 4D tensor of a 2D convolution, given the convolution descriptor, the input tensor descriptor and the filter descriptor

This function can help to setup the output tensor and allocate the proper amount of memory prior to launch the actual convolution.

Each dimension **h** and **w** of the output images is computed as followed:

```
outputDim = 1 + (inputDim + 2*pad - filterDim)/convolutionStride;
```

Param	In/out	Meaning
convDesc	input	Handle to a previously created convolution descriptor.
inputTensorDesc	input	Handle to a previously initialized tensor descriptor.
filterDesc	input	Handle to a previously initialized filter descriptor.
n	output	Number of output images.
c	output	Number of output feature maps per image.
h	output	Height of each output feature map.
w	output	Width of each output feature map.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_BAD_PARAM	One or more of the descriptors has not been created correctly or there is a mismatch between the feature maps of inputTensorDesc and filterDesc .
CUDNN_STATUS_SUCCESS	The object was set successfully.

4.34. cudnnSetConvolutionNdDescriptor_v2

```

cudnnStatus_t
cudnnSetConvolutionNdDescriptor_v2( cudnnConvolutionDescriptor_t convDesc,
                                     int arrayLength,
                                     int padA[],
                                     int filterStrideA[],
                                     int upscaleA[],
                                     cudnnConvolutionMode_t mode )

```

This function initializes a previously created generic convolution descriptor object into a n-D correlation. That same convolution descriptor can be reused in the backward path provided it corresponds to the same layer. The convolution computation will done in the same precision than the input and output tensors datatypes.



This routine is deprecated. `cudnnSetConvolutionNdDescriptor` should be used instead.

Param	In/out	Meaning
convDesc	input/ output	Handle to a previously created convolution descriptor.
arrayLength	input	Dimension of the convolution.
padA	input	Array of dimension <code>arrayLength</code> containing the zero-padding size for each dimension. For every dimension, the padding represents the number of extra zeros implicitly concatenated at the start and at the end of every element of that dimension .
filterStrideA	input	Array of dimension <code>arrayLength</code> containing the filter stride for each dimension. For every dimension, the filter stride represents the number of elements to slide to reach the next start of the filtering window of the next point.
upscaleA	input	Array of dimension <code>arrayLength</code> containing the upscale factor for each dimension.
mode	input	Selects between <code>CUDNN_CONVOLUTION</code> and <code>CUDNN_CROSS_CORRELATION</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The descriptor <code>convDesc</code> is nil. ▶ The <code>arrayLengthRequest</code> is negative. ▶ The enumerant <code>mode</code> has an invalid value. ▶ One of the elements of <code>padA</code> is strictly negative. ▶ One of the elements of <code>strideA</code> is negative or zero.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The <code>arrayLengthRequest</code> is greater than <code>CUDNN_DIM_MAX</code>. ▶ The array <code>upscaleA</code> contains an element different from 1.

4.35. cudnnGetConvolutionNdDescriptor_v2

```

cudnnStatus_t
cudnnGetConvolutionNdDescriptor_v2( const cudnnConvolutionDescriptor_t convDesc,
                                   int arrayLengthRequested,
                                   int *arrayLength,
                                   int padA[],

                                   int filterStrideA[],
                                   int upscaleA[],
                                   cudnnConvolutionMode_t *mode )

```

This function queries a previously initialized convolution descriptor object.



This routine is deprecated. `cudaGetConvolutionNdDescriptor` should be used instead.

Param	In/out	Meaning
<code>convDesc</code>	input/ output	Handle to a previously created convolution descriptor.
<code>arrayLengthRequested</code>	input	Dimension of the expected convolution descriptor. It is also the minimum size of the arrays <code>padA</code> , <code>filterStrideA</code> and <code>upscaleA</code> in order to be able to hold the results
<code>arrayLength</code>	output	actual dimension of the convolution descriptor.
<code>padA</code>	output	Array of dimension of at least <code>arrayLengthRequested</code> that will be filled with the padding parameters from the provided convolution descriptor.
<code>filterStrideA</code>	output	Array of dimension of at least <code>arrayLengthRequested</code> that will be filled with the filter stride from the provided convolution descriptor.
<code>upscaleA</code>	output	Array of dimension at least <code>arrayLengthRequested</code> that will be filled with the upscaling parameters from the provided convolution descriptor.
<code>mode</code>	output	convolution mode of the provided descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The descriptor <code>convDesc</code> is nil. ▶ The <code>arrayLengthRequest</code> is negative.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The <code>arrayLengthRequest</code> is greater than <code>CUDNN_DIM_MAX</code> .

4.36. `cudaSetConvolutionNdDescriptor`


```

cudaStatus_t
cudaSetConvolutionNdDescriptor( cudaConvolutionDescriptor_t convDesc,
                                int arrayLength,
                                int padA[],
                                int filterStrideA[],
                                int upscaleA[],
                                cudaConvolutionMode_t mode,
                                cudaDataType_t dataType )

```

This function initializes a previously created generic convolution descriptor object into a n-D correlation. That same convolution descriptor can be reused in the backward path

provided it corresponds to the same layer. The convolution computation will be done in the specified **dataType**, which can be potentially different from the input/output tensors.

 This routine replaces deprecated routine `cudaSetConvolutionNdDescriptor_v2`

Param	In/out	Meaning
<code>convDesc</code>	input/ output	Handle to a previously created convolution descriptor.
<code>arrayLength</code>	input	Dimension of the convolution.
<code>padA</code>	input	Array of dimension <code>arrayLength</code> containing the zero-padding size for each dimension. For every dimension, the padding represents the number of extra zeros implicitly concatenated at the start and at the end of every element of that dimension .
<code>filterStrideA</code>	input	Array of dimension <code>arrayLength</code> containing the filter stride for each dimension. For every dimension, the filter stride represents the number of elements to slide to reach the next start of the filtering window of the next point.
<code>upscaleA</code>	input	Array of dimension <code>arrayLength</code> containing the upscale factor for each dimension.
<code>mode</code>	input	Selects between <code>CUDNN_CONVOLUTION</code> and <code>CUDNN_CROSS_CORRELATION</code> .
<code>dataType</code>	input	Selects the datatype in which the computation will be done.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The descriptor <code>convDesc</code> is nil. ▶ The <code>arrayLengthRequest</code> is negative. ▶ The enumerant <code>mode</code> has an invalid value. ▶ The enumerant <code>dataType</code> has an invalid value. ▶ One of the elements of <code>padA</code> is strictly negative. ▶ One of the elements of <code>strideA</code> is negative or zero.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The <code>arrayLengthRequest</code> is greater than <code>CUDNN_DIM_MAX</code>. ▶ The array <code>upscaleA</code> contains an element different from 1.

4.37. cudnnGetConvolutionNdDescriptor

```

cudnnStatus_t
cudnnGetConvolutionNdDescriptor( const cudnnConvolutionDescriptor_t convDesc,
                                int arrayLengthRequested,
                                int *arrayLength,
                                int padA[],

                                int filterStrideA[],
                                int upscaleA[],
                                cudnnConvolutionMode_t *mode,
                                cudnnDataType_t *dataType )

```

This function queries a previously initialized convolution descriptor object.



This routine replaces deprecated routine `cudnnGetConvolutionNdDescriptor_v2`

Param	In/out	Meaning
convDesc	input/ output	Handle to a previously created convolution descriptor.
arrayLengthRequested	input	Dimension of the expected convolution descriptor. It is also the minimum size of the arrays <code>padA</code> , <code>filterStrideA</code> and <code>upscaleA</code> in order to be able to hold the results
arrayLength	output	actual dimension of the convolution descriptor.
padA	output	Array of dimension of at least <code>arrayLengthRequested</code> that will be filled with the padding parameters from the provided convolution descriptor.
filterStrideA	output	Array of dimension of at least <code>arrayLengthRequested</code> that will be filled with the filter stride from the provided convolution descriptor.
upscaleA	output	Array of dimension at least <code>arrayLengthRequested</code> that will be filled with the upscaling parameters from the provided convolution descriptor.
mode	output	convolution mode of the provided descriptor.
datatype	output	datatype of the provided descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The descriptor <code>convDesc</code> is nil. ▶ The <code>arrayLengthRequest</code> is negative.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The <code>arrayLengthRequest</code> is greater than <code>CUDNN_DIM_MAX</code>

4.38. cudnnSetConvolutionNdDescriptor_v3

```

cudnnStatus_t
cudnnSetConvolutionNdDescriptor_v3( cudnnConvolutionDescriptor_t convDesc,
                                     int arrayLength,
                                     int padA[],
                                     int filterStrideA[],
                                     int upscaleA[],
                                     cudnnConvolutionMode_t mode,
                                     cudnnDataType_t dataType )

```

This function is now equivalent to `cudnnSetConvolutionNdDescriptor`.

4.39. cudnnGetConvolutionNdDescriptor_v3

```

cudnnStatus_t
cudnnGetConvolutionNdDescriptor_v3( const cudnnConvolutionDescriptor_t convDesc,
                                     int arrayLengthRequested,
                                     int *arrayLength,
                                     int padA[],

                                     int filterStrideA[],
                                     int upscaleA[],
                                     cudnnConvolutionMode_t *mode,
                                     cudnnDataType_t *dataType )

```

This function is now equivalent to `cudnnGetConvolutionNdDescriptor`

4.40. cudnnGetConvolutionNdForwardOutputDim

```

cudnnStatus_t
cudnnGetConvolutionNdForwardOutputDim( const cudnnConvolutionDescriptor_t
    convDesc,
                                     const cudnnTensorDescriptor_t
    inputTensorDesc,
                                     const cudnnFilterDescriptor_t filterDesc,
                                     int nbDims,
                                     int tensorOutputDimA[] )

```

This function returns the dimensions of the resulting n-D tensor of a **nbDims-2-D** convolution, given the convolution descriptor, the input tensor descriptor and the filter descriptor. This function can help to setup the output tensor and allocate the proper amount of memory prior to launch the actual convolution.

Each dimension of the **(nbDims-2)-D** images of the output tensor is computed as followed:

```
outputDim = 1 + (inputDim + 2*pad - filterDim)/convolutionStride;
```

Param	In/out	Meaning
convDesc	input	Handle to a previously created convolution descriptor.
inputTensorDesc	input	Handle to a previously initialized tensor descriptor.

Param	In/out	Meaning
filterDesc	input	Handle to a previously initialized filter descriptor.
nbDims	input	Dimension of the output tensor
tensorOutputDims	output	Array of dimensions <code>nbDims</code> that contains on exit of this routine the sizes of the output tensor

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the parameters <code>convDesc</code>, <code>inputTensorDesc</code>, and <code>filterDesc</code>, is nil ▶ The dimension of the filter descriptor <code>filterDesc</code> is different from the dimension of input tensor descriptor <code>inputTensorDesc</code>. ▶ The dimension of the convolution descriptor is different from the dimension of input tensor descriptor <code>inputTensorDesc - 2</code>. ▶ The features map of the filter descriptor <code>filterDesc</code> is different from the one of input tensor descriptor <code>inputTensorDesc</code>. ▶ The size of the filter <code>filterDesc</code> is larger than the padded sizes of the input tensor. ▶ The dimension <code>nbDims</code> of the output array is negative or greater than the dimension of input tensor descriptor <code>inputTensorDesc</code>.
<code>CUDNN_STATUS_SUCCESS</code>	The routine exits successfully.

4.41. cudnnDestroyConvolutionDescriptor

```

cudnnStatus_t cudnnDestroyConvolutionDescriptor(cudnnConvolutionDescriptor_t
convDesc)

```

This function destroys a previously created convolution descriptor object.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was destroyed successfully.

4.42. cudnnFindConvolutionForwardAlgorithm

```

cudnnStatus_t
cudnnFindConvolutionForwardAlgorithm( cudnnHandle_t          handle,
                                     const cudnnTensorDescriptor_t xDesc,
                                     const cudnnFilterDescriptor_t wDesc,
                                     const cudnnConvolutionDescriptor_t convDesc,
                                     const cudnnTensorDescriptor_t yDesc,
                                     const int requestedAlgoCount,
                                     int *returnedAlgoCount,
                                     cudnnConvolutionFwdAlgoPerf_t *perfResults
                                     )

```

This function attempts all cuDNN algorithms and outputs performance metrics to a user-allocated array of `cudnnConvolutionFwdAlgoPerf_t`. These metrics are written in sorted fashion where the first element has the lowest compute time.



This function is host blocking.



It is recommend to run this function prior to allocating layer data; doing otherwise may needlessly inhibit some algorithm options due to resource usage.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
xDesc	input	Handle to the previously initialized input tensor descriptor.
wDesc	input	Handle to a previously initialized filter descriptor.
convDesc	input	Previously initialized convolution descriptor.
yDesc	input	Handle to the previously initialized output tensor descriptor.
requestedAlgoCount	input	The maximum number of elements to be stored in perfResults.
returnedAlgoCount	output	The number of output elements stored in perfResults.
perfResults	output	A user-allocated array to store performance metrics sorted ascending by compute time.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successful.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met:

Return Value	Meaning
	<ul style="list-style-type: none"> ▶ The handle is not allocated properly. ▶ The x, w or y descriptor is not allocated properly. ▶ The x, w or y descriptor has fewer than 1 dimension. ▶ Either returnedCount or perfResults is pointing to NULL. ▶ The requestedCount is less than 1.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	This function was unable to allocate memory to store sample input, filters and output.
<code>CUDNN_STATUS_INTERNAL_ERROR</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The function was unable to allocate necessary timing objects. ▶ The function was unable to deallocate necessary timing objects. ▶ The function was unable to deallocate sample input, filters and output.

4.43. cudnnGetConvolutionForwardAlgorithm

```

cudnnStatus_t
cudnnGetConvolutionForwardAlgorithm( cudnnHandle_t      handle,
                                     const cudnnTensorDescriptor_t xDesc,
                                     const cudnnFilterDescriptor_t  wDesc,
                                     const cudnnConvolutionDescriptor_t
convDesc,
                                     const cudnnTensorDescriptor_t  yDesc,
                                     cudnnConvolutionFwdPreference_t
preference,
                                     size_t
memoryLimitInbytes,
                                     cudnnConvolutionFwdAlgo_t      *algo
)
    
```

This function serves as a heuristic for obtaining the best suited algorithm for **cudnnConvolutionForward** for the given layer specifications. Based on the input preference, this function will either return the fastest algorithm or the fastest algorithm within a given memory limit. For an exhaustive search for the fastest algorithm, please use **cudnnFindConvolutionForwardAlgorithm**.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
xDesc	input	Handle to the previously initialized input tensor descriptor.
wDesc	input	Handle to a previously initialized convolution filter descriptor.
convDesc	input	Previously initialized convolution descriptor.
yDesc	input	Handle to the previously initialized output tensor descriptor.

Param	In/out	Meaning
preference	input	Enumerant to express the preference criteria in terms of memory requirement and speed.
memoryLimitInBytes	input	It is used when enumerant <code>preference</code> is set to <code>CUDNN_CONVOLUTION_FWD_SPECIFY_WORKSPACE_LIMIT</code> to specify the maximum amount of GPU memory the user is willing to use as a workspace
algo	output	Enumerant that specifies which convolution algorithm should be used to compute the results according to the specified preference

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successful.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the parameters <code>handle</code>, <code>xDesc</code>, <code>wDesc</code>, <code>convDesc</code>, <code>yDesc</code> is NULL. ▶ Either <code>yDesc</code> or <code>wDesc</code> have different dimensions from <code>xDesc</code>. ▶ The data types of tensors <code>xDesc</code>, <code>yDesc</code> or <code>wDesc</code> are not all the same. ▶ The number of feature maps in <code>xDesc</code> and <code>wDesc</code> differs. ▶ The tensor <code>xDesc</code> has a dimension smaller than 3.

4.44. cudnnGetConvolutionForwardWorkspaceSize

```

cudnnStatus_t
cudnnGetConvolutionForwardWorkspaceSize( cudnnHandle_t  handle,
                                         const  cudnnTensorDescriptor_t
                                         xDesc,
                                         const  cudnnFilterDescriptor_t
                                         wDesc,
                                         const  cudnnConvolutionDescriptor_t
                                         convDesc,
                                         const  cudnnTensor4dDescriptor_t
                                         yDesc,
                                         cudnnConvolutionFwdAlgo_t
                                         algo,
                                         size_t
                                         *sizeInBytes
                                         )

```

This function returns the amount of GPU memory workspace the user needs to allocate to be able to call `cudnnConvolutionForward` with the specified algorithm. The workspace allocated will then be passed to the routine `cudnnConvolutionForward`. The specified algorithm can be the result of the call to `cudnnGetConvolutionForwardAlgorithm` or can be chosen arbitrarily by the user. Note that not every algorithm is available for every configuration of the input tensor and/or every configuration of the convolution descriptor.

Param	In/ out	Meaning
handle	input	Handle to a previously created cuDNN context.
xDesc	input	Handle to the previously initialized x tensor descriptor.
wDesc	input	Handle to a previously initialized filter descriptor.
convDesc	input	Previously initialized convolution descriptor.
yDesc	input	Handle to the previously initialized y tensor descriptor.
algo	input	Enumerant that specifies the chosen convolution algorithm
sizeInBytes	output	Amount of GPU memory needed as workspace to be able to execute a forward convolution with the specified <code>algo</code>

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successful.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the parameters <code>handle</code>, <code>xDesc</code>, <code>wDesc</code>, <code>convDesc</code>, <code>yDesc</code> is NULL. ▶ The tensor <code>yDesc</code> or <code>wDesc</code> are not of the same dimension as <code>xDesc</code>. ▶ The tensor <code>xDesc</code>, <code>yDesc</code> or <code>wDesc</code> are not of the same data type. ▶ The numbers of feature maps of the tensor <code>xDesc</code> and <code>wDesc</code> differ. ▶ The tensor <code>xDesc</code> has a dimension smaller than 3.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The combination of the tensor descriptors, filter descriptor and convolution descriptor is not supported for the specified algorithm.

4.45. cudnnConvolutionForward

```

cudnnStatus_t
cudnnConvolutionForward( cudnnHandle_t          handle,
                        const void             *alpha,
                        const cudnnTensorDescriptor_t xDesc,
                        const void             *x,
                        const cudnnFilterDescriptor_t wDesc,
                        const void             *w,
                        const cudnnConvolutionDescriptor_t convDesc,
                        cudnnConvolutionFwdAlgo_t algo,
                        void                    *workspace,
                        size_t                  workspaceSizeInBytes,
                        const void             *beta,
                        const cudnnTensorDescriptor_t yDesc,
                        void                    *y )

```

This function executes convolutions or cross-correlations over **x** using filters specified with **w**, returning results in **y**. Scaling factors **alpha** and **beta** can be used to scale the input tensor and the output tensor respectively.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc	input	Handle to a previously initialized tensor descriptor.
x	input	Data pointer to GPU memory associated with the tensor descriptor xDesc .
wDesc	input	Handle to a previously initialized filter descriptor.
w	input	Data pointer to GPU memory associated with the filter descriptor wDesc .
convDesc	input	Previously initialized convolution descriptor.
algo	input	Enumerant that specifies which convolution algorithm should be used to compute the results
workSpace	input	Data pointer to GPU memory to a workspace needed to able to execute the specified algorithm. If no workspace is needed for a particular algorithm, that pointer can be nil
workSpaceSizeInBytes	input	Specifies the size in bytes of the provided workSpace
yDesc	input	Handle to a previously initialized tensor descriptor.
y	input/ output	Data pointer to GPU memory associated with the tensor descriptor yDesc that carries the result of the convolution.

This function supports only four specific combinations of data types for **xDesc**, **wDesc**, **convDesc** and **yDesc**. See the following for an exhaustive list of these configurations.

Data Type Configurations	xDesc 's, wDesc 's and yDesc 's Data Type	convDesc 's Data Type
TRUE_HALF_CONFIG	CUDNN_DATA_HALF	CUDNN_DATA_HALF
PSEUDO_HALF_CONFIG	CUDNN_DATA_HALF	CUDNN_DATA_FLOAT
FLOAT_CONFIG	CUDNN_DATA_FLOAT	CUDNN_DATA_FLOAT
DOUBLE_CONFIG	CUDNN_DATA_DOUBLE	CUDNN_DATA_DOUBLE



TRUE_HALF_CONFIG is only supported on architectures with true fp16 support (compute capability 5.3 and 6.0).

For this function, all algorithms perform deterministic computations. Specifying a separate algorithm can cause changes in performance and support. See the following for an exhaustive list of algorithm options and their respective supported parameters.

wDesc may only have format CUDNN_TENSOR_NHWC when all of the following are true:

- ▶ **algo** is CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_GEMM
- ▶ **xDesc** and **yDesc** is NHWC HWC-packed
- ▶ Data type configuration is PSEUDO_HALF_CONFIG or FLOAT_CONFIG
- ▶ The convolution is 2-dimensional

The following is an exhaustive list of algo support for 2-d convolutions.

- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_GEMM**
 - ▶ **xDesc** Format Support: All
 - ▶ **yDesc** Format Support: All
 - ▶ Data Type Config Support: All except TRUE_HALF_CONFIG
- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_PRECOMPUTED_GEMM**
 - ▶ **xDesc** Format Support: All
 - ▶ **yDesc** Format Support: All
 - ▶ Data Type Config Support: All
- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_GEMM**
 - ▶ **xDesc** Format Support: All
 - ▶ **yDesc** Format Support: All
 - ▶ Data Type Config Support: All except TRUE_HALF_CONFIG
- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_DIRECT**
 - ▶ This algorithm has no current implementation in cuDNN.
- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_FFT**
 - ▶ **xDesc** Format Support: NCHW HW-packed
 - ▶ **yDesc** Format Support: NCHW HW-packed
 - ▶ Data Type Config Support: PSEUDO_HALF_CONFIG, FLOAT_CONFIG
 - ▶ Notes:
 - ▶ **xDesc**'s feature map height + 2 * **convDesc**'s zero-padding height must equal 256 or less
 - ▶ **xDesc**'s feature map width + 2 * **convDesc**'s zero-padding width must equal 256 or less
 - ▶ **convDesc**'s vertical and horizontal filter stride must equal 1
 - ▶ **wDesc**'s filter height must be greater than **convDesc**'s zero-padding height
 - ▶ **wDesc**'s filter width must be greater than **convDesc**'s zero-padding width
- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_FFT_TILING**
 - ▶ **xDesc** Format Support: NCHW HW-packed
 - ▶ **yDesc** Format Support: NCHW HW-packed
 - ▶ Data Type Config Support: PSEUDO_HALF_CONFIG, FLOAT_CONFIG
 - ▶ Notes:
 - ▶ **wDesc**'s filter height must equal 32 or less
 - ▶ **wDesc**'s filter width must equal 32 or less

- ▶ **convDesc**'s vertical and horizontal filter stride must equal 1
- ▶ **wDesc**'s filter height must be greater than **convDesc**'s zero-padding height
- ▶ **wDesc**'s filter width must be greater than **convDesc**'s zero-padding width

The following is an exhaustive list of algo support for 3-d convolutions.

- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_GEMM**
 - ▶ **xDesc** Format Support: All
 - ▶ **yDesc** Format Support: All
 - ▶ Data Type Config Support: All except TRUE_HALF_CONFIG
- ▶ **CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_PRECOMPUTED_GEMM**
 - ▶ **xDesc** Format Support: NCDHW-fully-packed
 - ▶ **yDesc** Format Support: NCDHW-fully-packed
 - ▶ Data Type Config Support: All except TRUE_HALF_CONFIG

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The operation was launched successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ At least one of the following is NULL: handle, xDesc, wDesc, convDesc, yDesc, xData, w, yData, alpha, beta ▶ xDesc and yDesc have a non-matching number of dimensions ▶ xDesc and wDesc have a non-matching number of dimensions ▶ xDesc has fewer than three number of dimensions ▶ xDesc's number of dimensions is not equal to convDesc's array length + 2 ▶ xDesc and wDesc have a non-matching number of input feature maps per image ▶ xDesc, wDesc and yDesc have a non-matching data type ▶ For some spatial dimension, wDesc has a spatial size that is larger than the input spatial size (including zero-padding size)
CUDNN_STATUS_NOT_SUPPORTED	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ xDesc or yDesc have negative tensor striding ▶ xDesc, wDesc or yDesc has a number of dimensions that is not 4 or 5 ▶ The chosen algo does not support the parameters provided; see above for exhaustive list of parameter support for each algo
CUDNN_STATUS_MAPPING_ERROR	An error occurred during the texture binding of the filter data.
CUDNN_STATUS_EXECUTION_FAILED	The function failed to launch on the GPU.

4.46. cudnnConvolutionBackwardBias

```

cudnnStatus_t
cudnnConvolutionBackwardBias( cudnnHandle_t      handle,
                              const void        *alpha,
                              const cudnnTensorDescriptor_t dyDesc,
                              const void        *dy,
                              const void        *beta,
                              const cudnnTensorDescriptor_t dbDesc,
                              void              *db
                              )

```

This function computes the convolution function gradient with respect to the bias, which is the sum of every element belonging to the same feature map across all of the images of the input tensor. Therefore, the number of elements produced is equal to the number of features maps of the input tensor.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
dyDesc	input	Handle to the previously initialized input tensor descriptor.
dy	input	Data pointer to GPU memory associated with the tensor descriptor dyDesc .
dbDesc	input	Handle to the previously initialized output tensor descriptor.
db	output	Data pointer to GPU memory associated with the output tensor descriptor dbDesc .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The operation was launched successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the parameters n,height,width of the output tensor is not 1. ▶ The numbers of feature maps of the input tensor and output tensor differ. ▶ The dataType of the two tensor descriptors are different.

4.47. cudnnConvolutionBackwardFilter_v2

```

cudnnStatus_t
cudnnConvolutionBackwardFilter_v2( cudnnHandle_t          handle,
                                   const void             *alpha,
                                   const cudnnTensorDescriptor_t xDesc,
                                   const void             *xData,
                                   const cudnnTensorDescriptor_t dyDesc,
                                   const void             *dyData,
                                   const cudnnConvolutionDescriptor_t convDesc,
                                   const void             *beta,

                                   const cudnnFilterDescriptor_t dwDesc,
                                   void                   *dw );

```

This function is equivalent to running **cudnnConvolutionBackwardFilter** with the following additional parameters:

- ▶ algo: CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0
- ▶ workSpace: NULL
- ▶ workSpaceSizeInBytes: 0



This function is deprecated. **cudnnConvolutionBackwardFilter** should be used instead. Refer to **cudnnConvolutionBackwardFilter** for documentation.

4.48. cudnnFindConvolutionBackwardFilterAlgorithm

```

cudnnStatus_t
cudnnFindConvolutionBackwardFilterAlgorithm( cudnnHandle_t
  handle,

  cudnnTensorDescriptor_t xDesc,
  cudnnTensorDescriptor_t dyDesc,
  cudnnConvolutionDescriptor_t convDesc,
  cudnnFilterDescriptor_t dwDesc,

  const int requestedAlgoCount,
  int *returnedAlgoCount,
  cudnnConvolutionBwdFilterAlgoPerf_t *perfResults
)

```

This function attempts all cuDNN algorithms for **cudnnConvolutionBackwardFilter** and outputs performance metrics to a user-allocated array of

`cudaConvolutionBwdFilterAlgoPerf_t`. These metrics are written in sorted fashion where the first element has the lowest compute time.



This function is host blocking.



It is recommend to run this function prior to allocating layer data; doing otherwise may needlessly inhibit some algorithm options due to resource usage.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
xDesc	input	Handle to the previously initialized input tensor descriptor.
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
convDesc	input	Previously initialized convolution descriptor.
dwDesc	input	Handle to a previously initialized filter descriptor.
requestedAlgoCount	input	The maximum number of elements to be stored in perfResults.
returnedAlgoCount	output	The number of output elements stored in perfResults.
perfResults	output	A user-allocated array to store performance metrics sorted ascending by compute time.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successful.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The cudnn handle is not allocated properly. ▶ Any of the x, dy or dw descriptors were not allocated properly. ▶ Any of the xDesc, dyDesc or dwDesc descriptors have fewer than 1 dimension. ▶ Either returnedCount or perfResults is pointing to NULL. ▶ The requestedCount is less than 1.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	This function was unable to allocate memory to store sample input, filters and output.
<code>CUDNN_STATUS_INTERNAL_ERROR</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The function was unable to allocate necessary timing objects. ▶ The function was unable to deallocate necessary timing objects. ▶ The function was unable to deallocate sample input, filters and output.

4.49. cudnnGetConvolutionBackwardFilterAlgorithm

```

cudnnStatus_t
cudnnGetConvolutionBackwardFilterAlgorithm( cudnnHandle_t
    handle,
                                           const cudnnTensorDescriptor_t
    xDesc,
                                           const cudnnTensorDescriptor_t
    dyDesc,
                                           const cudnnConvolutionDescriptor_t
    convDesc,
                                           const cudnnFilterDescriptor_t
    dwDesc,

    cudnnConvolutionBwdFilterPreference_t preference,
    size_t
    memoryLimitInbytes,
                                           cudnnConvolutionBwdFilterAlgo_t
    *algo
    )
    
```

This function serves as a heuristic for obtaining the best suited algorithm for **cudnnConvolutionBackwardFilter** for the given layer specifications. Based on the input preference, this function will either return the fastest algorithm or the fastest algorithm within a given memory limit. For an exhaustive search for the fastest algorithm, please use **cudnnFindConvolutionBackwardFilterAlgorithm**.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
xDesc	input	Handle to the previously initialized input tensor descriptor.
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
convDesc	input	Previously initialized convolution descriptor.
dwDesc	input	Handle to a previously initialized filter descriptor.
preference	input	Enumerant to express the preference criteria in terms of memory requirement and speed.
memoryLimitInbytes	input	It is to specify the maximum amount of GPU memory the user is willing to use as a workspace. This is currently a placeholder and is not used.
algo	output	Enumerant that specifies which convolution algorithm should be used to compute the results according to the specified preference

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The query was successful.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The numbers of feature maps of the input tensor and output tensor differ. ▶ The dataType of the two tensor descriptors or the filter are different.

4.50. cudnnGetConvolutionBackwardFilterWorkspaceSize

```

cudnnStatus_t
cudnnGetConvolutionBackwardFilterWorkspaceSize( cudnnHandle_t  handle,
                                                const cudnnTensorDescriptor_t
                                                xDesc,
                                                const cudnnTensorDescriptor_t
                                                dyDesc,
                                                const
                                                cudnnConvolutionDescriptor_t  convDesc,
                                                const cudnnFilterDescriptor_t
                                                dwDesc,
                                                cudnnConvolutionFwdAlgo_t
                                                algo,
                                                size_t
                                                *sizeInBytes
                                                )

```

This function returns the amount of GPU memory workspace the user needs to allocate to be able to call **cudnnConvolutionBackwardFilter** with the specified algorithm. The workspace allocated will then be passed to the routine **cudnnConvolutionBackwardFilter**. The specified algorithm can be the result of the call to **cudnnGetConvolutionBackwardFilterAlgorithm** or can be chosen arbitrarily by the user. Note that not every algorithm is available for every configuration of the input tensor and/or every configuration of the convolution descriptor.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
xDesc	input	Handle to the previously initialized input tensor descriptor.
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
convDesc	input	Previously initialized convolution descriptor.
dwDesc	input	Handle to a previously initialized filter descriptor.
algo	input	Enumerant that specifies the chosen convolution algorithm
sizeInBytes	output	Amount of GPU memory needed as workspace to be able to execute a forward convolution with the specified <code>algo</code>

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The query was successful.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The numbers of feature maps of the input tensor and output tensor differ. ▶ The dataType of the two tensor descriptors or the filter are different.

Return Value	Meaning
CUDNN_STATUS_NOT_SUPPORTED	The combination of the tensor descriptors, filter descriptor and convolution descriptor is not supported for the specified algorithm.


4.51. cudnnConvolutionBackwardFilter

```

cudnnStatus_t
cudnnConvolutionBackwardFilter ( cudnnHandle_t          handle,
                                const void             *alpha,
                                const cudnnTensorDescriptor_t xDesc,
                                const void             *x,
                                const cudnnTensorDescriptor_t dyDesc,
                                const void             *dy,
                                const cudnnConvolutionDescriptor_t
convDesc,
                                cudnnConvolutionBwdFilterAlgo_t algo,
                                void
*workspace,
                                size_t
workspaceSizeInBytes,
                                const void             *beta,
                                const cudnnFilterDescriptor_t dwDesc,
                                void             *dw )

```

This function computes the convolution gradient with respect to filter coefficients using the specified **algo**, returning results in **gradDesc**. Scaling factors **alpha** and **beta** can be used to scale the input tensor and the output tensor respectively.

 This routine replaces `cudnnConvolutionBackwardFilter_v2`

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc	input	Handle to a previously initialized tensor descriptor.
x	input	Data pointer to GPU memory associated with the tensor descriptor <code>xDesc</code> .
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
dy	input	Data pointer to GPU memory associated with the backpropagation gradient tensor descriptor <code>dyDesc</code> .
convDesc	input	Previously initialized convolution descriptor.
algo	input	Enumerant that specifies which convolution algorithm should be used to compute the results

Param	In/out	Meaning
workSpace	input	Data pointer to GPU memory to a workspace needed to able to execute the specified algorithm. If no workspace is needed for a particular algorithm, that pointer can be nil
workSpaceSizeInBytes	input	Specifies the size in bytes of the provided <code>workSpace</code>
dwDesc	input	Handle to a previously initialized filter gradient descriptor.
dw	input/output	Data pointer to GPU memory associated with the filter gradient descriptor <code>dwDesc</code> that carries the result.

This function supports only three specific combinations of data types for `xDesc`, `dyDesc`, `convDesc` and `dwDesc`. See the following for an exhaustive list of these configurations.

Data Type Configurations	<code>xDesc</code> 's, <code>dyDesc</code> 's and <code>dwDesc</code> 's Data Type	<code>convDesc</code> 's Data Type
HALF_CONFIG	CUDNN_DATA_HALF	CUDNN_DATA_FLOAT
FLOAT_CONFIG	CUDNN_DATA_FLOAT	CUDNN_DATA_FLOAT
DOUBLE_CONFIG	CUDNN_DATA_DOUBLE	CUDNN_DATA_DOUBLE

Specifying a separate algorithm can cause changes in performance, support and computation determinism. See the following for an exhaustive list of algorithm options and their respective supported parameters and deterministic behavior.

`dwDesc` may only have format CUDNN_TENSOR_NHWC when all of the following are true:

- ▶ `algo` is CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0 or CUDNN_CONVOLUTION_BWD_FILTER_ALGO_1
- ▶ `xDesc` and `dyDesc` is NHWC HWC-packed
- ▶ Data type configuration is PSEUDO_HALF_CONFIG or FLOAT_CONFIG
- ▶ The convolution is 2-dimensional

The following is an exhaustive list of algo support for 2-d convolutions.

- ▶ CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0
 - ▶ Deterministic: No
 - ▶ `xDesc` Format Support: All
 - ▶ `dyDesc` Format Support: NCHW CHW-packed
 - ▶ Data Type Config Support: All
- ▶ CUDNN_CONVOLUTION_BWD_FILTER_ALGO_1
 - ▶ Deterministic: Yes
 - ▶ `xDesc` Format Support: All
 - ▶ `dyDesc` Format Support: NCHW CHW-packed
 - ▶ Data Type Config Support: All
- ▶ CUDNN_CONVOLUTION_BWD_FILTER_ALGO_FFT

- ▶ Deterministic: Yes
- ▶ **xDesc** Format Support: NCHW CHW-packed
- ▶ **dyDesc** Format Support: NCHW CHW-packed
- ▶ Data Type Config Support: HALF_CONFIG, FLOAT_CONFIG
- ▶ Notes:
 - ▶ **xDesc**'s feature map height + 2 * **convDesc**'s zero-padding height must equal 256 or less
 - ▶ **xDesc**'s feature map width + 2 * **convDesc**'s zero-padding width must equal 256 or less
 - ▶ **convDesc**'s vertical and horizontal filter stride must equal 1
 - ▶ **dwDesc**'s filter height must be greater than **convDesc**'s zero-padding height
 - ▶ **dwDesc**'s filter width must be greater than **convDesc**'s zero-padding width
- ▶ **CUDNN_CONVOLUTION_BWD_FILTER_ALGO_3**
 - ▶ Deterministic: No
 - ▶ **xDesc** Format Support: All
 - ▶ **dyDesc** Format Support: NCHW CHW-packed
 - ▶ Data Type Config Support: All

The following is an exhaustive list of algo support for 3-d convolutions.

- ▶ **CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0**
 - ▶ Deterministic: No
 - ▶ **xDesc** Format Support: All
 - ▶ **dyDesc** Format Support: NCDHW CDHW-packed
 - ▶ Data Type Config Support: All
- ▶ **CUDNN_CONVOLUTION_BWD_FILTER_ALGO_3**
 - ▶ Deterministic: No
 - ▶ **xDesc** Format Support: NCDHW-fully-packed
 - ▶ **dyDesc** Format Support: NCDHW-fully-packed
 - ▶ Data Type Config Support: All

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The operation was launched successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ At least one of the following is NULL: handle, xDesc, dyDesc, convDesc, dwDesc, xData, dyData, dwData, alpha, beta ▶ xDesc and dyDesc have a non-matching number of dimensions ▶ xDesc and dwDesc have a non-matching number of dimensions ▶ xDesc has fewer than three number of dimensions

Return Value	Meaning
	<ul style="list-style-type: none"> ▶ <code>xDesc</code>, <code>dyDesc</code> and <code>dwDesc</code> have a non-matching data type. ▶ <code>xDesc</code> and <code>dwDesc</code> have a non-matching number of input feature maps per image.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	<p>At least one of the following conditions are met:</p> <ul style="list-style-type: none"> ▶ <code>xDesc</code> or <code>dyDesc</code> have negative tensor striding ▶ <code>xDesc</code>, <code>dyDesc</code> or <code>dwDesc</code> has a number of dimensions that is not 4 or 5 ▶ The chosen algo does not support the parameters provided; see above for exhaustive list of parameter support for each algo
<code>CUDNN_STATUS_MAPPING_ERROR</code>	An error occurs during the texture binding of the filter data.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.52. cudnnConvolutionBackwardFilter_v3

```

cudnnStatus_t
cudnnConvolutionBackwardFilter_v3( cudnnHandle_t          handle,
                                   const void             *alpha,
                                   const cudnnTensorDescriptor_t xDesc,
                                   const void             *x,
                                   const cudnnTensorDescriptor_t dyDesc,
                                   const void             *dy,
                                   const cudnnConvolutionDescriptor_t
convDesc,
                                   cudnnConvolutionBwdFilterAlgo_t algo,
                                   void
*workspace,
                                   size_t
workspaceSizeInBytes,
                                   const void             *beta,
                                   const cudnnFilterDescriptor_t dwDesc,
                                   void                  *dw )

```

This function is now equivalent to `cudnnConvolutionBackwardFilter`

4.53. cudnnConvolutionBackwardData_v2

```

cudnnStatus_t
cudnnConvolutionBackwardData_v2( cudnnHandle_t          handle,
                                   const void             *alpha,
                                   const cudnnFilterDescriptor_t wDesc,
                                   const void             *w,
                                   const cudnnTensorDescriptor_t dyDesc,
                                   const void             *dy,
                                   const cudnnConvolutionDescriptor_t convDesc,
                                   const void             *beta,
                                   const cudnnTensorDescriptor_t dxDesc,
                                   void                  *dx );

```

This function is equivalent to running `cudaConvolutionBackwardData` with the following additional parameters:

- ▶ algo: CUDNN_CONVOLUTION_BWD_DATA_ALGO_0
- ▶ workSpace: NULL
- ▶ workSpaceSizeInBytes: 0



This function is deprecated. `cudaConvolutionBackwardData` should be used instead. Refer to `cudaConvolutionBackwardData` for documentation.

4.54. cudaFindConvolutionBackwardDataAlgorithm

```

cudaStatus_t
cudaFindConvolutionBackwardDataAlgorithm(cudaHandle_t
    handle,
                                        const cudaFilterDescriptor_t
    wDesc,
                                        const cudaTensorDescriptor_t
    dyDesc,
                                        const cudaConvolutionDescriptor_t
    convDesc,
                                        const cudaTensorDescriptor_t
    dxDesc,
                                        const int
    requestedAlgoCount,
                                        int
    *returnedAlgoCount,
                                        cudaConvolutionBwdFilterAlgoPerf_t
    *perfResults );

```

This function attempts all cuDNN algorithms for `cudaConvolutionBackwardData` and outputs performance metrics to a user-allocated array of `cudaConvolutionBwdDataAlgoPerf_t`. These metrics are written in sorted fashion where the first element has the lowest compute time.



This function is host blocking.



It is recommend to run this function prior to allocating layer data; doing otherwise may needlessly inhibit some algorithm options due to resource usage.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
wDesc	input	Handle to a previously initialized filter descriptor.
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
convDesc	input	Previously initialized convolution descriptor.
dxDesc	input	Handle to the previously initialized output tensor descriptor.
requestedAlgoCount	input	The maximum number of elements to be stored in perfResults.

Param	In/out	Meaning
returnedAlgoCount	output	The number of output elements stored in perfResults.
perfResults	output	A user-allocated array to store performance metrics sorted ascending by compute time.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successful.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The handle is not allocated properly. ▶ The w, dy or dx descriptor is not allocated properly. ▶ The w, dy or dx descriptor has fewer than 1 dimension. ▶ Either returnedCount or perfResults is pointing to NULL. ▶ The requestedCount is less than 1.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	This function was unable to allocate memory to store sample input, filters and output.
<code>CUDNN_STATUS_INTERNAL_ERROR</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The function was unable to allocate necessary timing objects. ▶ The function was unable to deallocate necessary timing objects. ▶ The function was unable to deallocate sample input, filters and output.

4.55. cudnnGetConvolutionBackwardDataAlgorithm

```

cudnnStatus_t
cudnnGetConvolutionBackwardDataAlgorithm(
    handle,
    wDesc,
    dyDesc,
    convDesc,
    dxDesc,
    preference,
    memoryLimitInbytes,
    *algo
)
    cudnnHandle_t
    const cudnnFilterDescriptor_t
    const cudnnTensorDescriptor_t
    const cudnnConvolutionDescriptor_t
    const cudnnTensorDescriptor_t
    cudnnConvolutionBwdDataPreference_t
    size_t
    cudnnConvolutionBwdDataAlgo_t

```

This function serves as a heuristic for obtaining the best suited algorithm for **cudaConvolutionBackwardData** for the given layer specifications. Based on the input preference, this function will either return the fastest algorithm or the fastest algorithm within a given memory limit. For an exhaustive search for the fastest algorithm, please use **cudaFindConvolutionBackwardDataAlgorithm**.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
wDesc	input	Handle to a previously initialized filter descriptor.
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
convDesc	input	Previously initialized convolution descriptor.
dxDesc	input	Handle to the previously initialized output tensor descriptor.
preference	input	Enumerant to express the preference criteria in terms of memory requirement and speed.
memoryLimitInBytes	input	It is to specify the maximum amount of GPU memory the user is willing to use as a workspace. This is currently a placeholder and is not used.
algo	output	Enumerant that specifies which convolution algorithm should be used to compute the results according to the specified preference

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The query was successful.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The numbers of feature maps of the input tensor and output tensor differ. ▶ The dataType of the two tensor descriptors or the filter are different.

4.56. cudaGetConvolutionBackwardDataWorkspaceSize

```

cudaStatus_t
cudaGetConvolutionBackwardDataWorkspaceSize(
    cudaHandle_t
        handle,
        const cudaFilterDescriptor_t
            wDesc,
        const cudaTensorDescriptor_t
            dyDesc,
        const
            cudaConvolutionDescriptor_t convDesc,
        const cudaTensorDescriptor_t
            dxDesc,
        cudaConvolutionFwdAlgo_t
            algo,
        size_t
            *sizeInBytes
    )

```

This function returns the amount of GPU memory workspace the user needs to allocate to be able to call `cudaDnnConvolutionBackwardData` with the specified algorithm. The workspace allocated will then be passed to the routine `cudaDnnConvolutionBackwardData`. The specified algorithm can be the result of the call to `cudaDnnGetConvolutionBackwardDataAlgorithm` or can be chosen arbitrarily by the user. Note that not every algorithm is available for every configuration of the input tensor and/or every configuration of the convolution descriptor.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
wDesc	input	Handle to a previously initialized filter descriptor.
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
convDesc	input	Previously initialized convolution descriptor.
dxDesc	input	Handle to the previously initialized output tensor descriptor.
algo	input	Enumerant that specifies the chosen convolution algorithm
sizeInBytes	output	Amount of GPU memory needed as workspace to be able to execute a forward convolution with the specified <code>algo</code>

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The query was successful.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The numbers of feature maps of the input tensor and output tensor differ. ▶ The <code>dataType</code> of the two tensor descriptors or the filter are different.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The combination of the tensor descriptors, filter descriptor and convolution descriptor is not supported for the specified algorithm.

4.57. cudnnConvolutionBackwardData

```

cudnnStatus_t
cudnnConvolutionBackwardData( cudnnHandle_t          handle,
                              const void            *alpha,
                              const cudnnFilterDescriptor_t wDesc,
                              const void            *w,
                              const cudnnTensorDescriptor_t dyDesc,
                              const void            *dy,
                              const cudnnConvolutionDescriptor_t convDesc,
                              cudnnConvolutionBwdDataAlgo_t algo,
                              void                  *workSpace,
                              size_t                workSpaceSizeInBytes,
                              const void            *beta,
                              const cudnnTensorDescriptor_t dxDesc,
                              void                  *dx );

```

This function computes the convolution gradient with respect to the output tensor using the specified **algo**, returning results in **gradDesc**. Scaling factors **alpha** and **beta** can be used to scale the input tensor and the output tensor respectively.



This routine replaces `cudnnConvolutionBackwardData_v2`

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
wDesc	input	Handle to a previously initialized filter descriptor.
w	input	Data pointer to GPU memory associated with the filter descriptor wDesc .
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
dy	input	Data pointer to GPU memory associated with the input differential tensor descriptor dyDesc .
convDesc	input	Previously initialized convolution descriptor.
algo	input	Enumerant that specifies which backward data convolution algorithm should be used to compute the results
workSpace	input	Data pointer to GPU memory to a workspace needed to able to execute the specified algorithm. If no workspace is needed for a particular algorithm, that pointer can be nil
workSpaceSizeInBytes	input	Specifies the size in bytes of the provided workSpace
dxDesc	input	Handle to the previously initialized output tensor descriptor.
dx	input/output	Data pointer to GPU memory associated with the output tensor descriptor dxDesc that carries the result.

This function supports only three specific combinations of data types for **wDesc**, **dyDesc**, **convDesc** and **dxDesc**. See the following for an exhaustive list of these configurations.

Data Type Configurations	wDesc's, dyDesc's and dxDesc's Data Type	convDesc's Data Type
HALF_CONFIG	CUDNN_DATA_HALF	CUDNN_DATA_FLOAT
FLOAT_CONFIG	CUDNN_DATA_FLOAT	CUDNN_DATA_FLOAT
DOUBLE_CONFIG	CUDNN_DATA_DOUBLE	CUDNN_DATA_DOUBLE

Specifying a separate algorithm can cause changes in performance, support and computation determinism. See the following for an exhaustive list of algorithm options and their respective supported parameters and deterministic behavior.

wDesc may only have format CUDNN_TENSOR_NHWC when all of the following are true:

- ▶ **algo** is CUDNN_CONVOLUTION_BWD_DATA_ALGO_1
- ▶ **dyDesc** and **dxDesc** is NHWC HWC-packed
- ▶ Data type configuration is PSEUDO_HALF_CONFIG or FLOAT_CONFIG
- ▶ The convolution is 2-dimensional

The following is an exhaustive list of algo support for 2-d convolutions.

- ▶ **CUDNN_CONVOLUTION_BWD_DATA_ALGO_0**
 - ▶ Deterministic: No
 - ▶ **dyDesc** Format Support: NCHW CHW-packed
 - ▶ **dxDesc** Format Support: All
 - ▶ Data Type Config Support: All
- ▶ **CUDNN_CONVOLUTION_BWD_DATA_ALGO_1**
 - ▶ Deterministic: Yes
 - ▶ **dyDesc** Format Support: NCHW CHW-packed
 - ▶ **dxDesc** Format Support: All
 - ▶ Data Type Config Support: All
- ▶ **CUDNN_CONVOLUTION_BWD_DATA_ALGO_FFT**
 - ▶ Deterministic: Yes
 - ▶ **dyDesc** Format Support: NCHW CHW-packed
 - ▶ **dxDesc** Format Support: NCHW CHW-packed
 - ▶ Data Type Config Support: HALF_CONFIG, FLOAT_CONFIG
 - ▶ Notes:
 - ▶ **dxDesc**'s feature map height + 2 * **convDesc**'s zero-padding height must equal 256 or less
 - ▶ **dxDesc**'s feature map width + 2 * **convDesc**'s zero-padding width must equal 256 or less
 - ▶ **convDesc**'s vertical and horizontal filter stride must equal 1
 - ▶ **wDesc**'s filter height must be greater than **convDesc**'s zero-padding height

- ▶ **wDesc**'s filter width must be greater than **convDesc**'s zero-padding width
- ▶ **CUDNN_CONVOLUTION_BWD_DATA_ALGO_FFT_TILING**
 - ▶ Deterministic: Yes
 - ▶ **dyDesc** Format Support: NCHW CHW-packed
 - ▶ **dxDesc** Format Support: NCHW CHW-packed
 - ▶ Data Type Config Support: HALF_CONFIG, FLOAT_CONFIG
 - ▶ Notes:
 - ▶ **wDesc**'s filter height must equal 32 or less
 - ▶ **wDesc**'s filter width must equal 32 or less
 - ▶ **convDesc**'s vertical and horizontal filter stride must equal 1
 - ▶ **wDesc**'s filter height must be greater than **convDesc**'s zero-padding height
 - ▶ **wDesc**'s filter width must be greater than **convDesc**'s zero-padding width

The following is an exhaustive list of algo support for 3-d convolutions.

- ▶ **CUDNN_CONVOLUTION_BWD_DATA_ALGO_0**
 - ▶ Deterministic: No
 - ▶ **dyDesc** Format Support: NCDHW CDHW-packed
 - ▶ **dxDesc** Format Support: All
 - ▶ Data Type Config Support: All
- ▶ **CUDNN_CONVOLUTION_BWD_DATA_ALGO_1**
 - ▶ Deterministic: Yes
 - ▶ **dyDesc** Format Support: NCDHW-fully-packed
 - ▶ **dxDesc** Format Support: NCDHW-fully-packed
 - ▶ Data Type Config Support: All

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The operation was launched successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ At least one of the following is NULL: handle, dyDesc, wDesc, convDesc, dxDesc, dy, w, dx, alpha, beta ▶ wDesc and dyDesc have a non-matching number of dimensions ▶ wDesc and dxDesc have a non-matching number of dimensions ▶ wDesc has fewer than three number of dimensions ▶ wDesc, dxDesc and dyDesc have a non-matching data type. ▶ wDesc and dxDesc have a non-matching number of input feature maps per image.

Return Value	Meaning
	<ul style="list-style-type: none"> ▶ <code>dyDescs</code>'s spatial sizes do not match with the expected size as determined by <code>cudaGetConvolutionNdForwardOutputDim</code>
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	<p>At least one of the following conditions are met:</p> <ul style="list-style-type: none"> ▶ <code>dyDesc</code> or <code>dxDesc</code> have negative tensor striding ▶ <code>dyDesc</code>, <code>wDesc</code> or <code>dxDesc</code> has a number of dimensions that is not 4 or 5 ▶ The chosen algo does not support the parameters provided; see above for exhaustive list of parameter support for each algo
<code>CUDNN_STATUS_MAPPING_ERROR</code>	An error occurs during the texture binding of the filter data or the input differential tensor data
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.58. `cudaConvolutionBackwardData_v3`

```

cudaStatus_t
cudaConvolutionBackwardData_v3( cudaHandle_t          handle,
                                const void           *alpha,
                                const cudaFilterDescriptor_t wDesc,
                                const void           *w,
                                const cudaTensorDescriptor_t dyDesc,
                                const void           *dy,
                                const cudaConvolutionDescriptor_t convDesc,
                                cudaConvolutionBwdDataAlgo_t algo,
                                void                 *workspace,
                                size_t               size_t,
                                workspaceSizeInBytes,
                                const void           *beta,
                                const cudaTensorDescriptor_t dxDesc,
                                void                 *dx );

```

This function is now equivalent to `cudaConvolutionBackwardData`.

4.59. `cudaSoftmaxForward`

```

cudaStatus_t
cudaSoftmaxForward( cudaHandle_t          handle,
                    cudaSoftmaxAlgorithm_t algorithm,
                    cudaSoftmaxMode_t     mode,
                    const void           *alpha,
                    const cudaTensorDescriptor_t xDesc,
                    const void           *x,
                    const void           *beta,
                    const cudaTensorDescriptor_t yDesc,
                    void                 *y );

```

This routine computes the softmax function.



All tensor formats are supported for all modes and algorithms with 4 and 5D tensors. Performance is expected to be highest with `NCHW` `fully-packed` tensors. For more than 5 dimensions tensors must be packed in their spatial dimensions

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
algorithm	input	Enumerant to specify the softmax algorithm.
mode	input	Enumerant to specify the softmax mode.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc	input	Handle to the previously initialized input tensor descriptor.
x	input	Data pointer to GPU memory associated with the tensor descriptor <code>xDesc</code> .
yDesc	input	Handle to the previously initialized output tensor descriptor.
y	output	Data pointer to GPU memory associated with the output tensor descriptor <code>yDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The dimensions <code>n, c, h, w</code> of the input tensor and output tensors differ. ▶ The <code>datatype</code> of the input tensor and output tensors differ. ▶ The parameters <code>algorithm</code> or <code>mode</code> have an invalid enumerant value.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.60. cudnnSoftmaxBackward

```

cudnnStatus_t
cudnnSoftmaxBackward( cudnnHandle_t          handle,
                      cudnnSoftmaxAlgorithm_t algorithm,
                      cudnnSoftmaxMode_t     mode,
                      const void             *alpha,
                      const cudnnTensorDescriptor_t yDesc,
                      const void             *yData,
                      const cudnnTensorDescriptor_t dyDesc,
                      const void             *dy,
                      const void             *beta,
                      const cudnnTensorDescriptor_t dxDesc,
                      void                   *dx );

```

This routine computes the gradient of the softmax function.



In-place operation is allowed for this routine; i.e., `dy` and `dx` pointers may be equal. However, this requires `dyDesc` and `dxDesc` descriptors to be identical (particularly, the strides of the input and output must match for in-place operation to be allowed).



All tensor formats are supported for all modes and algorithms with 4 and 5D tensors. Performance is expected to be highest with **NCHW fully-packed** tensors. For more than 5 dimensions tensors must be packed in their spatial dimensions

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
algorithm	input	Enumerant to specify the softmax algorithm.
mode	input	Enumerant to specify the softmax mode.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $\text{dstValue} = \text{alpha}[0] * \text{result} + \text{beta}[0] * \text{priorDstValue}$. Please refer to this section for additional details.
yDesc	input	Handle to the previously initialized input tensor descriptor.
y	input	Data pointer to GPU memory associated with the tensor descriptor <code>yDesc</code> .
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
dy	input	Data pointer to GPU memory associated with the tensor descriptor <code>dyData</code> .
dxDesc	input	Handle to the previously initialized output differential tensor descriptor.
dx	output	Data pointer to GPU memory associated with the output tensor descriptor <code>dxDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met:

Return Value	Meaning
	<ul style="list-style-type: none"> ▶ The dimensions <i>n, c, h, w</i> of the <i>yDesc</i>, <i>dyDesc</i> and <i>dxDesc</i> tensors differ. ▶ The strides <i>nStride</i>, <i>cStride</i>, <i>hStride</i>, <i>wStride</i> of the <i>yDesc</i> and <i>dyDesc</i> tensors differ. ▶ The <i>datatype</i> of the three tensors differs.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.61. cudnnCreatePoolingDescriptor

```

cudnnStatus_t cudnnCreatePoolingDescriptor( cudnnPoolingDescriptor_t*
poolingDesc )

```

This function creates a pooling descriptor object by allocating the memory needed to hold its opaque structure,

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was created successfully.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	The resources could not be allocated.

4.62. cudnnSetPooling2dDescriptor

```

cudnnStatus_t
cudnnSetPooling2dDescriptor( cudnnPoolingDescriptor_t poolingDesc,
cudnnPoolingMode_t mode,
int windowHeight,
int windowWidth,
int verticalPadding,
int horizontalPadding,
int verticalStride,
int horizontalStride )

```

This function initializes a previously created generic pooling descriptor object into a 2D description.

Param	In/out	Meaning
<code>poolingDesc</code>	input/output	Handle to a previously created pooling descriptor.
<code>mode</code>	input	Enumerant to specify the pooling mode.
<code>windowHeight</code>	input	Height of the pooling window.
<code>windowWidth</code>	input	Width of the pooling window.
<code>verticalPadding</code>	input	Size of vertical padding.
<code>horizontalPadding</code>	input	Size of horizontal padding
<code>verticalStride</code>	input	Pooling vertical stride.

Param	In/out	Meaning
horizontalStride	output	Pooling horizontal stride.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the parameters <code>windowHeight</code> , <code>windowWidth</code> , <code>verticalStride</code> , <code>horizontalStride</code> is negative or <code>mode</code> has an invalid enumerant value.

4.63. cudnnGetPooling2dDescriptor

```

cudnnStatus_t
cudnnGetPooling2dDescriptor( const cudnnPoolingDescriptor_t poolingDesc,
                             cudnnPoolingMode_t *mode,
                             int *windowHeight,
                             int *windowWidth,
                             int *verticalPadding,
                             int *horizontalPadding,
                             int *verticalStride,
                             int *horizontalStride )

```

This function queries a previously created 2D pooling descriptor object.

Param	In/out	Meaning
poolingDesc	input	Handle to a previously created pooling descriptor.
mode	output	Enumerant to specify the pooling mode.
windowHeight	output	Height of the pooling window.
windowWidth	output	Width of the pooling window.
verticalPadding	output	Size of vertical padding.
horizontalPadding	output	Size of horizontal padding.
verticalStride	output	Pooling vertical stride.
horizontalStride	output	Pooling horizontal stride.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.

4.64. cudnnSetPoolingNdDescriptor

```

cudnnStatus_t
cudnnSetPoolingNdDescriptor( cudnnPoolingDescriptor_t poolingDesc,
                             cudnnPoolingMode_t mode,
                             int nbDims,
                             int windowDimA[],
                             int paddingA[],
                             int strideA[] )

```

This function initializes a previously created generic pooling descriptor object.

Param	In/out	Meaning
poolingDesc	input/ output	Handle to a previously created pooling descriptor.
mode	input	Enumerant to specify the pooling mode.
nbDims	input	Dimension of the pooling operation.
windowDimA	output	Array of dimension <code>nbDims</code> containing the window size for each dimension.
paddingA	output	Array of dimension <code>nbDims</code> containing the padding size for each dimension.
strideA	output	Array of dimension <code>nbDims</code> containing the striding size for each dimension.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the elements of the arrays <code>windowDimA</code> , <code>paddingA</code> or <code>strideA</code> is negative or <code>mode</code> has an invalid enumerant value.

4.65. cudnnGetPoolingNdDescriptor

```

cudnnStatus_t
cudnnGetPoolingNdDescriptor( const cudnnPoolingDescriptor_t poolingDesc,
                             int nbDimsRequested,
                             cudnnPoolingMode_t *mode,
                             int *nbDims,
                             int windowDimA[],
                             int paddingA[],
                             int strideA[] )

```

This function queries a previously initialized generic pooling descriptor object.

Param	In/ out	Meaning
poolingDesc	input	Handle to a previously created pooling descriptor.

Param	In/out	Meaning
nbDimsRequested	input	Dimension of the expected pooling descriptor. It is also the minimum size of the arrays <code>windowDimA</code> , <code>paddingA</code> and <code>strideA</code> in order to be able to hold the results
mode	output	Enumerant to specify the pooling mode.
nbDims	output	Actual dimension of the pooling descriptor.
windowDimA	output	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the window parameters from the provided pooling descriptor.
paddingA	output	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the padding parameters from the provided pooling descriptor.
strideA	output	Array of dimension at least <code>nbDimsRequested</code> that will be filled with the stride parameters from the provided pooling descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was queried successfully.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The parameter <code>nbDimsRequested</code> is greater than <code>CUDNN_DIM_MAX</code> .

4.66. cudnnSetPooling2dDescriptor_v4

```

cudnnStatus_t
cudnnSetPooling2dDescriptor_v4( cudnnPoolingDescriptor_t poolingDesc,
                                cudnnPoolingMode_t mode,
                                cudnnNanPropagation_t maxpoolingNanOpt,
                                int windowHeight,
                                int windowWidth,
                                int verticalPadding,
                                int horizontalPadding,
                                int verticalStride,
                                int horizontalStride )

```

This function initializes a previously created generic pooling descriptor object into a 2D description.

Param	In/out	Meaning
poolingDesc	input/output	Handle to a previously created pooling descriptor.
mode	input	Enumerant to specify the pooling mode.
maxpoolingNanOpt	input	Enumerant to specify the Nan propagation mode.
windowHeight	input	Height of the pooling window.
windowWidth	input	Width of the pooling window.
verticalPadding	input	Size of vertical padding.

Param	In/out	Meaning
horizontalPadding	input	Size of horizontal padding
verticalStride	input	Pooling vertical stride.
horizontalStride	input	Pooling horizontal stride.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the parameters <code>windowHeight</code> , <code>windowWidth</code> , <code>verticalStride</code> , <code>horizontalStride</code> is negative or <code>mode</code> or <code>maxpoolingNanOpt</code> has an invalid enumerant value.

4.67. cudnnGetPooling2dDescriptor_v4

```

cudnnStatus_t
cudnnGetPooling2dDescriptor_v4( const cudnnPoolingDescriptor_t poolingDesc,
                                cudnnPoolingMode_t *mode,
                                cudnnNanPropagation_t *maxpoolingNanOpt,
                                int *windowHeight,
                                int *windowWidth,
                                int *verticalPadding,
                                int *horizontalPadding,
                                int *verticalStride,
                                int *horizontalStride )

```

This function queries a previously created 2D pooling descriptor object.

Param	In/out	Meaning
poolingDesc	input	Handle to a previously created pooling descriptor.
mode	output	Enumerant to specify the pooling mode.
maxpoolingNanOpt	output	Enumerant to specify the Nan propagation mode.
windowHeight	output	Height of the pooling window.
windowWidth	output	Width of the pooling window.
verticalPadding	output	Size of vertical padding.
horizontalPadding	output	Size of horizontal padding.
verticalStride	output	Pooling vertical stride.
horizontalStride	output	Pooling horizontal stride.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.

4.68. cudnnSetPoolingNdDescriptor_v4

```

cudnnStatus_t
cudnnSetPoolingNdDescriptor_v4( cudnnPoolingDescriptor_t poolingDesc,
                                cudnnPoolingMode_t mode,
                                cudnnNanPropagation_t maxpoolingNanOpt,
                                int nbDims,
                                int windowDimA[],
                                int paddingA[],
                                int strideA[] )

```

This function initializes a previously created generic pooling descriptor object.

Param	In/out	Meaning
poolingDesc	input/ output	Handle to a previously created pooling descriptor.
mode	input	Enumerant to specify the pooling mode.
maxpoolingNanOpt	input	Enumerant to specify the Nan propagation mode.
nbDims	input	Dimension of the pooling operation.
windowDimA	output	Array of dimension <code>nbDims</code> containing the window size for each dimension.
paddingA	output	Array of dimension <code>nbDims</code> containing the padding size for each dimension.
strideA	output	Array of dimension <code>nbDims</code> containing the striding size for each dimension.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the elements of the arrays <code>windowDimA</code> , <code>paddingA</code> or <code>strideA</code> is negative OR <code>mode</code> OR <code>maxpoolingNanOpt</code> has an invalid enumerant value.

4.69. cudnnGetPoolingNdDescriptor_v4

```

cudnnStatus_t
cudnnGetPoolingNdDescriptor_v4( const cudnnPoolingDescriptor_t poolingDesc,
                                int nbDimsRequested,
                                cudnnPoolingMode_t *mode,
                                cudnnNanPropagation_t *maxpoolingNanOpt,
                                int *nbDims,
                                int windowDimA[],
                                int paddingA[],
                                int strideA[] )

```

This function queries a previously initialized generic pooling descriptor object.

Param	In/out	Meaning
poolingDesc	input	Handle to a previously created pooling descriptor.
nbDimsRequested	input	Dimension of the expected pooling descriptor. It is also the minimum size of the arrays <code>windowDimA</code> , <code>paddingA</code> and <code>strideA</code> in order to be able to hold the results
mode	output	Enumerant to specify the pooling mode.
maxpoolingNanOutput	output	Enumerant to specify the Nan propagation mode.
nbDims	output	Actual dimension of the pooling descriptor.
windowDimA	output	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the window parameters from the provided pooling descriptor.
paddingA	output	Array of dimension of at least <code>nbDimsRequested</code> that will be filled with the padding parameters from the provided pooling descriptor.
strideA	output	Array of dimension at least <code>nbDimsRequested</code> that will be filled with the stride parameters from the provided pooling descriptor.

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was queried successfully.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The parameter <code>nbDimsRequested</code> is greater than <code>CUDNN_DIM_MAX</code> .

4.70. cudnnDestroyPoolingDescriptor

```

cudnnStatus_t cudnnDestroyPoolingDescriptor( cudnnPoolingDescriptor_t
poolingDesc )

```

This function destroys a previously created pooling descriptor object.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was destroyed successfully.

4.71. cudnnGetPooling2dForwardOutputDim

```

cudnnStatus_t
cudnnGetPooling2dForwardOutputDim( const cudnnPoolingDescriptor_t poolingDesc,
const cudnnTensorDescriptor_t inputDesc,
int *outN,
int *outC,
int *outH,
int *outW )

```

This function provides the output dimensions of a tensor after 2d pooling has been applied

Each dimension **h** and **w** of the output images is computed as followed:

```
outputDim = 1 + (inputDim + 2*padding - windowDim)/poolingStride;
```

Param	In/out	Meaning
poolingDesc	input	Handle to a previously initialized pooling descriptor.
inputDesc	input	Handle to the previously initialized input tensor descriptor.
N	output	Number of images in the output
C	output	Number of channels in the output
H	output	Height of images in the output
W	output	Width of images in the output

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The function launched successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ poolingDesc has not been initialized. ▶ poolingDesc or inputDesc has an invalid number of dimensions (2 and 4 respectively are required).

4.72. cudnnGetPoolingNdForwardOutputDim

```

cudnnStatus_t
cudnnGetPoolingNdForwardOutputDim( const cudnnPoolingDescriptor_t poolingDesc,
                                   const cudnnTensorDescriptor_t inputDesc,
                                   int nbDims,
                                   int outDimA[] )

```

This function provides the output dimensions of a tensor after Nd pooling has been applied

Each dimension of the **(nbDims-2) -D** images of the output tensor is computed as followed:

```
outputDim = 1 + (inputDim + 2*padding - windowDim)/poolingStride;
```

Param	In/out	Meaning
poolingDesc	input	Handle to a previously initialized pooling descriptor.
inputDesc	input	Handle to the previously initialized input tensor descriptor.
nbDims	input	Number of dimensions in which pooling is to be applied.
outDimA	output	Array of nbDims output dimensions

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ <code>poolingDesc</code> has not been initialized. ▶ The value of <code>nbDims</code> is inconsistent with the dimensionality of <code>poolingDesc</code> and <code>inputDesc</code>.

4.73. cudnnPoolingForward

```

cudnnStatus_t
cudnnPoolingForward( cudnnHandle_t      handle,
                    const cudnnPoolingDescriptor_t poolingDesc,
                    const void          *alpha,
                    const cudnnTensorDescriptor_t xDesc,
                    const void          *x,
                    const void          *beta,
                    const cudnnTensorDescriptor_t yDesc,
                    void                *y );

```

This function computes pooling of input values (i.e., the maximum or average of several adjacent values) to produce an output with smaller height and/or width.



All tensor formats are supported, best performance is expected when using **HW-packed** tensors. Only 2 and 3 spatial dimensions are allowed.

Param	In/out	Meaning
<code>handle</code>	input	Handle to a previously created cuDNN context.
<code>poolingDesc</code>	input	Handle to a previously initialized pooling descriptor.
<code>alpha, beta</code>	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
<code>xDesc</code>	input	Handle to the previously initialized input tensor descriptor.
<code>x</code>	input	Data pointer to GPU memory associated with the tensor descriptor <code>xDesc</code> .
<code>yDesc</code>	input	Handle to the previously initialized output tensor descriptor.
<code>y</code>	output	Data pointer to GPU memory associated with the output tensor descriptor <code>yDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.

Return Value	Meaning
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The dimensions <code>n</code>, <code>c</code> of the input tensor and output tensors differ. ▶ The <code>datatype</code> of the input tensor and output tensors differs.
CUDNN_STATUS_NOT_SUPPORTED	The <code>wStride</code> of input tensor or output tensor is not 1.
CUDNN_STATUS_EXECUTION_FAILED	The function failed to launch on the GPU.

4.74. cudnnPoolingBackward

```

cudnnStatus_t
cudnnPoolingBackward( cudnnHandle_t handle,
                      const cudnnPoolingDescriptor_t poolingDesc,
                      const void *alpha,
                      const cudnnTensorDescriptor_t yDesc,
                      const void *y,
                      const cudnnTensorDescriptor_t dyDesc,
                      const void *dy,
                      const cudnnTensorDescriptor_t xDesc,
                      const void *xData,
                      const void *beta,
                      const cudnnTensorDescriptor_t dxDesc,
                      void *dx )

```

This function computes the gradient of a pooling operation.



All tensor formats are supported, best performance is expected when using **HW-packed** tensors. Only 2 and 3 spatial dimensions are allowed

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
poolingDesc	input	Handle to the previously initialized pooling descriptor.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
yDesc	input	Handle to the previously initialized input tensor descriptor.
y	input	Data pointer to GPU memory associated with the tensor descriptor <code>yDesc</code> .
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
dy	input	Data pointer to GPU memory associated with the tensor descriptor <code>dyData</code> .
xDesc	input	Handle to the previously initialized output tensor descriptor.
x	input	Data pointer to GPU memory associated with the output tensor descriptor <code>xDesc</code> .

Param	In/out	Meaning
dxDesc	input	Handle to the previously initialized output differential tensor descriptor.
dx	output	Data pointer to GPU memory associated with the output tensor descriptor <code>dxDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The dimensions <code>n, c, h, w</code> of the <code>yDesc</code> and <code>dyDesc</code> tensors differ. ▶ The strides <code>nStride, cStride, hStride, wStride</code> of the <code>yDesc</code> and <code>dyDesc</code> tensors differ. ▶ The dimensions <code>n, c, h, w</code> of the <code>dxDesc</code> and <code>dxData</code> tensors differ. ▶ The strides <code>nStride, cStride, hStride, wStride</code> of the <code>xDesc</code> and <code>dxDesc</code> tensors differ. ▶ The <code>datatype</code> of the four tensors differ.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	The <code>wStride</code> of input tensor or output tensor is not 1.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.75. cudnnActivationForward

```

cudnnStatus_t
cudnnActivationForward( cudnnHandle_t          handle,
                       cudnnActivationMode_t  mode,
                       const void             *alpha,
                       const cudnnTensorDescriptor_t xDesc,
                       const void             *xData,
                       const void             *beta,
                       const cudnnTensorDescriptor_t yDesc,
                       void                   *yData )

```

This routine applies a specified neuron activation function element-wise over each input value.



In-place operation is allowed for this routine; i.e., `xData` and `yData` pointers may be equal. However, this requires `xDesc` and `yDesc` descriptors to be identical

(particularly, the strides of the input and output must match for in-place operation to be allowed).



All tensor formats are supported for 4 and 5 dimensions, however best performance is obtained when the strides of `xDesc` and `yDesc` are equal and `HW-packed`. For more than 5 dimensions the tensors must have their spatial dimensions packed.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
mode	input	Enumerant to specify the activation mode.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc	input	Handle to the previously initialized input tensor descriptor.
x	input	Data pointer to GPU memory associated with the tensor descriptor <code>xDesc</code> .
yDesc	input	Handle to the previously initialized output tensor descriptor.
y	output	Data pointer to GPU memory associated with the output tensor descriptor <code>yDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The parameter <code>mode</code> has an invalid enumerant value. ▶ The dimensions <code>n, c, h, w</code> of the input tensor and output tensors differ. ▶ The <code>datatype</code> of the input tensor and output tensors differs. ▶ The strides <code>nStride, cStride, hStride, wStride</code> of the input tensor and output tensors differ and in-place operation is used (i.e., <code>x</code> and <code>y</code> pointers are equal).
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.76. cudnnActivationBackward

```

cudnnStatus_t
cudnnActivationBackward( cudnnHandle_t          handle,
                        cudnnActivationMode_t  mode,
                        const void            *alpha,
                        cudnnTensorDescriptor_t yDesc,
                        const void            *y,
                        cudnnTensorDescriptor_t dyDesc,
                        const void            *dy,
                        cudnnTensorDescriptor_t xDesc,
                        const void            *x,
                        const void            *beta,
                        cudnnTensorDescriptor_t dxDesc,
                        void                  *dx );

```

This routine computes the gradient of a neuron activation function.



In-place operation is allowed for this routine; i.e. `dy` and `dx` pointers may be equal. However, this requires the corresponding tensor descriptors to be identical (particularly, the strides of the input and output must match for in-place operation to be allowed).



All tensor formats are supported for 4 and 5 dimensions, however best performance is obtained when the strides of `yDesc` and `xDesc` are equal and **HW-packed**. For more than 5 dimensions the tensors must have their spatial dimensions packed.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
mode	input	Enumerant to specify the activation mode.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
yDesc	input	Handle to the previously initialized input tensor descriptor.
y	input	Data pointer to GPU memory associated with the tensor descriptor <code>yDesc</code> .
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
dy	input	Data pointer to GPU memory associated with the tensor descriptor <code>dyDesc</code> .
xDesc	input	Handle to the previously initialized output tensor descriptor.
x	input	Data pointer to GPU memory associated with the output tensor descriptor <code>xDesc</code> .
dxDesc	input	Handle to the previously initialized output differential tensor descriptor.
dx	output	Data pointer to GPU memory associated with the output tensor descriptor <code>dxDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The parameter <code>mode</code> has an invalid enumerant value. ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the input differential tensor and output differential tensors differ and in-place operation is used.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The dimensions <code>n</code>, <code>c</code>, <code>h</code>, <code>w</code> of the input tensor and output tensors differ. ▶ The <code>datatype</code> of the input tensor and output tensors differs. ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the input tensor and the input differential tensor differ. ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the output tensor and the output differential tensor differ.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.77. cudnnCreateActivationDescriptor

```

cudnnStatus_t
cudnnCreateActivationDescriptor( cudnnActivationDescriptor_t
*activationDesc )

```

This function creates a activation descriptor object by allocating the memory needed to hold its opaque structure.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was created successfully.
<code>CUDNN_STATUS_ALLOC_FAILED</code>	The resources could not be allocated.

4.78. cudnnSetActivationDescriptor

```

cudnnStatus_t
cudnnSetActivationDescriptor( cudnnActivationDescriptor_t
activationDesc,
                             cudnnActivationMode_t      mode,
                             cudnnNanPropagation_t      reluNanOpt,
                             double                      reluCeiling )

```

This function initializes a previously created generic activation descriptor object.

Param	In/out	Meaning
activationDesc	input/output	Handle to a previously created pooling descriptor.
mode	input	Enumerant to specify the activation mode.
reluNanOpt	input	Enumerant to specify the Nan propagation mode.
reluCeiling	input	floating point number to specify the clipping threshod when the activation mode is set to <code>CUDNN_ACTIVATION_CLIPPED_RELU</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was set successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	<code>reluNanOpt</code> has an invalid enumerant value.

4.79. cudnnGetActivationDescriptor

```

cudnnStatus_t
cudnnGetActivationDescriptor( const cudnnActivationDescriptor_t
activationDesc,
                             cudnnActivationMode_t           *mode,
                             cudnnNanPropagation_t
*reluNanOpt,
                             double
*reluCeiling )
    
```

This function queries a previously initialized generic activation descriptor object.

Param	In/out	Meaning
activationDesc	input	Handle to a previously created activation descriptor.
mode	output	Enumerant to specify the activation mode.
reluNanOpt	output	Enumerant to specify the Nan propagation mode.
reluCeiling	output	floating point number to specify the clipping threshod when the activation mode is set to <code>CUDNN_ACTIVATION_CLIPPED_RELU</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was queried successfully.

4.80. cudnnDestroyActivationDescriptor

```

cudnnStatus_t
cudnnDestroyActivationDescriptor( cudnnActivationDescriptor_t
activationDesc )
    
```


This function destroys a previously created activation descriptor object.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was destroyed successfully.

4.81. cudnnActivationForward_v4

```

cudnnStatus_t
cudnnActivationForward_v4( cudnnHandle_t handle,
                          cudnnActivationDescriptor_t activationDesc,
                          const void *alpha,
                          const cudnnTensorDescriptor_t srcDesc,
                          const void *srcData,
                          const void *beta,
                          const cudnnTensorDescriptor_t destDesc,
                          void *destData )

```

This routine applies a specified neuron activation function element-wise over each input value.



In-place operation is allowed for this routine; i.e., `xData` and `yData` pointers may be equal. However, this requires `xDesc` and `yDesc` descriptors to be identical (particularly, the strides of the input and output must match for in-place operation to be allowed).



All tensor formats are supported for 4 and 5 dimensions, however best performance is obtained when the strides of `xDesc` and `yDesc` are equal and **HW-packed**. For more than 5 dimensions the tensors must have their spatial dimensions packed.

Param	In/ out	Meaning
handle	input	Handle to a previously created cuDNN context.
activationDesc	input	Activation descriptor.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: $dstValue = alpha[0]*result + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc	input	Handle to the previously initialized input tensor descriptor.
x	input	Data pointer to GPU memory associated with the tensor descriptor <code>xDesc</code> .
yDesc	input	Handle to the previously initialized output tensor descriptor.
y	output	Data pointer to GPU memory associated with the output tensor descriptor <code>yDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The function launched successfully.


Return Value	Meaning
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The parameter <code>mode</code> has an invalid enumerant value. ▶ The dimensions <code>n</code>, <code>c</code>, <code>h</code>, <code>w</code> of the input tensor and output tensors differ. ▶ The <code>datatype</code> of the input tensor and output tensors differs. ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the input tensor and output tensors differ and in-place operation is used (i.e., <code>x</code> and <code>y</code> pointers are equal).
CUDNN_STATUS_EXECUTION_FAILED	The function failed to launch on the GPU.


4.82. cudnnActivationBackward_v4

```

cudnnStatus_t
cudnnActivationBackward_v4( cudnnHandle_t          handle,
                           cudnnActivationDescriptor_t activationDesc,
                           const void             *alpha,
                           const cudnnTensorDescriptor_t srcDesc,
                           const void             *srcData,
                           const cudnnTensorDescriptor_t srcDiffDesc,
                           const void             *srcDiffData,
                           const cudnnTensorDescriptor_t destDesc,
                           const void             *destData,
                           const void             *beta,
                           const cudnnTensorDescriptor_t destDiffDesc,
                           void                   *destDiffData)
    
```

This routine computes the gradient of a neuron activation function.

 In-place operation is allowed for this routine; i.e. `dy` and `dx` pointers may be equal. However, this requires the corresponding tensor descriptors to be identical (particularly, the strides of the input and output must match for in-place operation to be allowed).

 All tensor formats are supported for 4 and 5 dimensions, however best performance is obtained when the strides of `yDesc` and `xDesc` are equal and `HW-packed`. For more than 5 dimensions the tensors must have their spatial dimensions packed.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN context.
activationDesc	input	Activation descriptor.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the computation result with prior value in the output layer as follows: <code>dstValue = alpha[0]*result + beta[0]*priorDstValue</code> . Please refer to this section for additional details.

Param	In/out	Meaning
yDesc	input	Handle to the previously initialized input tensor descriptor.
y	input	Data pointer to GPU memory associated with the tensor descriptor <code>yDesc</code> .
dyDesc	input	Handle to the previously initialized input differential tensor descriptor.
dy	input	Data pointer to GPU memory associated with the tensor descriptor <code>dyDesc</code> .
xDesc	input	Handle to the previously initialized output tensor descriptor.
x	input	Data pointer to GPU memory associated with the output tensor descriptor <code>xDesc</code> .
dxDesc	input	Handle to the previously initialized output differential tensor descriptor.
dx	output	Data pointer to GPU memory associated with the output tensor descriptor <code>dxDesc</code> .

The possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The function launched successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the input differential tensor and output differential tensors differ and in-place operation is used.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ The dimensions <code>n</code>, <code>c</code>, <code>h</code>, <code>w</code> of the input tensor and output tensors differ. ▶ The <code>datatype</code> of the input tensor and output tensors differs. ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the input tensor and the input differential tensor differ. ▶ The strides <code>nStride</code>, <code>cStride</code>, <code>hStride</code>, <code>wStride</code> of the output tensor and the output differential tensor differ.
<code>CUDNN_STATUS_EXECUTION_FAILED</code>	The function failed to launch on the GPU.

4.83. cudnnCreateLRNDescriptor

```
cudaStatus_t cudnnCreateLRNDescriptor( cudnnLRNDescriptor_t* poolingDesc )
```

This function allocates the memory needed to hold the data needed for LRN and DivisiveNormalization layers operation and returns a descriptor used with subsequent layer forward and backward calls.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The object was created successfully.

Return Value	Meaning
CUDNN_STATUS_ALLOC_FAILED	The resources could not be allocated.

4.84. cudnnSetLRNDescriptor

```

cudnnStatus_t
CUDNNWINAPI cudnnSetLRNDescriptor( cudnnLRNDescriptor_t  normDesc,
                                   unsigned              lrnN,
                                   double                 lrnAlpha,
                                   double                 lrnBeta,
                                   double                 lrnK );

```

This function initializes a previously created LRN descriptor object.



Macros CUDNN_LRN_MIN_N, CUDNN_LRN_MAX_N, CUDNN_LRN_MIN_K, CUDNN_LRN_MIN_BETA defined in cudnn.h specify valid ranges for parameters.



Values of double parameters will be cast down to the tensor datatype during computation.

Param	In/out	Meaning
normDesc	output	Handle to a previously created LRN descriptor.
lrnN	input	Normalization window width in elements. LRN layer uses a window [center-lookBehind, center+lookAhead], where lookBehind = floor((lrnN-1)/2), lookAhead = lrnN-lookBehind-1. So for n=10, the window is [k-4...k+5] with a total of 10 samples. For DivisiveNormalization layer the window has the same extents as above in all 'spatial' dimensions (dimA[2], dimA[3], dimA[4]). By default lrnN is set to 5 in cudnnCreateLRNDescriptor.
lrnAlpha	input	Value of the alpha variance scaling parameter in the normalization formula. Inside the library code this value is divided by the window width for LRN and by (window width)^#spatialDimensions for DivisiveNormalization. By default this value is set to 1e-4 in cudnnCreateLRNDescriptor.
lrnBeta	input	Value of the beta power parameter in the normalization formula. By default this value is set to 0.75 in cudnnCreateLRNDescriptor.
lrnK	input	Value of the k parameter in normalization formula. By default this value is set to 2.0.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was set successfully.
CUDNN_STATUS_BAD_PARAM	One of the input parameters was out of valid range as described above.

4.85. cudnnGetLRNDescriptor

```

cudnnStatus_t
CUDNNWINAPI cudnnGetLRNDescriptor( cudnnLRNDescriptor_t  normDesc,
                                   unsigned               *lrnN,
                                   double                 *lrnAlpha,
                                   double                 *lrnBeta,
                                   double                 *lrnK );

```

This function retrieves values stored in the previously initialized LRN descriptor object.

Param	In/out	Meaning
normDesc	output	Handle to a previously created LRN descriptor.
lrnN, lrnAlpha, lrnBeta, lrnK	output	Pointers to receive values of parameters stored in the descriptor object. See cudnnSetLRNDescriptor for more details. Any of these pointers can be NULL (no value is returned for the corresponding parameter).

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	Function completed successfully.

4.86. cudnnDestroyLRNDescriptor

```

cudnnStatus_t cudnnDestroyLRNDescriptor(cudnnLRNDescriptor_t lrnDesc)

```

This function destroys a previously created LRN descriptor object.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The object was destroyed successfully.

4.87. cudnnLRNCrossChannelForward

```

cudnnStatus_t CUDNNWINAPI cudnnLRNCrossChannelForward(
    cudnnHandle_t          handle,
    cudnnLRNDescriptor_t  normDesc,
    cudnnLRNMode_t        lrnMode,
    const void*            alpha,
    const cudnnTensorDescriptor_t xDesc,
    const void*            *x,
    const void*            *beta,
    const cudnnTensorDescriptor_t yDesc,
    void*                  *y);

```

This function performs the forward LRN layer computation.



Supported formats are: positive-strided, NCHW for 4D x and y, and only NCDHW DHW-packed for 5D (for both x and y). Only non-overlapping 4D and 5D tensors are supported.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN library descriptor.
normDesc	input	Handle to a previously initialized LRN parameter descriptor.
lrnMode	input	LRN layer mode of operation. Currently only CUDNN_LRN_CROSS_CHANNEL_DIM1 is implemented. Normalization is performed along the tensor's dimA[1].
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the layer output value with prior value in the destination tensor as follows: $dstValue = alpha[0]*resultValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc, yDesc	input	Tensor descriptor objects for the input and output tensors.
x	input	Input tensor data pointer in device memory.
y	output	Output tensor data pointer in device memory.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The computation was performed successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the tensor pointers x, y is NULL. ▶ Number of input tensor dimensions is 2 or less. ▶ LRN descriptor parameters are outside of their valid ranges. ▶ One of tensor parameters is 5D but is not in NCDHW DHW-packed format.
CUDNN_STATUS_NOT_SUPPORTED	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ Any of the input tensor datatypes is not the same as any of the output tensor datatype. ▶ x and y tensor dimensions mismatch. ▶ Any tensor parameters strides are negative.

4.88. cudnnLRNCrossChannelBackward

```

cudnnStatus_t CUDNNWINAPI cudnnLRNCrossChannelBackward(
    cudnnHandle_t             handle,
    cudnnLRNDescriptor_t     normDesc,
    cudnnLRNMode_t          lrnMode,
    const void*              alpha,
    const cudnnTensorDescriptor_t yDesc,
    const void*              *y,
    const cudnnTensorDescriptor_t dyDesc,
    const void*              *dy,
    const cudnnTensorDescriptor_t xDesc,
    const void*              *x,
    const void*              *beta,
    const cudnnTensorDescriptor_t dxDesc,
    void*                    *dx);

```

This function performs the backward LRN layer computation.



Supported formats are: positive-strided, NCHW for 4D x and y, and only NCDHW DHW-packed for 5D (for both x and y). Only non-overlapping 4D and 5D tensors are supported.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN library descriptor.
normDesc	input	Handle to a previously initialized LRN parameter descriptor.
lrnMode	input	LRN layer mode of operation. Currently only CUDNN_LRN_CROSS_CHANNEL_DIM1 is implemented. Normalization is performed along the tensor's dimA[1].
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the layer output value with prior value in the destination tensor as follows: $dstValue = alpha[0]*resultValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
yDesc, y	input	Tensor descriptor and pointer in device memory for the layer's y data.
dyDesc, dy	input	Tensor descriptor and pointer in device memory for the layer's input cumulative loss differential data dy (including error backpropagation).
xDesc, x	input	Tensor descriptor and pointer in device memory for the layer's x data. Note that these values are not modified during backpropagation.
dxDesc, dx	output	Tensor descriptor and pointer in device memory for the layer's resulting cumulative loss differential data dx (including error backpropagation).

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The computation was performed successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met:

Return Value	Meaning
	<ul style="list-style-type: none"> ▶ One of the tensor pointers x, y is NULL. ▶ Number of input tensor dimensions is 2 or less. ▶ LRN descriptor parameters are outside of their valid ranges. ▶ One of tensor parameters is 5D but is not in NCDHW DHW-packed format.
<code>CUDNN_STATUS_NOT_SUPPORTED</code>	<p>At least one of the following conditions are met:</p> <ul style="list-style-type: none"> ▶ Any of the input tensor datatypes is not the same as any of the output tensor datatype. ▶ Any pairwise tensor dimensions mismatch for x, y, dx, dy. ▶ Any tensor parameters strides are negative.

4.89. cudnnDivisiveNormalizationForward

```

cudnnStatus_t CUDNNWINAPI cudnnDivisiveNormalizationForward(
    cudnnHandle_t          handle,
    cudnnLRNDescriptor_t   normDesc,
    cudnnDivNormMode_t     mode,
    const void             *alpha,
    const cudnnTensorDescriptor_t xDesc,
    const void             *x,
    const void             *means,
    void                  *temp,
    void                  *temp2,
    const void             *beta,
    const cudnnTensorDescriptor_t yDesc,
    void                  *y );

```

This function performs the forward spatial DivisiveNormalization layer computation. It divides every value in a layer by the standard deviation of its spatial neighbors as described in *"What is the Best Multi-Stage Architecture for Object Recognition"*, Jarrett 2009, Local Contrast Normalization Layer section. Note that Divisive Normalization only implements the $x/\max(c, \sigma_x)$ portion of the computation, where σ_x is the variance over the spatial neighborhood of x . The full LCN (Local Contrastive Normalization) computation can be implemented as a two-step process:

```

x_m = x-mean(x);
y = x_m/max(c, sigma(x_m));

```

The "x-mean(x)" which is often referred to as "subtractive normalization" portion of the computation can be implemented using cuDNN average pooling layer followed by a call to addTensor.



Supported tensor formats are NCHW for 4D and NCDHW for 5D with any non-overlapping non-negative strides. Only 4D and 5D tensors are supported.

Param	In/out	Meaning
handle	input	Handle to a previously created cuDNN library descriptor.
normDesc	input	Handle to a previously initialized LRN parameter descriptor. This descriptor is used for both LRN and DivisiveNormalization layers.
divNormMode	input	DivisiveNormalization layer mode of operation. Currently only CUDNN_DIVNORM_PRECOMPUTED_MEANS is implemented. Normalization is performed using the means input tensor that is expected to be precomputed by the user.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the layer output value with prior value in the destination tensor as follows: $dstValue = alpha[0]*resultValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc, yDesc	input	Tensor descriptor objects for the input and output tensors. Note that xDesc is shared between x, means, temp and temp2 tensors.
x	input	Input tensor data pointer in device memory.
means	input	Input means tensor data pointer in device memory. Note that this tensor can be NULL (in that case it's values are assumed to be zero during the computation). This tensor also doesn't have to contain means, these can be any values, a frequently used variation is a result of convolution with a normalized positive kernel (such as Gaussian).
temp, temp2	workspace	Temporary tensors in device memory. These are used for computing intermediate values during the forward pass. These tensors do not have to be preserved as inputs from forward to the backward pass. Both use xDesc as their descriptor.
y	output	Pointer in device memory to a tensor for the result of the forward DivisiveNormalization computation.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The computation was performed successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the tensor pointers <code>x</code>, <code>y</code>, <code>temp</code>, <code>temp2</code> is NULL. ▶ Number of input tensor or output tensor dimensions is outside of [4,5] range. ▶ A mismatch in dimensions between any two of the input or output tensors. ▶ For in-place computation when pointers <code>x == y</code>, a mismatch in strides between the input data and output data tensors. ▶ Alpha or beta pointer is NULL. ▶ LRN descriptor parameters are outside of their valid ranges. ▶ Any of the tensor strides are negative.
CUDNN_STATUS_UNSUPPORTED	At least one of the following conditions are met:

Return Value	Meaning
	► Any of the input and output tensor strides mismatch (for the same dimension).

4.90. cudnnDivisiveNormalizationBackward

```

cudnnStatus_t
CUDNNWINAPI cudnnDivisiveNormalizationBackward(
    cudnnHandle_t          handle,
    cudnnLRNDescriptor_t   normDesc,
    cudnnDivNormMode_t     mode,
    const void             *alpha,
    const cudnnTensorDescriptor_t xDesc,
    const void             *x,
    const void             *means,
    const void             *dy,
    void                   *temp,
    void                   *temp2,
    const void             *beta,
    const cudnnTensorDescriptor_t dxDesc,
    void                   *dx,
    void                   *dMeans );

```

This function performs the backward DivisiveNormalization layer computation.



Supported tensor formats are NCHW for 4D and NCDHW for 5D with any non-overlapping non-negative strides. Only 4D and 5D tensors are supported.

Param	In/ out	Meaning
handle	input	Handle to a previously created cuDNN library descriptor.
normDesc	input	Handle to a previously initialized LRN parameter descriptor (this descriptor is used for both LRN and DivisiveNormalization layers).
mode	input	DivisiveNormalization layer mode of operation. Currently only CUDNN_DIVNORM_PRECOMPUTED_MEANS is implemented. Normalization is performed using the means input tensor that is expected to be precomputed by the user.
alpha, beta	input	Pointers to scaling factors (in host memory) used to blend the layer output value with prior value in the destination tensor as follows: $dstValue = alpha[0]*resultValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc, x, means	input	Tensor descriptor and pointers in device memory for the layer's x and means data. Note: the means tensor is expected to be precomputed by the user. It can also contain any valid values (not required to be actual means, and can be for instance a result of a convolution with a Gaussian kernel).
dy	input	Tensor pointer in device memory for the layer's dy cumulative loss differential data (error backpropagation).
temp, temp2	workspace	Temporary tensors in device memory. These are used for computing intermediate values during the backward pass. These tensors do not have to be preserved from forward to backward pass. Both use xDesc as a descriptor.

Param	In/out	Meaning
dxDesc	input	Tensor descriptor for dx and dMeans.
dx, dMeans	output	Tensor pointers (in device memory) for the layer's resulting cumulative gradients dx and dMeans (dLoss/dx and dLoss/dMeans). Both share the same descriptor.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The computation was performed successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the tensor pointers <code>x</code>, <code>dx</code>, <code>temp</code>, <code>temp2</code>, <code>dy</code> is NULL. ▶ Number of any of the input or output tensor dimensions is not within the [4,5] range. ▶ Either alpha or beta pointer is NULL. ▶ A mismatch in dimensions between <code>xDesc</code> and <code>dxDesc</code>. ▶ LRN descriptor parameters are outside of their valid ranges. ▶ Any of the tensor strides is negative.
<code>CUDNN_STATUS_UNSUPPORTED</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ Any of the input and output tensor strides mismatch (for the same dimension).





4.91. cudnnBatchNormalizationForwardInference

```

cudnnStatus_t CUDNNWINAPI cudnnBatchNormalizationForwardInference(
    cudnnHandle_t          handle,
    cudnnBatchNormMode_t  mode,
    const void             *alpha,
    const void             *beta,
    const cudnnTensorDescriptor_t xDesc,
    const void             *x,
    const cudnnTensorDescriptor_t yDesc,
    void                  *y,
    const cudnnTensorDescriptor_t bnScaleBiasMeanVarDesc,
    const void             *bnScale,
    const void             *bnBias,
    const void             *estimatedMean,
    const void             *estimatedInvVariance,
    double                 epsilon );

```

This function performs the forward BatchNormalization layer computation for inference phase. This layer is based on the paper "*Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*", S. Ioffe, C. Szegedy, 2015.

-  Only 4D and 5D tensors are supported.
-  The epsilon value has to be the same during training, backpropagation and inference.
-  For training phase use `cudaBatchNormalizationForwardTraining`.
-  Much higher performance when HW-packed tensors are used for all of `x`, `dy`, `dx`.

Param	Meaning
<code>handle</code>	Input. Handle to a previously created cuDNN library descriptor.
<code>mode</code>	Input. Mode of operation (spatial or per-activation). cudaBatchNormMode_t
<code>alpha</code> , <code>beta</code>	Inputs. Pointers to scaling factors (in host memory) used to blend the layer output value with prior value in the destination tensor as follows: $dstValue = alpha[0]*resultValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
<code>xDesc</code> , <code>yDesc</code> , <code>x</code> , <code>y</code>	Tensor descriptors and pointers in device memory for the layer's <code>x</code> and <code>y</code> data.
<code>bnScaleBiasMeanVarDesc</code> , <code>bnScaleData</code> , <code>bnBiasData</code>	Inputs. Tensor descriptor and pointers in device memory for the batch normalization scale and bias parameters (in the original paper bias is referred to as <code>beta</code> and scale as <code>gamma</code>).
<code>estimatedMean</code> , <code>estimatedVariance</code>	Inputs. Mean and variance tensors (these have the same descriptor as the bias and scale). It is suggested that <code>resultRunningMean</code> , <code>resultRunningVariance</code> from the <code>cudaBatchNormalizationForwardTraining</code> call accumulated during the training phase are passed as inputs here.
<code>epsilon</code>	Input. Epsilon value used in the batch normalization formula. Minimum allowed value is <code>CUDNN_BN_MIN_EPSILON</code> defined in <code>cuda.h</code> .

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
<code>CUDNN_STATUS_SUCCESS</code>	The computation was performed successfully.
<code>CUDNN_STATUS_BAD_PARAM</code>	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ One of the pointers <code>alpha</code>, <code>beta</code>, <code>x</code>, <code>y</code>, <code>bnScaleData</code>, <code>bnBiasData</code>, <code>estimatedMean</code>, <code>estimatedInvVariance</code> is NULL.


Return Value	Meaning
	<ul style="list-style-type: none"> ▶ Number of xDesc or yDesc tensor descriptor dimensions is not within the [4,5] range. ▶ bnScaleBiasMeanVarDesc dimensions are not 1xC(x1)x1x1 for spatial or 1xC(xD)xHxW for per-activation mode (parenthesis for 5D). ▶ epsilon value is less than CUDNN_BN_MIN_EPSILON ▶ Dimensions or data types mismatch for xDesc, yDesc


4.92. cudnnBatchNormalizationForwardTraining


```


cudnnStatus_t CUDNNWINAPI cudnnBatchNormalizationForwardTraining(
    cudnnHandle_t          handle,
    cudnnBatchNormMode_t  mode,
    const void             *alpha,
    const void             *beta,
    const cudnnTensorDescriptor_t xDesc,
    const void             *x,
    const cudnnTensorDescriptor_t yDesc,
    void                  *y,
    const cudnnTensorDescriptor_t bnScaleBiasMeanVarDesc,
    const void             *bnScale,
    const void             *bnBias,
    double                 exponentialAverageFactor,
    void                  *resultRunningMean,
    void                  *resultRunningInvVariance,
    double                 epsilon,
    void                  *resultSaveMean,
    void                  *resultSaveInvVariance );
    
```

This function performs the forward BatchNormalization layer computation for training phase.

 Only 4D and 5D tensors are supported.

 The epsilon value has to be the same during training, backpropagation and inference.

 For inference phase use cudnnBatchNormalizationForwardInference.

 Much higher performance for HW-packed tensors for both x and y.

Param	Meaning
handle	Handle to a previously created cuDNN library descriptor.
mode	Mode of operation (spatial or per-activation). cudnnBatchNormMode_t

Param	Meaning
alpha, beta	Inputs. Pointers to scaling factors (in host memory) used to blend the layer output value with prior value in the destination tensor as follows: $dstValue = alpha[0]*resultValue + beta[0]*priorDstValue$. Please refer to this section for additional details.
xDesc, yDesc, x, y	Tensor descriptors and pointers in device memory for the layer's x and y data.
bnScaleBiasMeanVarDesc	Shared tensor descriptor desc for all the 6 tensors below in the argument list. The dimensions for this tensor descriptor are dependent on the normalization mode.
bnScale, bnBias	Inputs. Pointers in device memory for the batch normalization scale and bias parameters (in original paper bias is referred to as beta and scale as gamma). Note that bnBias parameter can replace the previous layer's bias parameter for improved efficiency.
exponentialAverageFactor	Input. Factor used in the moving average computation $runningMean = newMean*factor + runningMean*(1-factor)$. Use a $factor=1/(1+n)$ at N-th call to the function to get Cumulative Moving Average (CMA) behavior $CMA[n] = (x[1]+...+x[n])/n$. Since $CMA[n+1] = (n*CMA[n]+x[n+1])/(n+1) = ((n+1)*CMA[n]-CMA[n])/(n+1) + x[n+1]/(n+1) = CMA[n]*(1-1/(n+1))+x[n+1]*1/(n+1)$
resultRunningMean, resultRunningInvVariance	Inputs/outputs. Running mean and variance tensors (these have the same descriptor as the bias and scale). If these tensors are initially uninitialized, it is required that exponentialAverageFactor=1 is used for the very first call of a complete training cycle. This is necessary to properly initialize the moving average. Both of these pointers can be NULL but only at the same time. The value stored in resultRunningInvVariance (or passed as an input in inference mode) is the moving average of the expression $1 / \sqrt{eps+variance[x]}$ where variance is computed either over batch or spatial+batch dimensions depending on the mode. Note that both of these pointers can be NULL but only at the same time.
epsilon	Epsilon value used in the batch normalization formula. Minimum allowed value is CUDNN_BN_MIN_EPSILON defined in cudnn.h. Same epsilon value should be used in forward and backward functions.
resultSaveMean, resultSaveInvVariance	Outputs. Optional cache to save intermediate results computed during the forward pass - these can then be reused to speed up the backward pass. For this to work correctly, the bottom layer data has to remain unchanged until the backward function is called. Note that both of these parameters can be NULL but only at the same time. It is recommended to use this cache since memory overhead is relatively small because these tensors have a much lower product of dimensions than the data tensors.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The computation was performed successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> One of the pointers <code>alpha</code>, <code>beta</code>, <code>x</code>, <code>y</code>, <code>bnScaleData</code>, <code>bnBiasData</code> is NULL.


Return Value	Meaning
	<ul style="list-style-type: none"> ▶ Number of xDesc or yDesc tensor descriptor dimensions is not within the [4,5] range. ▶ bnScaleBiasMeanVarDesc dimensions are not 1xC(x1)x1x1 for spatial or 1xC(xD)xHxW for per-activation mode (parens for 5D). ▶ Exactly one of resultSaveMean, resultSaveInvVariance pointers is NULL. ▶ Exactly one of resultRunningMean, resultRunningInvVariance pointers is NULL. ▶ epsilon value is less than CUDNN_BN_MIN_EPSILON ▶ Dimensions or data types mismatch for xDesc, yDesc


4.93. cudnnBatchNormalizationBackward


```

cudnnStatus_t CUDNNWINAPI cudnnBatchNormalizationBackward(
    cudnnHandle_t          handle,
    cudnnBatchNormMode_t  mode,
    const void             *alphaDataDiff,
    const void             *betaDataDiff,
    const void             *alphaParamDiff,
    const void             *betaParamDiff,
    const cudnnTensorDescriptor_t xDesc,
    const void             *x,
    const cudnnTensorDescriptor_t dyDesc,
    const void             *dy,
    const cudnnTensorDescriptor_t dxDesc,
    void                  *dx,
    const cudnnTensorDescriptor_t bnScaleBiasDiffDesc,
    const void             *bnScale,
    void                  *resultBnScaleDiff,
    void                  *resultBnBiasDiff,
    double                 epsilon,
    const void             *savedMean,
    const void             *savedInvVariance
);
    
```

This function performs the backward BatchNormalization layer computation.

 Only 4D and 5D tensors are supported.

 The epsilon value has to be the same during training, backpropagation and inference.

 Much higher performance when HW-packed tensors are used for all of x, dy, dx.

Param	Meaning
handle	Handle to a previously created cuDNN library descriptor.

Param	Meaning
mode	Mode of operation (spatial or per-activation). cudnnBatchNormMode_t
alphaDataDiff, betaDataDiff	Inputs. Pointers to scaling factors (in host memory) used to blend the gradient output dx with a prior value in the destination tensor as follows: $\text{dstValue} = \alpha[0] * \text{resultValue} + \beta[0] * \text{priorDstValue}$. Please refer to this section for additional details.
alphaParamDiff, betaParamDiff	Inputs. Pointers to scaling factors (in host memory) used to blend the gradient outputs dBnScaleResult and dBnBiasResult with prior values in the destination tensor as follows: $\text{dstValue} = \alpha[0] * \text{resultValue} + \beta[0] * \text{priorDstValue}$. Please refer to this section for additional details.
xDesc, x, dyDesc, dy, dxDesc, dx	Tensor descriptors and pointers in device memory for the layer's x data, backpropagated differential dy (inputs) and resulting differential with respect to x, dx (output).
bnScaleBiasDiffDesc	Shared tensor descriptor for all the 5 tensors below in the argument list (bnScale, resultBnScaleDiff, resultBnBiasDiff, savedMean, savedInvVariance). The dimensions for this tensor descriptor are dependent on normalization mode. Note: The data type of this tensor descriptor must be 'float' for FP16 and FP32 input tensors, and 'double' for FP64 input tensors.
bnScale	Input. Pointers in device memory for the batch normalization scale parameter (in original paper bias is referred to as gamma). Note that bnBias parameter is not needed for this layer's computation.
resultBnScaleDiff, resultBnBiasDiff	Outputs. Pointers in device memory for the resulting scale and bias differentials computed by this routine. Note that scale and bias gradients are not backpropagated below this layer (since they are dead-end computation DAG nodes).
epsilon	Epsilon value used in batch normalization formula. Minimum allowed value is CUDNN_BN_MIN_EPSILON defined in cudnn.h. Same epsilon value should be used in forward and backward functions.
savedMean, savedInvVariance	Inputs. Optional cache parameters containing saved intermediate results computed during the forward pass. For this to work correctly, the layer's x and bnScale, bnBias data has to remain unchanged until the backward function is called. Note that both of these parameters can be NULL but only at the same time. It is recommended to use this cache since the memory overhead is relatively small.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The computation was performed successfully.
CUDNN_STATUS_BAD_PARAM	At least one of the following conditions are met: <ul style="list-style-type: none"> ▶ Any of the pointers <code>alpha</code>, <code>beta</code>, <code>x</code>, <code>dy</code>, <code>dx</code>, <code>bnScale</code>, <code>resultBnScaleDiff</code>, <code>resultBnBiasDiff</code> is NULL. ▶ Number of xDesc or yDesc or dxDesc tensor descriptor dimensions is not within the [4,5] range. ▶ bnScaleBiasMeanVarDesc dimensions are not 1xC(x1)x1x1 for spatial or 1xC(xD)xHxW for per-activation mode (parentheses for 5D).

Return Value	Meaning
	<ul style="list-style-type: none"> ▶ Exactly one of savedMean, savedInvVariance pointers is NULL. ▶ epsilon value is less than CUDNN_BN_MIN_EPSILON ▶ Dimensions or data types mismatch for any pair of xDesc, dyDesc, dxDesc


4.94. cudnnDeriveBNTensorDescriptor


```


cudnnStatus_t CUDNNWINAPI cudnnDeriveBNTensorDescriptor(
    cudnnTensorDescriptor_t derivedBnDesc,
    const cudnnTensorDescriptor_t xDesc,
    cudnnBatchNormMode_t mode);

```

Derives a secondary tensor descriptor for BatchNormalization scale, invVariance, bnBias, bnScale subtensors from the layer's x data descriptor. Use the tensor descriptor produced by this function as the bnScaleBiasMeanVarDesc and bnScaleBiasDiffDesc parameters in Spatial and Per-Activation Batch Normalization forward and backward functions. Resulting dimensions will be **1xC(x1)x1x1** for BATCHNORM_MODE_SPATIAL and **1xC(xD)xHxW** for BATCHNORM_MODE_PER_ACTIVATION (parentheses for 5D). For HALF input data type the resulting tensor descriptor will have a FLOAT type. For other data types it will have the same type as the input data.

 Only 4D and 5D tensors are supported.

 derivedBnDesc has to be first created using cudnnCreateTensorDescriptor

 xDesc is the descriptor for the layer's x data and has to be setup with proper dimensions prior to calling this function.

Param	In/out	Meaning
derivedBnDesc	output	Handle to a previously created tensor descriptor.
xDesc	input	Handle to a previously created and initialized layer's x data descriptor.
mode	input	Batch normalization layer mode of operation.

Possible error values returned by this function and their meanings are listed below.

Return Value	Meaning
CUDNN_STATUS_SUCCESS	The computation was performed successfully.
CUDNN_STATUS_BAD_PARAM	Invalid Batch Normalization mode.

Chapter 5.

ACKNOWLEDGMENTS

Some of the cuDNN library routines were derived from code developed by others and are subject to the following:

5.1. University of Tennessee

Copyright (c) 2010 The University of Tennessee.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer listed in this license in the documentation and/or other materials provided with the distribution.
- * Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

5.2. University of California, Berkeley

COPYRIGHT

All contributions by the University of California:
Copyright (c) 2014, The Regents of the University of California (Regents)
All rights reserved.

All other contributions:
Copyright (c) 2014, the respective contributors
All rights reserved.

Caffe uses a shared copyright model: each contributor holds copyright over their contributions to Caffe. The project versioning records all such contribution and copyright details. If a contributor wants to further mark their specific copyright on a particular contribution, they should indicate their copyright solely in the commit message of the change when it is committed.

LICENSE

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

CONTRIBUTION AGREEMENT

By contributing to the BVLC/caffe repository through pull-request, comment, or otherwise, the contributor releases their content to the license and copyright terms herein.

5.3. Facebook AI Research, New York

Copyright (c) 2014, Facebook, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name Facebook nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Additional Grant of Patent Rights

"Software" means fbcunn software distributed by Facebook, Inc.

Facebook hereby grants you a perpetual, worldwide, royalty-free, non-exclusive, irrevocable (subject to the termination provision below) license under any rights in any patent claims owned by Facebook, to make, have made, use, sell, offer to sell, import, and otherwise transfer the Software. For avoidance of doubt, no license is granted under Facebook's rights in any patent claims that are infringed by (i) modifications to the Software made by you or a third party, or (ii) the Software in combination with any software or other technology provided by you or a third party.

The license granted hereunder will terminate, automatically and without notice, for anyone that makes any claim (including by filing any lawsuit, assertion or other action) alleging (a) direct, indirect, or contributory infringement or inducement to infringe any patent: (i) by Facebook or any of its subsidiaries or affiliates, whether or not such claim is related to the Software, (ii) by any party if such claim arises in whole or in part from any software, product or service of Facebook or any of its subsidiaries or affiliates, whether or not such claim is related to the Software, or (iii) by any party relating to the Software; or (b) that any right in any patent claim of Facebook is invalid or unenforceable.

Notice

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 NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA and the NVIDIA logo are trademarks 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

© 2007-2015 NVIDIA Corporation. All rights reserved.