



**International
Standard**

ISO/IEC 15938-17

**Information technology —
Multimedia content description
interface —**

Part 17:

**Compression of neural networks for
multimedia content description and
analysis**

*Technologies de l'information — Interface de description du
contenu multimédia —*

*Partie 17: Compression des réseaux neuronaux pour la
description et l'analyse du contenu multimédia*

**Second edition
2024-01**

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-17:2024



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2024

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

Foreword	v
Introduction	vi
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 Abbreviated terms, conventions and symbols	3
4.1 General	3
4.2 Abbreviated terms	3
4.3 List of symbols	3
4.4 Number formats and computation conventions	6
4.5 Arithmetic operators	6
4.6 Logical operators	7
4.7 Relational operators	7
4.8 Bit-wise operators	7
4.9 Assignment operators	8
4.10 Range notation	8
4.11 Mathematical functions	8
4.12 Array functions	9
4.13 Order of operation precedence	11
4.14 Variables, syntax elements and tables	11
5 Overview	13
5.1 General	13
5.2 Compression tools	13
5.3 Creating encoding pipelines	14
6 Syntax and semantics	15
6.1 Specification of syntax and semantics	15
6.1.1 Method of specifying syntax in tabular form	15
6.1.2 Bit ordering	16
6.1.3 Specification of syntax functions and data types	16
6.1.4 Semantics	17
6.2 General bitstream syntax elements	18
6.2.1 NNR unit	18
6.2.2 Aggregate NNR unit	18
6.2.3 Composition of NNR bitstream	19
6.3 NNR bitstream syntax	20
6.3.1 NNR unit syntax	20
6.3.2 NNR unit size syntax	20
6.3.3 NNR unit header syntax	20
6.3.4 NNR unit payload syntax	25
6.3.5 Byte alignment syntax	31
6.4 Semantics	31
6.4.1 General	31
6.4.2 NNR unit size semantics	31
6.4.3 NNR unit header semantics	31
6.4.4 NNR unit payload semantics	39
7 Decoding process	45
7.1 General	45
7.2 NNR decompressed data formats	46
7.3 Decoding methods	47
7.3.1 General	47
7.3.2 Decoding method for NNR compressed payloads of type NNR_PT_INT	47
7.3.3 Decoding method for NNR compressed payloads of type NNR_PT_FLOAT	48

7.3.4	Decoding method for NNR compressed payloads of type NNR_PT_RAW_FLOAT	48
7.3.5	Decoding method for NNR compressed payloads of type NNR_PT_BLOCK	49
7.3.6	Decoding process for an integer weight tensor	50
8	Parameter reduction	51
8.1	General	51
8.2	Methods	51
8.2.1	Batchnorm folding	51
8.3	Syntax and semantics	52
8.3.1	Sparsification using compressibility loss	52
8.3.2	Sparsification using micro-structured pruning	52
8.3.3	Combined pruning and sparsification	52
8.3.4	Unstructured statistics-adaptive sparsification	53
8.3.5	Structured sparsification (global and local approach)	53
8.3.6	Weight unification	53
8.3.7	Low rank/low displacement rank for convolutional and fully connected layers	54
8.3.8	Batchnorm folding	54
8.3.9	Local scaling adaptation (LSA)	54
9	Parameter quantization	55
9.1	General	55
9.2	Methods	55
9.2.1	Uniform quantization method	55
9.2.2	Codebook-based method	55
9.2.3	Dependent scalar quantization method	55
9.2.4	Predictive residual encoding (PRE)	55
9.3	Syntax and semantics	55
9.3.1	Uniform quantization method	55
9.3.2	Codebook-based method	56
9.3.3	Dependent scalar quantization method	56
10	Entropy coding	56
10.1	Methods	56
10.1.1	DeepCABAC	56
10.2	Syntax and semantics	58
10.2.1	DeepCABAC syntax	58
10.3	Entropy decoding process	64
10.3.1	General	64
10.3.2	Initialization process	64
10.3.3	Binarization process	65
10.3.4	Decoding process flow	66
Annex A (normative)	Implementation for NNEF	73
Annex B (informative)	Implementation for ONNX®	75
Annex C (informative)	Implementation for PyTorch®	77
Annex D (informative)	Implementation for TensorFlow®	79
Annex E (informative)	Recommendation for carriage of NNR bitstreams in other containers	81
Annex F (informative)	Recommendation for naming method regarding performance metric type	83
Annex G (informative)	Encoding side information for selected compression tools	84
Bibliography		95

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC had received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and <https://patents.iec.ch>. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 29, *Coding of audio, picture, multimedia and hypermedia information*.

This second edition cancels and replaces the first edition (ISO/IEC 15938-17:2022), which has been technically revised.

The main changes are as follows:

- Support for incremental compression of updates of neural networks respective to a base model,
- Additional sparsification tools,
- Additional entropy coding tools, leveraging dependencies in incremental updates,
- Additional quantization tools, including representation as residuals of updates, and
- Additional high-level syntax, covering the new coding tools as well as more metadata (e.g. performance metrics).

A list of all parts in the ISO/IEC 15938 series can be found on the ISO and IEC websites.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

Introduction

Artificial neural networks have been adopted for a broad range of tasks in multimedia analysis and processing, media coding, data analytics and many other fields. Their recent success is based on the feasibility of processing much larger and complex neural networks (deep neural networks, DNNs) than in the past, and the availability of large-scale training data sets. As a consequence, trained neural networks contain a large number of parameters and weights, resulting in a quite large size (e.g. several hundred MBs). Many applications require the deployment of a particular trained network instance, potentially to a larger number of devices, which may have limitations in terms of processing power and memory (e.g. mobile devices or smart cameras), and also in terms of communication bandwidth. Any use case, in which a trained neural network (or its updates) needs to be deployed to a number of devices thus benefits from a standard for the compressed representation of neural networks.

Considering the fact that compression of neural networks is likely to have a hardware dependent and hardware independent component, this document is designed as a toolbox of compression technologies. Some of these technologies require specific representations in an exchange format (i.e. sparse representations, adaptive quantization), and thus a normative specification for representing outputs of these technologies is defined. Others do not at all materialize in a serialized representation (e.g. pruning), however, also for the latter ones required metadata is specified. This document is independent of a particular neural network exchange format, and interoperability with common formats is described in the annexes.

This document thus defines a high-level syntax that specifies required metadata elements and related semantics. In cases where the structure of binary data is to be specified (e.g. decomposed matrices) this document also specifies the actual bitstream syntax of the respective block. Annexes to the document specify the requirements and constraints of compressed neural network representations; as defined in this document; and how they are applied.

- [Annex A](#) specifies the implementation of this document with the Neural Network Exchange Format (NNEF¹⁾), defining the use of NNEF to represent network topologies in a compressed neural network bitstream.
- [Annex B](#) provides recommendations for the implementation of this document with the Open Neural Network Exchange Format (ONNX²⁾), defining the use of ONNX to represent network topologies in a compressed neural network bitstream.
- [Annex C](#) provides recommendations for the implementation of this document with the PyTorch³⁾ format, defining the reference to PyTorch elements in the network topology description of a compressed neural network bitstream.
- [Annex D](#) provides recommendations for the implementation of this document with the Tensorflow⁴⁾ format, defining the reference to Tensorflow elements in the network topology description of a compressed neural network bitstream.
- [Annex E](#) provides recommendations for the carriage of tensors compressed according to this document in third party container formats.
- [Annex F](#) provides recommendations for the naming of common performance metrics to specify the metric that was used for validation.

1) NNEF is the trademark of a product owned by The Khronos® Group. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO/IEC of the product named.

2) ONNX is the trademark of a product owned by LF PROJECTS, LLC. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO/IEC of the product named.

3) PyTorch is the trademark of a product supplied by Facebook, Inc. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO/IEC of the product named.

4) TensorFlow is the trademark of a product supplied by Google LLC. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO/IEC of the product named.

- [Annex G](#) provides recommendations for implementing the encoding side of some of the compression tools.

The compression tools described in this document have been selected and evaluated for neural networks used in applications for multimedia description, analysis and processing. However, they may be useful for the compression of neural networks used in other applications and applied to other types of data.

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-17:2024

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-17:2024

Information technology — Multimedia content description interface —

Part 17:

Compression of neural networks for multimedia content description and analysis

1 Scope

This document specifies Neural Network Coding (NNC) as a compressed representation of the parameters/weights of a trained neural network and a decoding process for the compressed representation, complementing the description of the network topology in existing (exchange) formats for neural networks. It establishes a toolbox of compression methods, specifying (where applicable) the resulting elements of the compressed bitstream. Most of these tools can be applied to the compression of entire neural networks, and some of them can also be applied to the compression of differential updates of neural networks with respect to a base network. Such differential updates are for example useful when models are redistributed after fine-tuning or transfer learning, or when providing versions of a neural network with different compression ratios.

This document does not specify a complete protocol for the transmission of neural networks, but focuses on compression of network parameters. Only the syntax format, semantics, associated decoding process requirements, parameter sparsification, parameter transformation methods, parameter quantization, entropy coding method and integration/signalling within existing exchange formats are specified, while other matters such as pre-processing, system signalling and multiplexing, data loss recovery and post-processing are considered to be outside the scope of this document. Additionally, the internal processing steps performed within a decoder are also considered to be outside the scope of this document; only the externally observable output behaviour is required to conform to the specifications of this document.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 10646, *Information technology — Universal coded character set (UCS)*

ISO/IEC 60559, *Information technology — Microprocessor Systems — Floating-Point arithmetic*

IETF RFC 1950, *ZLIB Compressed Data Format Specification version 3.3*

NNEF-v1.0.3⁵⁾, *Neural Network Exchange Format*, The Khronos NNEF Working Group, Version 1.0.3, 2020

FIPS PUB 180-4:2015, *Secure Hash Standard (SHS)*

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

5) Available from: <https://www.khronos.org/registry/NNEF/specs/1.0/nnef-1.0.3.pdf>

ISO and IEC maintain terminology databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <https://www.iso.org/obp>
- IEC Electropedia: available at <https://www.electropedia.org/>

3.1

aggregate NNR unit

NNR unit which carries multiple NNR units in its payload

3.2

base neural network

neural network serving as reference for a differential update

3.3

compressed neural network representation

NNR

representation of a neural network with model parameters encoded using compression tools

3.4

decomposition

transformation to express a tensor as product of two tensors

3.5

hyperparameter

parameter whose value is used to control the learning process

3.6

layer

collection of nodes operating together at a specific depth within a neural network

3.7

model parameter

coefficients of the neural network model such as weights and biases

3.8

NNR unit

data structure for carrying (compressed or uncompressed) neural network data and related metadata

3.9

parameter identifier

value that uniquely identifies a parameter throughout different incremental updates

Note 1 to entry: Parameters having the same parameter identifier are at the same position in the same tensor in different incremental updates. This means they are co-located.

3.10

pruning

reduction of parameters in (a part of) the neural network

3.11

sparsification

increase of the number of zero-valued entries of a tensor

3.12

tensor

multidimensional structure grouping related model parameters

3.13**updated neural network**

neural network resulting from modifying the base neural network

Note 1 to entry: The updated neural network is reconstructed by applying the differential update to the base neural network.

4 Abbreviated terms, conventions and symbols**4.1 General**

This subclause contains the definition of operators, notations, functions, textual conventions and processes used throughout this document.

The mathematical operators used in this document are similar to those used in the C programming language. However, the results of integer division and arithmetic shift operations are specified more precisely, and additional operations are specified, such as exponentiation and real-valued division. Numbering and counting conventions generally begin from 0, e.g. "the first" is equivalent to the 0-th, "the second" is equivalent to the 1-th, etc.

4.2 Abbreviated terms

DeepCABAC	Context-adaptive binary arithmetic coding for deep neural networks
LDR	Low displacement rank
LPS	Layer parameter set
LR	Low-rank
LSA	Local scaling adaptation
LSB	Least significant bit
MPS	Model parameter set
MSB	Most significant bit
MSE	Mean square error
NN	Neural network
NNC	Neural network coding
NDU	NNR compressed data unit
NNEF	Neural network exchange format
QP	Quantization parameter
PRE	Predictive residual encoding
SBT	Stochastic binary-ternary quantization
SVD	Singular value decomposition

4.3 List of symbols

This document defines the following symbols:

A	Input tensor
B	Output tensor
B_{jl}^k	Block in superblock j of layer k .
b	Bias parameter
C_i	Number of input channels of a convolutional layer
C_o	Number of output channels of a convolutional layer
c_j^k	Number of channels in dimension j of tensor in layer k
$c_j^{k'}$	Derived number of channels in dimension j of tensor in layer k
d_j^k	Depth dimension of tensor at layer k
e	Parameter of f-circulant matrix Z_e
F	Parameter tensor of a convolutional layer
f	Parameter of f-circulant matrix Z_f
G_k	Left-hand side matrix of Low Rank decomposed representation of matrix W_k
H_k	Right-hand side matrix of Low Rank decomposed representation of matrix W_k
h_j^k	Height dimension of tensor for layer k
K	Dimension of a convolutional kernel
L	Loss function
L_c	Compressibility loss
L_d	Diversity loss
L_s	Task loss
L_t	Training loss
M	Feature matrix
M_k	Pruning mask for layer k
m	Sparsification hyperparameter
m_i	i -th row of feature matrix M
n_j^k	Kernel size of tensor at layer k .
n^k	Dimension resulting from a product over n_j^k
P	Stochastic transition matrix
p	Pruning ratio hyperparameter
p_{ij}	Elements of transition matrix P
q	Sparsification ratio hyperparameter

q_b	Binary quantization
q_t	Ternary quantization
S	Importance of parameters for pruning
S_j^k	Superblock j in layer k
s	Local scaling factors
s_j^k	Size of superblock j in layer k
T	Topology element
T^q	Quantizable topology element
u	Unification ratio hyperparameter
W	Parameter tensor
ΔW	Difference of parameter tensor
W_l	Weight tensor of l -th layer
W_k	Parameter tensor of layer k
$\hat{W}_k$	Low Rank approximation of W_k
w	Parameter vector
$w_{l,i}$	Vector of weights for the i -th filter in the l -th layer
$w'_{l,i}$	Vector of normalized weights for the i -th filter in the l -th layer
y, y_{ref}	Coding performance, reference coding performance
y_d	Coding performance difference
X	Input to a batch-normalization layer
Z_e	f -circulant matrix
Z_f	f -circulant matrix
α	Folded batch normalization parameter
α'	Combined value for folded batch normalization parameter and local scaling factors
β	Batch normalization parameter
β_u	Updated batch normalization parameter
γ_c	Compressibility loss multiplier
γ	Batch normalization parameter
γ_u	Updated batch normalization parameter
δ	Folded batch normalization parameter
δ_f	Sparsification threshold (mean of filter means)

δ_s	Scaling factor for sparsification
ϵ	Scalar close to zero to avoid division by zero in batch normalization
λ	Eigenvector
λ_c	Compressibility loss weight
λ_d	Diversity loss weight
μ	Batch normalization parameter
v_j^k	Width dimension of tensor for layer k .
π	Equilibrium probability of P
π_t	Probability of applying ternary quantization
ρ	Parameter
σ	Batch normalization parameter
τ	Threshold (sparsification, ternary-binary quantization)
θ_ρ	Weight magnitude threshold
φ	Smoothing factor

4.4 Number formats and computation conventions

This document defines the following number formats:

integer	Integer number which may be arbitrarily small or large. Integers are also referred to as signed integers.
unsigned integer	Unsigned integer that may be zero or arbitrarily large.
float	Floating point number according to ISO/IEC 60559.

If not specified otherwise, outcomes of all operators and mathematical functions are mathematically exact. Whenever an outcome shall be a float, it is explicitly specified.

4.5 Arithmetic operators

The following arithmetic operators are defined:

+	Addition
−	Subtraction (as a two-argument operator) or negation (as a unary prefix operator)
*	Multiplication, including matrix multiplication
◦	Element-wise multiplication of two transposed vectors or element-wise multiplication of a transposed vector with rows of a matrix or Hadamard product of two matrices with identical dimensions
x^y	Exponentiation. Specifies x to the power of y . In other contexts, such notation is used for superscripting not intended for interpretation as exponentiation.

/	Integer division with truncation of the result toward zero. For example, $7 / 4$ and $-7 / -4$ are truncated to 1 and $-7 / 4$ and $7 / -4$ are truncated to -1.
$\div$	Used to denote division in mathematical equations where no truncation or rounding is intended.
$\frac{x}{y}$	Used to denote division in mathematical equations where no truncation or rounding is intended, including element-wise division of two transposed vectors or element-wise division of a transposed vector with rows of a matrix.
$\sum_{i=x}^y f(i)$	The summation of $f(i)$ with i taking all integer values from x up to and including y .
$\prod_{i=x}^y f(i)$	The product of $f(i)$ with i taking all integer values from x up to and including y .
$x \% y$	Modulus. Remainder of x divided by y , defined only for integers x and y with $x \geq 0$ and $y > 0$.

4.6 Logical operators

The following logical operators are defined:

$x \&\& y$	Boolean logical "and" of x and y
$x y$	Boolean logical "or" of x and y
!	Boolean logical "not"
$x ? y : z$	If x is TRUE or not equal to 0, evaluates to the value of y ; otherwise, evaluates to the value of z .

4.7 Relational operators

The following relational operators are defined as follows:

>	Greater than
$\geq$	Greater than or equal to
<	Less than
$\leq$	Less than or equal to
==	Equal to
!=	Not equal to

When a relational operator is applied to a syntax element or variable that has been assigned the value "na" (not applicable), the value "na" is treated as a distinct value for the syntax element or variable. The value "na" is considered not to be equal to any other value.

4.8 Bit-wise operators

The following bit-wise operators are defined as follows:

&	Bit-wise "and". When operating on integer arguments, operates on a two's complement representation of the integer value. When operating on a binary argument that contains fewer bits than another argument, the shorter argument is extended by adding more significant bits equal to 0.
---	---

	Bit-wise "or". When operating on integer arguments, operates on a two's complement representation of the integer value. When operating on a binary argument that contains fewer bits than another argument, the shorter argument is extended by adding more significant bits equal to 0.
^	Bit-wise "exclusive or". When operating on integer arguments, operates on a two's complement representation of the integer value. When operating on a binary argument that contains fewer bits than another argument, the shorter argument is extended by adding more significant bits equal to 0.
$x \gg y$	Arithmetic right shift of a two's complement integer representation of x by y binary digits. This function is defined only for non-negative integer values of y . Bits shifted into the MSBs as a result of the right shift have a value equal to the MSB of x prior to the shift operation.
$x \ll y$	Arithmetic left shift of a two's complement integer representation of x by y binary digits. This function is defined only for non-negative integer values of y . Bits shifted into the LSBs as a result of the left shift have a value equal to 0.
!	Bit-wise not operator returning 1 if applied to 0 and 0 if applied to 1.

4.9 Assignment operators

The following arithmetic operators are defined as follows:

=	Assignment operator
++	Increment, i.e. $x++$ is equivalent to $x = x + 1$; when used in an array index, evaluates to the value of the variable prior to the increment operation.
--	Decrement, i.e. $x--$ is equivalent to $x = x - 1$; when used in an array index, evaluates to the value of the variable prior to the decrement operation.
+=	Increment by amount specified, i.e. $x += 3$ is equivalent to $x = x + 3$, and $x += (-3)$ is equivalent to $x = x + (-3)$.
-=	Decrement by amount specified, i.e. $x -= 3$ is equivalent to $x = x - 3$, and $x -= (-3)$ is equivalent to $x = x - (-3)$.

4.10 Range notation

The following notation is used to specify a range of values:

$x = y..z$	x takes on integer values starting from y to z , inclusive, with x , y , and z being integer numbers and z being greater than y .
$\text{array}[x..y]$	a sub-array containing the elements of array comprised between position x and y included. If x is greater than y , the resulting sub-array is empty.

4.11 Mathematical functions

The following mathematical functions are defined:

$\text{Ceil}(x)$	the smallest integer greater than or equal to x
$\text{Floor}(x)$	the largest integer less than or equal to x
$\text{Log2}(x)$	the base-2 logarithm of x

$$\text{Min}(x, y) = \begin{cases} x & ; x \leq y \\ y & ; x > y \end{cases}$$

$$\text{Max}(x, y) = \begin{cases} x & ; x \geq y \\ y & ; x < y \end{cases}$$

4.12 Array functions

`Size( arrayName[] )` returns the number of elements contained in the array or tensor named `arrayName`. If `arrayName[]` is a tensor this corresponds to the product of all dimensions of the tensor.

`Prod( arrayName[] )` returns the product of all elements of array `arrayName[]`.

`TensorReshape( arrayName[], tensorDimension[])` returns the reshaped tensor `array_name[]` with the specified `tensorDimension[]`, without changing its data.

`IndexToXY(w, h, i, bs)` returns an array with two elements. The first element is an x coordinate and the second element is a y coordinate pointing into a 2D array of width w and height h . x and y point to the position that corresponds to scan index i when the block is scanned in blocks of size bs times bs . x and y are derived as follows:

A variable `fullRowOfBlocks` is set to $w * bs$

A variable `blockY` is set to $i / \text{fullRowOfBlocks}$

A variable `iOff` is set to $i \% \text{fullRowOfBlocks}$

A variable `currBlockH` is set to $\text{Min}(bs, h - \text{blockY} * bs)$

A variable `fullBlocks` is set to $bs * \text{currBlockH}$

A variable `blockX` is set to $iOff / \text{fullBlocks}$

A variable `blockOff` is set to $iOff \% \text{fullBlocks}$

A variable `currBlockW` is set to $\text{Min}(bs, w - \text{blockX} * bs)$

A variable `posX` is set to $\text{blockOff} \% \text{currBlockW}$

A variable `posY` is set to $\text{blockOff} / \text{currBlockW}$

The variable x is set to $\text{blockX} * bs + \text{posX}$

The variable y is set to $\text{blockY} * bs + \text{posY}$

`TensorIndex( tensorDimensions[], i, scan )` returns an array with the same number of dimensions as `tensorDimensions[]` where the elements of the array are set to integer values so that the array can be used as an index pointing to an element of a tensor with dimensions `tensorDimensions[]` as follows:

If variable `scan` is equal to 0:

The returned array points to the i -th element in row-major scan order of a tensor with dimensions `tensorDimensions[]`.

If variable `scan` is greater than 0:

A variable `bs` is set to $4 \ll \text{scan_order}$.

A variable `h` is set to `tensorDimensions[0]`.

A variable `w` is set to $\text{Prod}(\text{tensorDimensions}) / h$.

Two variables x and y are set to the first and second element of the array that is returned by calling `IndexToXY(w, h, i, bs)`, respectively.

The returned array is `TensorIndex(tensorDimensions, y * w + x, 0)`.

NOTE Variable *scan* usually corresponds to syntax element *scan_order*.

`GetEntryPointIdx( tensorDimensions[], i, scan )` returns -1 if index *i* doesn't point to the first position of an entry point. If index *i* points to the first position of an entry point, it returns the entry point index within the tensor. To determine the positions and indexes of entry points, the following applies:

A variable *w* is set to $\text{Prod}(\text{tensorDimensions}) / \text{tensorDimensions}[0]$.

A variable *epIdx* is set to $i / (w * (4 \ll \text{scan})) - 1$.

If $i > 0$ and $i \% (w * (4 \ll \text{scan}))$ is equal to 0, index *i* points to the first position of an entry point and the entry point index is equal to *epIdx*.

Otherwise, index *i* doesn't point to the first position of an entry point.

`ShiftArrayIndex( inputArray[], shiftIndexPosition )` returns an array *outputArray[]* which is a copy of *inputArray[]* but with the element at position 0 of the *inputArray* shifted to *shiftIndexPosition* as follows:

A variable *outputArray[]* is initialized with a copy of *inputArray[]*.

If *shiftIndexPosition* is greater than 0:

The first element of *outputArray[]* is erased from *outputArray[]*.

The first element of *inputArray[]* is inserted into *outputArray[]* before the element with position *shiftIndexPosition* and after element with position *shiftIndexPosition*-1.

`DimensionShift( inputTensor[], tensorDimensions[], firstDimensionShift )` returns a tensor *reorderedTensor[]* which is a copy of *inputTensor[]* with the same number of dimensions, but where the dimensions are rearranged such that the first dimension of *inputTensor[]* specified by *tensorDimensions[]* is shifted to position *firstDimensionsShift* as follows:

A variable *reorderedTensor* is initialized with dimensions equal to `ShiftArrayIndex( tensorDimensions, firstDimensionShift )`.

The elements of variable *reorderedTensor* are set as follows:

```
for( i = 0; i < Prod( tensorDimensions ); i++ ){
    idxA = TensorIndex( tensorDimensions, i, 0 )
    idxB = ShiftArrayIndex( idxA, firstDimensionShift )
    reorderedTensor[idxB] = inputTensor[idxA]
}
```

`AxisSwap( inputTensor[], tensorDimensions[], numberOfDimensions, axis0, axis1 )` returns a tensor which is derived from *inputTensor* (with dimensions *tensorDimensions* and number of dimensions as *numberOfDimensions*) and where values in the axis indexes *axis0* and *axis1* of the *inputTensor* are swapped.

`TensorSplit( inputTensor[], splitIndices, splitAxis )` returns an array of tensors *subTensors* that is derived by splitting tensor *inputTensor* into $N = \text{Size}(\text{splitIndices}) + 1$ tensors using the provided array of indices *splitIndices* along the provided axis *splitAxis* as follows:

An array *inputDims* is set to the dimensions of tensor *inputTensor*.

An element with value 0 is inserted into *splitIndices* before the first element and an element with value *inputDims[splitAxis]* is inserted into *splitIndices* after the last element.

Tensor *subTensors[x]* (with *x* being an integer from 0 to *N*) is derived as follows:

An array *subTensorDims* is set to *inputDims*.

Element *subTensorDims*[*splitAxis*] is replaced with value *splitIndices*[*x* + 1] – *splitIndices*[*x*].

The elements of *subTensors*[*x*] are set as follows:

```
for( i = 0; i < Prod( subTensorDims ); i++ ) {
    subIdx = TensorIndex( subTensorDims, i, 0 )
    inputIdx = TensorIndex( inputDims, i, 0 )
    inputIdx[splitAxis] += splitIndices[x]
    subTensors[subIdx] = inputTensor[inputIdx]
```

4.13 Order of operation precedence

When the order of precedence in an expression is not indicated explicitly by use of parentheses, the following rules apply:

- Operations of a higher precedence are evaluated before any operation of a lower precedence.
- Operations of the same precedence are evaluated sequentially from left to right.

[Table 1](#) specifies the precedence of operations from highest to lowest; a higher position in the table indicates a higher precedence.

NOTE For those operators that are also used in the C programming language, the order of precedence used in this document is the same as used in the C programming language.

Table 1 — Operation precedence from highest (at top of table) to lowest (at bottom of table)

operations (with operands <i>x</i> , <i>y</i> , and <i>z</i>)
" <i>x</i> ++", " <i>x</i> --"
"! <i>x</i> ", " <i>-x</i> " (as a unary prefix operator)
" <i>x</i> ^{<i>y</i>} "
" <i>x</i> * <i>y</i> ", " <i>x</i> / <i>y</i> ", " <i>x</i> ÷ <i>y</i> ", " $\frac{x}{y}$ ", " <i>x</i> % <i>y</i> ", " $\prod_{i=x}^y f(i)$ ", " <i>x</i> ◦ <i>y</i> "
" <i>x</i> + <i>y</i> ", " <i>x</i> – <i>y</i> " (as a two-argument operator), " $\sum_{i=x}^y f(i)$ "
" <i>x</i> << <i>y</i> ", " <i>x</i> >> <i>y</i> "
" <i>x</i> < <i>y</i> ", " <i>x</i> ≤ <i>y</i> ", " <i>x</i> > <i>y</i> ", " <i>x</i> ≥ <i>y</i> "
" <i>x</i> == <i>y</i> ", " <i>x</i> != <i>y</i> "
" <i>x</i> & <i>y</i> "
" <i>x</i> <i>y</i> "
" <i>x</i> && <i>y</i> "
" <i>x</i> <i>y</i> "
" <i>x</i> ? <i>y</i> : <i>z</i> "
" <i>x</i> . <i>y</i> "
" <i>x</i> = <i>y</i> ", " <i>x</i> += <i>y</i> ", " <i>x</i> -= <i>y</i> "

4.14 Variables, syntax elements and tables

Syntax elements in the bitstream are represented in **bold** type. Each syntax element is described by its name (all lower-case letters with underscore characters), and one data type for its method of coded representation. The decoding process behaves according to the value of the syntax element and to the values of previously decoded syntax elements. When a value of a syntax element is used in the syntax tables or the text, it appears in regular (i.e. not bold) type.

In some cases the syntax tables may use the values of other variables derived from syntax elements values. Such variables appear in the syntax tables, or text, named by a mixture of lower-case and upper-case letter and without any underscore characters (camel case notation). Variables starting with an upper-case letter are derived for the decoding of the current syntax structure and all depending syntax structures. Variables starting with an upper-case letter may be used in the decoding process for later syntax structures without mentioning the originating syntax structure of the variable. Variables starting with a lower-case letter are only used within the (sub)clause in which they are derived.

In some cases, "mnemonic" names for syntax element values or variable values are used interchangeably with their numerical values. Sometimes "mnemonic" names are used without any associated numerical values. The association of values and names is specified in the text. The names are constructed from one or more groups of letters separated by an underscore character. Each group starts with an upper-case letter and may contain more upper-case letters.

NOTE The syntax is described in a manner that closely follows the C-language syntactic constructs.

Functions that specify properties of the current position in the bitstream are referred to as syntax functions. These functions are specified in [subclause 6.3](#) and assume the existence of a bitstream pointer with an indication of the position of the next bit to be read by the decoding process from the bitstream. Syntax functions are described by their names, which are constructed as syntax element names and end with left and right round parentheses including zero or more variable names (for definition) or values (for usage), separated by commas (if more than one variable).

Functions that are not syntax functions (including mathematical functions specified in [subclause 4.11](#) and array functions specified in [subclause 4.12](#)) are described by their names, which start with an upper-case letter, contain a mixture of lower and upper-case letters without any underscore character, and end with left and right parentheses including zero or more variable names (for definition) or values (for usage) separated by commas (if more than one variable).

A one-dimensional array is referred to as a list. A two-dimensional array is referred to as a matrix. Arrays can either be syntax elements or variables. Subscripts or square parentheses are used for the indexing of arrays. In reference to a visual depiction of a matrix, the first subscript is used as a row (vertical) index and the second subscript is used as a column (horizontal) index. The indexing order is reversed when using square parentheses rather than subscripts for indexing. Thus, an element of a matrix s at horizontal position x and vertical position y may be denoted either as $s[x][y]$ or as s_{yx} . A single column of a matrix may be referred to as a list and denoted by omission of the row index. Thus, the column of a matrix s at horizontal position x may be referred to as the list $s[x]$.

A multi-dimensional array is a variable with a number of dimensions. An element of the multi-dimensional array is either indexed by specifying all required indexes like e.g. `variable[x][y][z]` or by a single index variable that itself is a one-dimensional array specifying the indexes. For example `variable[i]` with i being a one-dimensional array with elements $[x, y, z]$. Multi-dimensional arrays are, for example, used to specify tensors.

A specification of values of the entries in rows and columns of an array may be denoted by $\{ \{ \dots \} \{ \dots \} \}$, where each inner pair of brackets specifies the values of the elements within a row in increasing column order and the rows are ordered in increasing row order. Thus, setting a matrix s equal to $\{ \{ 1 \ 6 \} \{ 4 \ 9 \} \}$ specifies that $s[0][0]$ is set equal to 1, $s[1][0]$ is set equal to 6, $s[0][1]$ is set equal to 4, and $s[1][1]$ is set equal to 9.

Binary notation is indicated by enclosing the string of bit values by single quote marks. For example, '01000001' represents an eight-bit string having only its second and its last bits (counted from the most to the least significant bit) equal to 1.

Hexadecimal notation, indicated by prefixing the hexadecimal number by "0x", may be used instead of binary notation when the number of bits is an integer multiple of 4. For example, 0x41 represents an eight-bit string having only its second and its last bits (counted from the most to the least significant bit) equal to 1.

Numerical values not enclosed in single quotes and not prefixed by "0x" are decimal values.

A value equal to 0 represents a FALSE condition in a test statement. The value TRUE is represented by any value different from zero.

5 Overview

5.1 General

This clause provides an overview of the compression tools defined in this document and describes how they can be combined to encoding pipelines.

5.2 Compression tools

This document contains the following groups of compression tools.

Parameter reduction methods process a model to obtain a compact representation. Examples of such methods include, *parameter sparsification*, *parameter pruning*, *weight unification*, and *decomposition methods*.

Sparsification processes parameters or groups of parameters to produce a sparse representation of the model (e.g. by replacing some weight values with zeros). The sparsification can generate additional metadata (e.g. masks). The sparsification can be structured or unstructured. This document includes methods for unstructured sparsification with compressibility loss (information about encoding provided in [subclause G.1.2](#)), structured sparsification using micro-structured sparsification (information about encoding provided in [subclause G.1.3](#)), unstructured statistics-adaptive sparsification (information about encoding provided in [subclause G.1.5](#)), and structured sparsification (information about encoding provided in [subclause G.1.6](#)).

Unification processes the parameters to produce group of similar parameters. Unification does not eliminate or constrain the weights to be zero, but it lowers the entropy of model parameters by making them similar to each other. This document includes a method for parameter unification (information about encoding provided in [subclause G.1.7](#)).

Pruning reduces the number of parameters by eliminating parameters or groups of parameters. The procedure results in a dense representation which has less parameters in comparison to the original model, e.g. by removing some redundant convolution filters from the layers. This document includes a a method for combined pruning and sparsification (information about encoding provided in [subclause G.1.4](#))

Decomposition performs a matrix decomposition operation to change the structure of the weights of a model. This document includes a method for low rank/low displacement rank for convolutional and fully connected layers (information about encoding provided in [subclause G.1.8](#)).

Along with the reduction methods mentioned above, this document includes decomposition methods that are introduced and tested as part of a parameter quantization technique. Examples of such methods are batchnorm folding ([subclause 8.2.1](#) and information about encoding provided in [subclause G.1.9](#)) and local scaling adaptation (information about encoding provided in [subclause G.1.10](#)).

The parameter reduction methods can be combined or applied in sequence to produce a compact model.

Parameter quantization methods reduce the precision of the representation of parameters. If supported by the inference engine, the quantized representation can be used for more efficient inference. This document includes methods for uniform quantization ([subclause 9.2.1](#)), codebook-based quantization ([subclause 9.2.2](#)), dependent scalar quantization ([subclause 9.2.3](#)), iterative QP optimization (information about encoding provided in [subclause G.2.2](#)) and stochastic binary-ternary quantization (information about encoding provided in [subclause G.2.3](#)).

Predictive residual encoding (PRE, [subclause 9.2.4](#)) enables to code residuals based on a previously decoded update of the model rather than the complete weight update.

Entropy coding methods encode the results of parameter quantization methods. This document includes DeepCABAC ([subclause 10.1.1](#)) as entropy encoding method. Supported extensions for DeepCABAC include Row Skipping and Temporal Context Modeling.

Row Skipping reduces the number of bins to be decoded and also the bitstream size by skipping decoding of matrix rows that are entirely zero. The method is described in [subclause 10.1.1.2](#). *Temporal Context*

Modeling uses information from previously decoded incremental updates to improve the context modeling of DeepCABAC and thus increases the coding efficiency. The method is described in [subclause 10.1.1.4](#).

5.3 Creating encoding pipelines

The compression tools in this document can be combined to form different encoding pipelines. Some of the tools are alternatives for addressing neural network models with different types of characteristics, while other tools are designed to work in sequence.

[Figure 1](#) shows an overview of encoding pipelines that can be assembled using the compression tools in this document. From the group of parameter transformation tools, multiple tools can be applied in sequence. Parameter quantization can be applied to source models as well as to the outputs of transformation with parameter reduction methods. Entropy coding is usually applied to the output of quantization. Raw outputs of earlier steps without applying entropy coding can be serialized if needed.




Figure 1 — NNR encoding pipelines

The following encoding pipelines are considered typical examples of using this document:

1. Dependent scalar quantization ([subclause 9.2.3](#)) – DeepCABAC ([subclause 10.1.1](#))
2. Sparsification (information about encoding provided in [subclause G.1.2](#)) – Dependent scalar quantization ([subclause 9.2.3](#)) – DeepCABAC ([subclause 10.1.1](#))
3. Low-rank decomposition (information about encoding provided in [subclause G.1.8](#)) – Dependent scalar quantization ([subclause 9.2.3](#)) – DeepCABAC ([subclause 10.1.1](#))
4. Codebook-based quantization ([subclause 9.2.2](#)) – DeepCABAC ([subclause 10.1.1](#))
5. Unification (information about encoding provided in [subclause G.1.7](#)) – DeepCABAC ([subclause 10.1.1](#))
6. Stochastic binary-ternary quantization (information about encoding provided in [subclause G.2.3](#)) – Predictive residual encoding ([subclause 9.2.4](#)) – DeepCABAC ([subclause 10.1.1](#))

This list is non-exhaustive.

The following coding tools are only applicable to updates of neural network parameters:

- Predictive residual encoding ([subclause 9.2.4](#));

— Temporal Context Modeling ([subclause 10.1.1.4](#)).

6 Syntax and semantics

6.1 Specification of syntax and semantics

6.1.1 Method of specifying syntax in tabular form

The syntax tables specify a superset of the syntax of all allowed bitstreams. Additional constraints on the syntax may be specified, either directly or indirectly, in other clauses.

[Table 2](#) lists examples of the syntax specification format. When **syntax_element** appears, it specifies that a syntax element is parsed from the bitstream and the bitstream pointer is advanced to the next position beyond the syntax element in the bitstream parsing process.

Table 2 — Examples of the syntax specification format

Syntax	Type/Clause
/* A statement can be a syntax element with an associated data type or can be an expression used to specify conditions for the existence, type and quantity of syntax elements, as in the following two examples */	
syntax_element	st(v)
conditioning statement	
/*A group of statements enclosed in curly brackets is a compound statement and is treated functionally as a single statement. */	
{	
statement	
statement	
...	
}	
/* A "while" structure specifies a test of whether a condition is true, and if true, specifies evaluation of a statement (or compound statement) repeatedly until the condition is no longer true */	
while(condition)	
statement	
/* A "do ... while" structure specifies evaluation of a statement once, followed by a test of whether a condition is true, and if true, specifies repeated evaluation of the statement until the condition is no longer true */	
do	
statement	
while(condition)	
/* An "if ... else" structure specifies a test of whether a condition is true and, if the condition is true, specifies evaluation of a primary statement, otherwise, specifies evaluation of an alternative statement. The "else" part of the structure and the associated alternative statement is omitted if no alternative statement evaluation is needed */	
if(condition)	
primary statement	
else	
alternative statement	

Table 2 (continued)

Syntax	Type/Clause
/* A "for" structure specifies evaluation of an initial statement, followed by a test of a condition, and if the condition is true, specifies repeated evaluation of a primary statement followed by a subsequent statement until the condition is no longer true. */	
for(initial statement; condition; subsequent statement)	
primary statement	

6.1.2 Bit ordering

For bit-oriented delivery, the bit order of syntax fields in the syntax tables is specified to start with the MSB and proceed to the LSB.

6.1.3 Specification of syntax functions and data types

The functions presented here are used in the syntactical description. These functions are expressed in terms of the value of a bitstream pointer that indicates the position of the next bit to be read by the decoding process from the bitstream.

byte_aligned() is specified as follows:

- If the current position in the bitstream is on a byte boundary, i.e. the next bit in the bitstream is the first bit in a byte, the return value of byte_aligned() is equal to TRUE.
- Otherwise, the return value of byte_aligned() is equal to FALSE.

read_bits(n) reads the next n bits from the bitstream and advances the bitstream pointer by n bit positions. When n is equal to 0, read_bits(n) is specified to return a value equal to 0 and to not advance the bitstream pointer.

get_bit_pointer() returns the position of the bitstream pointer relative to the beginning of the current NNR unit as unsigned integer value. get_bit_pointer() >> 3 points to the current byte of the bitstream pointer. get_bit_pointer() & 7 points to the current bit in the current byte of the bitstream pointer where a value of 0 indicates the most significant bit.

set_bit_pointer(pos) sets the position of the bitstream pointer such that get_bit_pointer() equals pos.

The following data types specify the parsing process of each syntax element:

- ae(v): context-adaptive arithmetic entropy-coded syntax element. The parsing process for this data type is specified in [subclause 10.3.4.3.2](#).
- at(v) : arithmetic entropy-coded termination syntax. The parsing process for this data type is specified in [subclause 10.3.4.3.5](#).
- iae(n): signed integer using n arithmetic entropy-coded bits using the bypass mode of DeepCABAC as specified in [subclause 10.3.4.3.4](#). The read bypass bins are interpreted as a two's complement integer representation with most significant bit written first.
- uae(n): unsigned integer using n arithmetic entropy-coded bits using the bypass mode of DeepCABAC as specified in [subclause 10.3.4.3.4](#). The read bypass bins are interpreted as a binary representation of an unsigned integer with most significant bit written first. When n=0, uae(n) does not decode any bins and returns 0.
- f(n): fixed-pattern bit string using n bits written (from left to right) with the left bit first. The parsing process for this data type is specified by the return value of the function read_bits(n).
- i(n): signed integer using n bits. When n is "v" in the syntax table, the number of bits varies in a manner dependent on the value of other syntax elements. The parsing process for this data type is specified by

the return value of the function `read_bits( n )` interpreted as a two's complement integer representation with most significant bit written first.

- `u(n)`: unsigned integer using `n` bits. When `n` is “v” in the syntax table, the number of bits varies in a manner dependent on the value of other syntax elements. The parsing process for this data type is specified by the return value of the function `read_bits( n )` interpreted as a binary representation of an unsigned integer with most significant bit written first.
- `ue(k)`: unsigned integer `k`-th order Exp-Golomb-coded syntax element. The parsing process for this descriptor is according to the following pseudo-code with `x` as result:

```
x = 0
bit = 1
while( bit ) {
    bit = 1 - u( 1 )
    x += bit << k
    k += 1
}
k -= 1
if( k > 0 )
    x += u( k )
```

- `ie(k)`: signed integer `k`-th order Exp-Golomb-coded syntax element. The parsing process for this descriptor is according to the following pseudo-code with `x` as result:

```
val = ue( k )
if( (val & 1) != 0 )
    x = ((val+1)>>1)
else
    x = - (val>>1)
```

- `flt(n)`: Floating point value using `n` bits where `n` may be 32, 64, or 128 in little-endian byte order as specified in ISO/IEC 60559 as `binary32`, `binary64`, or `binary128`, respectively.
- `st(v)`: null-terminated string, which shall be encoded as UTF-8 characters in accordance with ISO/IEC 10646. The parsing process is specified as follows: `st(v)` begins at a byte-aligned position in the bitstream and reads and returns a series of bytes from the bitstream, beginning at the current position and continuing up to but not including the next byte-aligned byte that is equal to `0x00`, and advances the bitstream pointer by $(\text{stringLength} + 1) * 8$ bit positions, where `stringLength` is equal to the number of bytes returned.

NOTE The `st(v)` and `flt(n)` syntax descriptors are only used in this document when the current position in the bitstream is a byte-aligned position.

- `bs(v)`: Byte-sequence specifies a sequence of bytes of variable length, starting at byte-aligned position. The length of the sequence is determined from the size of the NNR unit containing the byte sequence.

6.1.4 Semantics

Semantics associated with the syntax structures and with the syntax elements within each structure are specified in a subclause following the subclause containing the syntax structures.

The following definitions apply to the semantics specification.

unspecified is used to specify some values of a particular *syntax element* to indicate that the values have no specified meaning in this document and will not have a specified meaning in the future as an integral part of future versions of this document.

reserved is used to specify that some values of a particular *syntax element* are for future use by ISO/IEC and shall not be used in *bitstreams* conforming to this version of this document, but may be used in bitstreams conforming to future extensions of this document by ISO/IEC.

nnr_reserved_zero_0bit shall be an element of length 0. Decoders shall ignore the value of **nnr_reserved_zero_0bit**.

nnr_reserved_zero_1bit, when present, shall be equal to 0 in bitstreams conforming to this version of this document. Other values for **nnr_reserved_zero_1bit** are reserved for future use by ISO/IEC. Decoders shall ignore the value of **nnr_reserved_zero_1bit**.

nnr_reserved_zero_2bits, when present, shall be equal to 0 in bitstreams conforming to this version of this document. Other values for **nnr_reserved_zero_2bits** are reserved for future use by ISO/IEC. Decoders shall ignore the value of **nnr_reserved_zero_2bits**.

nnr_reserved_zero_3bits, when present, shall be equal to 0 in bitstreams conforming to this version of this document. Other values for **nnr_reserved_zero_3bits** are reserved for future use by ISO/IEC. Decoders shall ignore the value of **nnr_reserved_zero_3bits**.

nnr_reserved_zero_5bits, when present, shall be equal to 0 in bitstreams conforming to this version of this document. Other values for **nnr_reserved_zero_5bits** are reserved for future use by ISO/IEC. Decoders shall ignore the value of **nnr_reserved_zero_5bits**.

nnr_reserved_zero_7bits, when present, shall be equal to 0 in bitstreams conforming to this version of this document. Other values for **nnr_reserved_zero_7bits** are reserved for future use by ISO/IEC. Decoders shall ignore the value of **nnr_reserved_zero_7bits**.

6.2 General bitstream syntax elements

6.2.1 NNR unit

NNR unit is the data structure for carrying neural network data and related metadata which is compressed or represented using this document.

NNR units carry compressed or uncompressed information about neural network metadata, topology information, complete or partial layer data, filters, kernels, biases, quantized weights, tensors or alike.

An NNR unit consists of the following data elements (shown in [Figure 2](#)):

- **NNR unit size**: This data element signals the total byte size of the NNR unit, including the NNR unit size.
- **NNR unit header**: This data element contains information about the NNR unit type and related metadata.
- **NNR unit payload**: This data element contains compressed or uncompressed data related to the neural network.




Figure 2 — NNR Unit data structure

6.2.2 Aggregate NNR unit

An aggregate NNR unit is an NNR unit which carries multiple NNR units in its payload. Aggregate NNR units provide a grouping mechanism for several NNR units which are related to each other and benefit from aggregation under a single NNR unit (shown in [Figure 3](#)).




Figure 3 — Aggregate NNR unit data structure

6.2.3 Composition of NNR bitstream

NNR bitstream is composed of a sequence of NNR units (shown in [Figure 4](#)).




Figure 4 — NNR bitstream data structure

In an NNR bitstream; the following constraints apply unless otherwise stated in this document or defined by NNR profiles:

(NNR_STR, NNR_MPS, NNR_NDU, NNR_LPS, NNR_TPL and NNR_QNT are NNR unit types as specified in [Table 3](#) of [subclause 6.4.3](#))

- An NNR bitstream shall start with an NNR start unit (NNR_STR) ([subclause 6.4.3](#))
- There shall be a single NNR model parameter set (NNR_MPS) ([subclause 6.4.3](#)) in an NNR bitstream which shall precede any NNR_NDU ([subclause 6.4.3](#)) in the NNR bitstream
- NNR layer parameter sets (NNR_LPS) shall be active until the next NNR layer parameter set in the NNR bitstream or until the boundary of an Aggregate NNR unit is reached.
- **topology_elem_id** and **topology_elem_id_index** ([subclause 6.4.3.7](#)) values shall be unique in the NNR bitstream.
- NNR_TPL or NNR_QNT units; if present in the NNR bitstream; shall precede any NNR_NDUs that reference their data structures (e.g. **topology_elem_id**).

6.3 NNR bitstream syntax

6.3.1 NNR unit syntax

nnr_unit() {	Descriptor
nnr_unit_size()	
nnr_unit_header()	
nnr_unit_payload()	
}	

6.3.2 NNR unit size syntax

nnr_unit_size() {	Descriptor
nnr_unit_size_flag	u(1)
nnr_unit_size	u(15 + nnr_unit_size_flag*16)
}	

6.3.3 NNR unit header syntax

6.3.3.1 General

nnr_unit_header() {	Descriptor
nnr_unit_type	u(6)
independently_decodable_flag	u(1)
partial_data_counter_present_flag	u(1)
if(partial_data_counter_present_flag)	
partial_data_counter	u(8)
if(nnr_unit_type == NNR_MPS)	
nnr_model_parameter_set_unit_header()	
if(nnr_unit_type == NNR_LPS)	
nnr_layer_parameter_set_unit_header()	
if(nnr_unit_type == NNR_TPL)	
nnr_topology_unit_header()	
if(nnr_unit_type == NNR_QNT)	
nnr_quantization_unit_header()	
if(nnr_unit_type == NNR_NDU)	
nnr_compressed_data_unit_header()	
if(nnr_unit_type == NNR_STR)	
nnr_start_unit_header()	
if(nnr_unit_type == NNR_AGG)	
nnr_aggregate_unit_header()	
}	

6.3.3.2 NNR start unit header syntax

nnr_start_unit_header() {	Descriptor
general_profile_idc	u(8)
}	

6.3.3.3 NNR model parameter set unit header syntax

nnr_model_parameter_set_unit_header() {	Descriptor
nnr_reserved_zero_0bit	u(0)
}	

6.3.3.4 NNR layer parameter set unit header syntax

nnr_layer_parameter_set_unit_header() {	Descriptor
lps_self_contained_flag	u(1)
nnr_reserved_zero_7_bits	u(7)
}	

6.3.3.5 NNR topology unit header syntax

nnr_topology_unit_header() {	Descriptor
topology_storage_format	u(8)
topology_compression_format	u(8)
}	

6.3.3.6 NNR quantization unit header syntax

nnr_quantization_unit_header() {	Descriptor
quantization_storage_format	u(8)
quantization_compression_format	u(8)
}	

6.3.3.7 NNR compressed data unit header syntax

nnr_compressed_data_unit_header() {	Descriptor
nnr_compressed_data_unit_payload_type	u(5)
nnr_multiple_topology_elements_present_flag	u(1)
nnr_decompressed_data_format_present_flag	u(1)
input_parameters_present_flag	u(1)
if(nnr_multiple_topology_elements_present_flag == 1)	
topology_elements_ids_list(mps_topology_indexed_reference_flag)	
else {	
if(!mps_topology_indexed_reference_flag)	
topology_elem_id	st(v)
else	
topology_elem_id_index	ue(7)
}	

if(general_profile_idc == 1) {	
node_id_present_flag	u(1)
if(node_id_present_flag) {	
device_id	ue(1)
parameter_id	ue(5)
put_node_depth	ue(4)
}	
if(mps_parent_signalling_enabled_flag == 1) {	
parent_node_id_present_flag	u(1)
if(parent_node_id_present_flag) {	
parent_node_id_type	u(2)
temporal_context_modeling_flag	u(1)
if(parent_node_id_type == ICNN_NDU_ID) {	
parent_device_id	ue(1)
if(!node_id_present_flag) {	
parameter_id	ue(5)
put_node_depth	ue(4)
}	
} else if(parent_node_id_type == ICNN_NDU_PL_SHA256)	
parent_node_payload_sha256	u(256)
else if(parent_node_id_type == ICNN_NDU_PL_SHA512)	
parent_node_payload_sha512	u(512)
}	
}	
}	
if(nnr_compressed_data_unit_payload_type == NNR_PT_FLOAT nnr_compressed_data_unit_payload_type == NNR_PT_BLOCK) {	
codebook_present_flag	u(1)
if(codebook_present_flag)	
integer_codebook(CbZeroOffset, Codebook, CbSize)	
}	
if(nnr_compressed_data_unit_payload_type == NNR_PT_INT nnr_compressed_data_unit_payload_type == NNR_PT_FLOAT nnr_compressed_data_unit_payload_type == NNR_PT_BLOCK)	
dq_flag	u(1)
if(nnr_decompressed_data_format_present_flag == 1)	
nnr_decompressed_data_format	u(7)
if(input_parameters_present_flag == 1) {	
tensor_dimensions_flag	u(1)
cabac_unary_length_flag	u(1)
compressed_parameter_types	u(4)
if((compressed_parameter_types & NNR_CPT_DC) != 0) {	
decomposition_rank	ue(3)
g_number_of_rows	ue(3)
}	

if(tensor_dimensions_flag == 1)	
tensor_dimension_list()	
if (nnr_compressed_data_unit_payload_type != NNR_PT_BLOCK)	
if(nnr_multiple_topology_elements_present_flag == 1)	
topology_tensor_dimension_mapping()	
if(cabac_unary_length_flag == 1)	
cabac_unary_length_minus1	u(8)
}	
if(nnr_compressed_data_unit_payload_type == NNR_PT_BLOCK && (compressed_parameter_types & NNR_CPT_DC) != 0 && codebook_present_flag)	
integer_codebook(CbZeroOffsetDC, CodebookDC, CbSizeDC)	
if(count_tensor_dimensions > 1) {	
if(general_profile_idc == 1)	
first_tensor_dimension_shift	ue(1)
scan_order	u(4)
if(scan_order > 0) {	
for(j=0; j < NumBlockRowsMinus1; j++) {	
cabac_offset_list[j]	u(8)
if(dq_flag)	
dq_state_list[j]	u(3)
if(j == 0) {	
bit_offset_delta1	ue(11)
BitOffsetList[j] = bit_offset_delta1	
}	
else {	
bit_offset_delta2	ie(7)
BitOffsetList[j] = BitOffsetList[j-1] + bit_offset_delta2	
}	
}	
}	
}	
byte_alignment()	
}	

integer_codebook() is defined as follows:

integer_codebook[cbZeroOffset, integerCodebook, cbSize] {	Descriptor
codebook_egk	u(4)
codebook_size	ue(2)
cbSize = codebook_size	
codebook_centre_offset	ie(2)
cbZeroOffset = (codebook_size >> 1) + codebook_centre_offset	
codebook_zero_value	ie(7)
integerCodebook[cbZeroOffset] = codebook_zero_value	
previousValue = integerCodebook[cbZeroOffset]	

for(j = cbZeroOffset - 1; j >= 0; j--) {	
codebook_delta_left	ue(codebook_egk)
integerCodebook[j] = previousValue - codebook_delta_left - 1	
previousValue = integerCodebook[j]	
}	
previousValue = integerCodebook[cbZeroOffset]	
for(j = cbZeroOffset + 1; j < codebook_size; j++) {	
codebook_delta_right	ue(codebook_egk)
integerCodebook[j] = previousValue + codebook_delta_right + 1	
previousValue = integerCodebook[j]	
}	
}	

tensor_dimension_list() is defined as follows:

tensor_dimension_list() {	Descriptor
count_tensor_dimensions	ue(1)
for(j = 0; j < count_tensor_dimensions; j++)	
tensor_dimensions[j]	ue(7)
}	

topology_elements_ids_list(topologyIndexedFlag) is defined as follows:

topology_elements_ids_list(topologyIndexedFlag) {	Descriptor
count_topology_elements_minus2	ue(7)
if(topologyIndexedFlag == 0)	
byte_alignment()	
for(j = 0; j < count_topology_elements_minus2 + 2; j++) {	
if (topologyIndexedFlag == 0) {	
topology_elem_id_list[j]	st(v)
}	
else {	
topology_elem_id_index_list[j]	ue(7)
}	
}	
if (topologyIndexedFlag == 1)	
byte_alignment()	
}	

topology_tensor_dimension_mapping() is defined as follows:

topology_tensor_dimension_mapping () {	Descriptor
concatentation_axis_index	u(8)
for(j = 0; j < count_topology_elements_minus2 + 1 ; j++) {	
split_index[j]	ue(7)
}	
for(k = 0; k < count_topology_elements_minus2 + 2 ; k++) {	

number_of_shifts [k]	ue(1)
for(i = 0; i < number_of_shifts[k] ; i++) {	
shift_index [k][i]	ue(7)
shift_value [k][i]	ue(1)
}	
}	
}	

6.3.3.8 NNR aggregate unit header syntax

nnr_aggregate_unit_header() {	Descriptor
nnr_aggregate_unit_type	u(8)
entry_points_present_flag	u(1)
nnr_reserved_zero_7bits	u(7)
num_of_nnr_units_minus2	u(16)
if(entry_points_present_flag)	
for(i = 0; i < num_of_nnr_units_minus2 + 2; i++) {	
nnr_unit_type [i]	u(6)
nnr_unit_entry_point [i]	u(34)
}	
for(i = 0; i < num_of_nnr_units_minus2 + 2; i++) {	
quant_bitdepth [i]	u(5)
if(mps_unification_flag lps_unification_flag){	
ctu_scan_order [i]	u(1)
nnr_reserved_zero_2bits	u(2)
}	
else	
nnr_reserved_zero_3bits	u(3)
}	
}	

6.3.4 NNR unit payload syntax

6.3.4.1 General

nnr_unit_payload() {	Descriptor
if(nnr_unit_type == NNR_MPS)	
nnr_model_parameter_set_unit_payload()	
if(nnr_unit_type == NNR_LPS)	
nnr_layer_parameter_set_unit_payload()	
if(nnr_unit_type == NNR_TPL)	
nnr_topology_unit_payload()	
if(nnr_unit_type == NNR_QNT)	
nnr_quantization_unit_payload()	
if(nnr_unit_type == NNR_NDU)	
nnr_compressed_data_unit_payload()	

if(nnr_unit_type == NNR_STR)	
nnr_start_unit_payload()	
if(nnr_unit_type == NNR_AGG)	
nnr_aggregate_unit_payload()	
}	

6.3.4.2 NNR start unit payload syntax

nnr_start_unit_payload() {	Descriptor
nnr_reserved_zero_0bit	u(0)
}	

6.3.4.3 NNR model parameter set unit payload syntax

nnr_model_parameter_set_unit_payload() {	Descriptor
topology_carriage_flag	u(1)
mps_sparsification_flag	u(1)
mps_pruning_flag	u(1)
mps_unification_flag	u(1)
mps_decomposition_performance_map_flag	u(1)
mps_quantization_method_flags	u(3)
mps_topology_indexed_reference_flag	u(1)
if(general_profile_idc == 1) {	
base_model_id_present_flag	u(1)
validation_set_performance_present_flag	u(1)
metric_type_performance_map_valid_flag	u(1)
mps_parent_signalling_enabled_flag	u(1)
if(mps_parent_signalling_enabled_flag == 1)	
nnr_pre_flag	u(1)
else	
nnr_reserved_zero_1bit	u(1)
nnr_reserved_zero_2bits	u(2)
if(base_model_id_present_flag == 1)	
base_model_id	st(v)
if(validation_set_performance_present_flag == 1 metric_type_performance_map_valid_flag == 1)	
performance_metric_type	st(v)
}	
else	
nnr_reserved_zero_7bits	u(7)
if((mps_quantization_method_flags & NNR_QSU) == NNR_QSU (mps_quantization_method_flags & NNR_QCB) == NNR_QCB) {	
mps_qp_density	u(3)
mps_quantization_parameter	i(13)
}	
if(mps_sparsification_flag == 1)	

sparsification_performance_map()	
if(mps_pruning_flag == 1)	
pruning_performance_map()	
if(mps_unification_flag == 1)	
unification_performance_map()	
if(mps_decomposition_performance_map_flag == 1)	
decomposition_performance_map()	
if(general_profile_idc == 1 && validation_set_performance_present_flag == 1)	
validation_set_performance	flt(32)
byte_alignment()	
}	

sparsification_performance_map() is defined as follows:

sparsification_performance_map() {	Descriptor
spm_count_thresholds	u(8)
for(i = 0; i < (spm_count_thresholds-1); i++) {	
sparsification_threshold[i]	flt(32)
non_zero_ratio[i]	flt(32)
spm_nn_accuracy[i]	flt(32)
spm_count_classes[i]	u(8)
spm_class_bitmask[i]	ue(7)
for (j = 0; j < spm_count_classes[i]; j++)	
spm_nn_class_accuracy[i][j]	flt(32)
}	
}	

pruning_performance_map() is defined as follows:

pruning_performance_map() {	Descriptor
ppm_count_pruning_ratios	u(8)
for(i = 0; i < (ppm_count_pruning_ratios-1); i++) {	
pruning_ratio[i]	flt(32)
ppm_nn_accuracy[i]	flt(32)
ppm_count_classes[i]	u(8)
ppm_class_bitmask[i]	ue(7)
for(j = 0; j < ppm_count_classes[i]; j++)	
ppm_nn_class_accuracy[i][j]	flt(32)
}	
}	

unification_performance_map() is defined as follows:

unification_performance_map() {	Descriptor
upm_count_thresholds	u(8)
for(i = 0; i < (upm_count_thresholds-1); i++) {	
count_resaped_tensor_dimension	ue(1)
for(j = 0; j < (count_resaped_tensor_dimension-1); j++)	
reshaped_tensor_dimensions[j]	ue(7)
byte_alignment()	
count_super_block_dimension	u(8)
for(j = 0; j < (count_super_block_dimension-1); j++)	
super_block_dimensions[j]	u(8)
count_block_dimension	u(8)
for(j = 0; j < (count_block_dimension-1); j++)	
block_dimensions[j]	u(8)
unification_threshold[i]	flt(32)
upm_nn_accuracy[i]	flt(32)
upm_count_classes[i]	u(8)
upm_class_bitmask[i]	ue(7)
for(j = 0; j < upm_count_classes[i]; j++)	
upm_nn_class_accuracy[i][j]	flt(32)
}	
}	

Decomposition_performance_map() is defined as follows:

decomposition_performance_map() {	Descriptor
dpm_count_thresholds	u(8)
for(i = 0; i < (dpm_count_thresholds-1); i++) {	
mse_threshold[i]	flt(32)
dpm_nn_accuracy[i]	flt(32)
nn_reduction_ratio[i]	flt(32)
dpm_count_classes[i]	u(16)
for(j = 0; j < dpm_count_classes[i]; j++)	
dpm_nn_class_accuracy[i][j]	flt(32)
}	
}	

6.3.4.4 NNR layer parameter set unit payload syntax

nnr_layer_parameter_set_unit_payload() {	Descriptor
nnr_reserved_zero_1_bit	u(1)
lps_sparsification_flag	u(1)
lps_pruning_flag	u(1)
lps_unification_flag	u(1)
lps_quantization_method_flags	u(3)
nnr_reserved_zero_1bit	u(1)

if((lps_quantization_method_flags & NNR_QCB) == NNR_QCB (lps_quantization_method_flags & NNR_QSU) == NNR_QSU) {	
lps_qp_density	u(3)
lps_quantization_parameter	i(13)
}	
if(lps_sparsification_flag == 1)	
sparsification_performance_map()	
if(lps_pruning_flag == 1)	
pruning_performance_map()	
if(lps_unification_flag == 1)	
unification_performance_map()	
byte_alignment()	
}	

6.3.4.5 NNR topology unit payload syntax

nnr_topology_unit_payload() {	Descriptor
if(topology_storage_format == NNR_TPL_PRUN)	
nnr_pruning_topology_container()	
else if(topology_storage_format == NNR_TPL_REFLIST)	
topology_elements_ids_list(0)	
else	
topology_data	bs(v)
}	

nnr_pruning_topology_container() is specified as follows:

nnr_pruning_topology_container() {	Descriptor
nnr_rep_type	u(2)
prune_flag	u(1)
order_flag	u(1)
sparse_flag	u(1)
nnr_reserved_zero_3bits	u(3)
if (prune_flag == 1) {	
if (nnr_rep_type == NNR_TPL_BMSK)	
bit_mask()	
else if(nnr_rep_type == NNR_TPL_DICT) {	
count_ids	ue(7)
if (!mps_topology_indexed_reference_flag) {	
byte_alignment()	
for (j = 0; j < count_ids; j++) {	
element_id[j]	st(v)
}	
}	
else {	
for (j = 0; j < count_ids; j++) {	

element_id_index [j]	ue(7)
}	
}	
for (j = 0; j < count_ids; j++) {	
count_dims [j]	ue(1)
for(k = 0; k < count_dims[j]; k++){	
dim [j][k]	ue(7)
}	
}	
byte_alignment()	
}	
}	
if (sparse_flag == 1) {	
bit_mask()	
}	
}	

bit_mask() is specified as follows:

bit_mask() {	Descriptor
count_bits	u(32)
for(j = 0; j < count_bits; j++) {	
bit_mask_value [j]	u(1)
}	
byte_alignment()	
}	

6.3.4.6 NNR quantization unit payload syntax

nnr_quantization_unit_payload() {	Descriptor
quantization_data	bs(v)
}	

6.3.4.7 NNR compressed data unit payload syntax

nnr_compressed_data_unit_payload() {	Descriptor
if(nnr_compressed_data_unit_payload_type == NNR_PT_RAW_FLOAT)	
for(i = 0; i < Prod(TensorDimensions); i++)	
raw_float32_parameter [TensorIndex(TensorDimensions, i , 0)]	flt(32)
decode_compressed_data_unit_payload()	
}	

decode_compressed_data_unit_payload() invokes the decoding process as specified in [subclause 7.3](#).

6.3.4.8 NNR aggregate unit payload syntax

nnr_aggregate_unit_payload() {	Descriptor
for(i = 0; i < num_of_nnr_units_minus2 + 2; i++)	
nnr_unit()	
}	

6.3.5 Byte alignment syntax

byte_alignment() {	Descriptor
alignment_bit_equal_to_one /* equal to 1 */	f(1)
while(!byte_aligned())	
alignment_bit_equal_to_zero /* equal to 0 */	f(1)
}	

6.4 Semantics

6.4.1 General

Semantics associated with the syntax structures and elements within these structures are specified in this subclause. When the semantics of a syntax element are specified using a table or a set of tables, any values that are not specified in the table(s) shall not be present in the bitstream unless otherwise specified in this document.

6.4.2 NNR unit size semantics

nnr_unit_size_flag specifies the number of bits used as the data type of the `nnr_unit_size`. If this value is 0, then `nnr_unit_size` is a 15 bits unsigned integer value, otherwise it is 31 bits unsigned integer value.

nnr_unit_size specifies the size of the NNR unit, which is the sum of byte sizes of `nnr_unit_size()`, `nnr_unit_header()` and `nnr_unit_payload()`.

6.4.3 NNR unit header semantics

6.4.3.1 General

nnr_unit_type specifies the type of the NNR unit, as specified in [Table 3](#).

Table 3 — NNR unit Types

nnr_unit_type	Identifier	NNR unit Type	Description
0	NNR_STR	NNR start unit	Compressed neural network bitstream start indicator
1	NNR_MPS	NNR model parameter set data unit	Neural network global metadata and information
2	NNR_LPS	NNR layer parameter set data unit	Metadata related to a partial representation of neural network
3	NNR_TPL	NNR topology data unit	Neural network topology information
4	NNR_QNT	NNR quantization data unit	Neural network quantization information
5	NNR_NDU	NNR compressed data unit	Compressed neural network data
6	NNR_AGG	NNR aggregate unit	NNR unit with payload containing multiple NNR units
7..31	NNR_RSVD	Reserved	ISO/IEC-reserved range
32..63	NNR_UNSP	Unspecified	Unspecified range

The values in the range NNR_RSVD are reserved for use in future versions of this or related specifications. Encoders shall not use these values. Decoders conforming to this version of the specification may ignore NNR units using these values. The values in the range NNR_UNSP are not specified, their use is outside the scope of this specification. Decoders conforming to this version of the specification may ignore NNR units using these values.

independently_decodable_flag specifies whether this compressed data unit is independently decodable. A value of 1 indicates an independently decodable NNR unit. A value of 0 indicates that this NNR unit is not independently decodable and its payload should be combined with other NNR units for successful decodability/decompressibility. The value of `independently_decodable_flag` shall be the same for all NNR units which refer to the same `topology_elem_id` or `topology_elem_id_index` value or the same `topology_elem_id_list`.

partial_data_counter_present_flag equal to 1 specifies that the syntax element `partial_data_counter` is present in NNR unit header. `partial_data_counter_present_flag` equal to 0 specifies that the syntax element `partial_data_counter` is not present in NNR unit header.

partial_data_counter specifies the index of the partial data carried in the payload of this NNR data unit with respect to the whole data for a certain topology element. A value of 0 indicates no partial information (i.e. the data in this NNR unit is all data associated to a topology element and it is complete), a value bigger than 0 indicates the index of the partial information (i.e. data in this NNR unit should be concatenated with the data in accompanying NNR units until `partial_data_counter` of an NNR unit reaches 1). This counter counts backwards to indicate initially the total number of partitions. If not present, the value of `partial_data_counter` is inferred to be equal to 0. If the value of `independently_decodable_flag` is equal to 0, the value of `partial_data_counter_present_flag` shall be equal to 1 and the value of `partial_data_counter` shall be greater than 0. If the value of `independently_decodable_flag` is equal to 1, the values of `partial_data_counter_present_flag` and `partial_data_counter` are undefined, in this version of this document.

NOTE In future versions of this document, if the value of `independently_decodable_flag` is equal to 1 and if `partial_data_counter_present_flag` is equal to 1, `partial_data_counter` can have non-zero values, based on the assumption that multiple independently decodable NNR units are combined to construct a model.

6.4.3.2 NNR start unit header semantics

general_profile_idc indicates a profile to which NNR bitstream conforms as specified in this document.

Value	Semantics
0	Bitstream conforming to the base feature set (as defined below).
1	Bitstream conforming to the extended feature set (as defined below).

The base feature set contains the following compression tools:

- unstructured sparsification with compressibility loss
- structured sparsification using micro-structured sparsification
- parameter unification
- combined pruning and sparsification
- low rank/low displacement rank for convolutional and fully connected layers
- batchnorm folding
- local scaling adaptation
- uniform quantization
- codebook-based quantization
- dependent scalar quantization

— DeepCABAC

The syntax of the base feature set excludes all syntax elements gated with `general_profile_idc > 0`.

NOTE The base feature set corresponds to the now obsolete first edition.

In addition, the extended feature set contains the following compression tools (i.e. all compression tools specified in this document):

- unstructured statistics-adaptive sparsification
- structured sparsification
- iterative QP optimization
- stochastic binary-ternary quantization
- predictive residual encoding
- row skipping and temporal context modeling for DeepCABAC

Some compression tools of the base feature set have been adapted for the compression of differential updates.

The syntax of the extended feature set includes the complete syntax defined in this document, in particular, syntax elements to reference components of base neural networks to which updates shall be applied.

6.4.3.3 NNR model parameter set unit header semantics

Header elements of the model parameter set (reserved for future use).

6.4.3.4 NNR layer parameter set unit header semantics

lps_self_contained_flag equal to 1 specifies that NNR units that refer to the layer parameter set are a full or partial NN model and shall be successfully reconstructable with the NNR units. A value of 0 indicates that the NNR units that refer to the layer parameter set should be combined with NNR units that refer to other layer parameter sets for successful reconstruction of a full or partial NN model.

6.4.3.5 NNR topology unit header semantics

topology_storage_format specifies the format of the stored neural network topology information, as specified in [Table 4](#).

Table 4 — Topology storage format identifiers

topology_storage_format value	Identifier	Description
0	NNR_TPL_UNREC	Unrecognized topology format
1		Topology format shall be represented as specified in Annex A .
2..4		See Annexes B to D for further information.
5	NNR_TPL_PRUN	Topology pruning information
6	NNR_TPL_REFLIST	Topology element reference list information
7..127	NNR_TPL_RSVD	ISO/IEC-reserved range
128..255	NNR_TPL_UNSP	Unspecified range

The value NNR_TPL_UNREC indicates that the topology format is unknown. Encoders may use this value if the topology format used is not among the set of formats for which identifiers are specified. Decoders conforming to this version of the specification may ignore NNR units using this value or may attempt to recognize the format by parsing the start of the topology payload.

The values in the range NNR_TPL_RSVD are reserved for use in future versions of this or related specifications. Encoders shall not use these values. Decoders conforming to this version of the specification may ignore NNR units using these values. The values in the range NNR_TPL_UNSP are not specified, their use is outside the scope of this specification. Decoders conforming to this version of the specification may ignore NNR units using these values.

topology_compression_format specifies that one of the compression formats defined in [Table 5](#) is applied on the stored topology data in topology_data.

Table 5 — Topology compression format identifiers

topology_compression_format	Identifier	Description
0x00	NNR_PT_RAW	Uncompressed
0x01	NNR_DFL	Deflate method, shall be implemented according to IETF RFC 1950
0x02-0xFF		Reserved

6.4.3.6 NNR quantization unit header semantics

quantization_storage_format specifies the format of the stored neural network quantization information, as specified in [Table 6](#).

Table 6 — Quantization storage format identifiers

quantization_storage_format value	Identifier	Description
0	NNR_QNT_UNREC	Unrecognized quantization format.
1		Topology format shall be represented as specified in Annex A .
2..4		See Annexes B to D for further information.
5..127	NNR_QNT_RSVD	ISO/IEC-reserved range
128..255	NNR_QNT_UNSP	Unspecified range

The value NNR_QNT_UNREC indicates that the quantization format is unknown. Encoders may use this value if the quantization format used is not among the set of formats for which identifiers are specified. Decoders conforming to this version of the specification may ignore NNR units using this value or may attempt to recognize the format by parsing the start of the topology payload.

The values in the range NNR_QNT_RSVD are reserved for use in future versions of this or related specifications. Encoders shall not use these values. Decoders conforming to this version of the specification may ignore NNR units using these values. The values in the range NNR_QNT_UNSP are not specified, their use is outside the scope of this specification. Decoders conforming to this version of the specification may ignore NNR units using these values.

quantization_compression_format specifies that one of the compression formats defined in [Table 7](#) is applied on the stored quantization data in quantization_data.

Table 7 — Quantization compression format identifiers

quantization_compression_format	Identifier	Description
0x00	NNR_PT_RAW	Uncompressed
0x01	NNR_DFL	Deflate method, shall be implemented according to IETF RFC 1950
0x02-0xFF		Reserved

6.4.3.7 NNR compressed data unit header semantics

nnr_compressed_data_unit_payload_type is as defined in [Table 16](#) of [subclause 7.3.1](#).

nnr_multiple_topology_elements_present_flag specifies whether multiple topology units are present in the bitstream. In case there are multiple units, the list of their IDs is included. When **nnr_compressed_data_unit_payload_type** is set to NNR_PT_BLOCK, this flag shall be set to 1 and **topology_elements_ids_list()** in the NNR compressed data unit header shall list the topology elements or topology element indexes of RecWeight, RecWeightG, RecWeightH, RecLS, RecBeta, RecGamma, RecMean, RecVar and RecBias, in the given order and based on their presence as indicated by the value of **compressed_parameter_type** in the NNR compressed data unit header.

nnr_decompressed_data_format_present_flag specifies whether the data format to be obtained after decompression is present in the bitstream.

input_parameters_present_flag specifies whether the group of elements including tensor dimensions, DeepCABAC unary length and compressed parameter types is present in the bitstream.

topology_elem_id specifies a unique identifier for the topology element to which an NNR compressed data unit refers. The semantic interpretation of this field is context dependent.

topology_elem_id_index specifies a unique index value of a topology element which is signalled in topology information of payload type NNR_TPL_REFLIST. The first index shall be 0 (i.e. 0-indexed).

node_id_present_flag equal to 1 indicates that syntax elements **device_id**, **parameter_id**, and **put_node_depth** are present.

device_id uniquely identifies the context in which the current NDU has been generated. The use of this identifier is defined by the application, and multiple such identifiers may be used on the same device.

parameter_id uniquely identifies the parameter of the model to which the tensors stored in the NDU relate to. If **parent_node_id_type** is equal to ICNN_NDU_ID, **parameter_id** shall equal the **parameter_id** of the associated parent NDU.

put_node_depth determines the order of incremental updates of NDUs. If **parent_node_id_type** is equal to ICNN_NDU_ID, **put_node_depth** – 1 shall equal the **put_node_depth** of the associated parent NDU.

parent_node_id_present_flag indicates whether the NDU represents a differential update of a base neural network. It shall be set to 1 for NDUs representing differential updates, and to 0 otherwise. A value of 1 also indicates that syntax element **parent_node_id_type** is present. If **parent_node_id_present_flag** is not present, it is inferred to be 0.

parent_node_id_type specifies the parent node id type. It indicates which further syntax elements for uniquely identifying the parent node are present. The allowed values for **parent_node_id_type** are defined in [Table 8](#).

Table 8 — Parent node id type identifiers

parent_node_id_type	Identifier	Description
0	ICNN_NDU_ID	Indicates that syntax elements parent_device_id , parameter_id , and put_node_depth are present
1	ICNN_NDU_PL_SHA256	Indicates that syntax element parent_node_payload_sha256 is present
2	ICNN_NDU_PL_SHA512	Indicates that syntax element parent_node_payload_sha512 is present
3	ICNN_NDU_EXT	Indicates that the parent node is not part of an NNC bitstream. topology_elem_id or topology_elem_id_index are used to identify the corresponding tensor in the base neural network.

temporal_context_modeling_flag specifies whether temporal context modeling is enabled. A `temporal_context_modeling_flag` equal to 1 indicates that temporal context modeling is enabled. If `temporal_context_modeling_flag` is not present, it is inferred to be 0.

parent_device_id is equal to syntax element `device_id` of the parent NDU.

parent_node_payload_sha256 shall conform to a SHA* hash as specified in FIPS PUB 180-4, of the `nnr_compressed_data_unit_payload` of the parent NDU.

parent_node_payload_sha512 shall conform to a SHA* hash as specified in FIPS PUB 180-4, of the `nnr_compressed_data_unit_payload` of the parent NDU.

count_topology_elements_minus2 + 2 specifies the number of topology elements for which this NNR compressed data unit carries data in the payload.

codebook_present_flag specifies whether codebooks are used. If `codebook_present_flag` is not present, it is inferred to be 0.

dq_flag specifies whether the quantization method is dependent scalar quantization according to [subclause 9.2.3](#) or uniform quantization according to [subclause 9.2.1](#). A `dq_flag` equal to 0 indicates that the uniform quantization method is used. A `dq_flag` equal to 1 indicates that the dependent scalar quantization method is used. If `dq_flag` is not present, it is inferred to be 0.

nnr_decompressed_data_format is defined in [Table 15](#) of [subclause 7.2](#). If `general_profile_idc` has the value 0, `nnr_decompressed_data_format` may only take the value 0 and 1. If `nnr_compressed_data_unit_payload_type` is equal to `NNR_PT_INT`, `nnr_decompressed_data_format` shall be set to a format with data type `TENSOR_INT`, otherwise to a format with data type `TENSOR_FLOAT`.

tensor_dimensions_flag specifies whether the tensor dimensions are defined in the bitstream. If they are not included in the bitstream, they shall be obtained from the model topology description.

cabac_unary_length_flag specifies whether the length of the unary part in the DeepCABAC binarization is included in the bitstream.

compressed_parameter_types specifies the compressed parameter types present in the current topology element to which an NNR compressed data unit refers. If multiple compressed parameter types are specified, they are combined by OR. The compressed parameter types are defined in [Table 9](#).

Table 9 — Compressed parameter type identifiers

Compressed parameter type	Compressed parameter type ID	Bit in <code>compressed_parameter_types</code>
Decomposition present	<code>NNR_CPT_DC</code>	0x01
Local scaling present	<code>NNR_CPT_LS</code>	0x02
Batch norm parameters present	<code>NNR_CPT_BN</code>	0x04
Bias present	<code>NNR_CPT_BI</code>	0x08

When decomposition is present, the tensors G and H represent the result of decomposing the original tensor. If $(\text{compressed_parameter_types} \ \& \ \text{NNR_CPT_DC}) \neq 0$ the variables `TensorDimensionsG` and `TensorDimensionsH` are derived as follows:

- Variable `TensorDimensionsG` is set to `[g_number_of_rows, decomposition_rank]`.
- Variable `TensorDimensionsH` is set to `[decomposition_rank, hNumberOfColumns]` where `hNumberOfColumns` is defined as

$$\text{hNumberOfColumns} = \frac{\prod_{i=0}^{\text{count_tensor_dimensions}-1} \text{tensor_dimensions}[i]}{\text{g_number_of_rows}}$$

If (compressed_parameter_types & NNR_CPT_DC) != 0 and if nnr_compressed_data_unit_payload_type != NNR_PT_BLOCK, the NNR unit contains a decomposed tensor G and the next NNR unit in the bitstream contains the corresponding decomposed tensor H.

A variable TensorDimensions is derived as follows:

- If an NNR unit contains a decomposed tensor G and nnr_compressed_data_unit_payload_type != NNR_PT_BLOCK, TensorDimensions is set to TensorDimensionsG.
- Otherwise, if an NNR unit contains a decomposed tensor H and nnr_compressed_data_unit_payload_type != NNR_PT_BLOCK, TensorDimensions is set to TensorDimensionsH.
- Otherwise, TensorDimensions is set to tensor_dimensions.

A variable NumBlockRowsMinus1 is defined as follows:

- If scan_order is equal to 0, NumBlockRowsMinus1 is set to 0.
- Otherwise, if nnr_compressed_data_unit_payload_type == NNR_PT_BLOCK and (compressed_parameter_types & NNR_CPT_DC) != 0, NumBlockRowsMinus1 is set to $((\text{TensorDimensionsG}[0] + (4 \ll \text{scan_order}) - 1) \gg (2 + \text{scan_order})) + ((\text{TensorDimensionsH}[0] + (4 \ll \text{scan_order}) - 1) \gg (2 + \text{scan_order})) - 2$.
- Otherwise, NumBlockRowsMinus1 is set to $((\text{TensorDimensions}[0] + (4 \ll \text{scan_order}) - 1) \gg (2 + \text{scan_order})) - 1$.

decomposition_rank specifies the rank of the low-rank decomposed weight tensor components relative to tensor_dimensions.

g_number_of_rows specifies the number of rows of matrix G in the case where the reconstruction is performed for decomposed tensors in an NNR unit of type NNR_PT_BLOCK.

cabac_unary_length_minus1 specifies the length of the unary part in the DeepCABAC binarization minus 1.

first_tensor_dimension_shift specifies the shift of the first tensor dimension for tensor dimension reordering and shall be smaller than the value of count_tensor_dimensions. If first_tensor_dimension_shift is not present, it is inferred to be 0.

scan_order specifies the block scanning order for parameters with more than one dimension according to the following table:

- 0: No block scanning
- 1: 8x8 blocks
- 2: 16x16 blocks
- 3: 32x32 blocks
- 4: 64x64 blocks

cabac_offset_list specifies a list of values to be used to initialize variable lvlOffset at the beginning of entry points.

dq_state_list specifies a list of values to be used to initialize variable stateId at the beginning of entry points.

bit_offset_delta1 specifies the first element of list BitOffsetList.

bit_offset_delta2 specifies elements of list BitOffsetList except for the first element, as difference to the previous element of list BitOffsetList.

Variable BitOffsetList is a list of bit offsets to be used to set the bitstream pointer position at the beginning of entry points.

codebook_egk specifies the Exp-Golomb parameter k for decoding of syntax elements `codebook_delta_left` and `codebook_delta_right`.

codebook_size specifies the number of elements in the codebook. It is used for setting variable `CbSize`.

codebook_centre_offset specifies an offset for accessing elements in the codebook relative to the centre of the codebook. It is used for calculating variable `CbZeroOffset`.

codebook_zero_value specifies the value of the codebook at position `CbZeroOffset`. It is involved in creating variable `Codebook` (the array representing the codebook).

codebook_delta_left specifies the difference between a codebook value and its right neighbour minus 1 for values left to the centre position. It is involved in creating variable `Codebook` (the array representing the codebook).

codebook_delta_right specifies the difference between a codebook value and its left neighbour minus 1 for values right to the centre position. It is involved in creating variable `Codebook` (the array representing the codebook).

count_tensor_dimensions specifies a counter of how many dimensions are specified. For example, for a 4-dimensional tensor, `count_tensor_dimensions` is 4. If it is not included in the bitstream, it shall be obtained from the model topology description.

tensor_dimensions specifies an array or list of dimension values. For example, for a convolutional layer, `tensor_dimensions` is an array or list of length 4. For NNR units carrying elements `G` or `H` of a decomposed tensor, `tensor_dimensions` is set to the dimensions of the original tensor. The actual tensor dimensions of `G` and `H` for the decoding methods are derived from `tensor_dimensions`, `decomposition_rank`, and `g_number_of_rows`. If it is not included in the bitstream, it shall be obtained from the model topology description.

topology_elem_id_list specifies a list of unique identifiers related to the topology element to which an NNR compressed data unit refers. Elements of `topology_elem_id_list` are semantically equivalent to syntax element `topology_elem_id` or the index of it when listed in topology payload of type `NNR_TPL_REFLIST`. The semantic interpretation of this field is context dependent.

topology_elem_id_index_list specifies a list of unique indexes related to the topology elements listed in topology information with payload type `NNR_TPL_REFLIST`. The first element in the topology shall have the index value of 0.

concatentation_axis_index indicates the 0-based concatenation axis.

split_index[] indicates the tensor splitting index along the concatenation axis indicated by `concatentation_axis_index` in order to generate each individual tensor which is concatenated.

number_of_shifts[] indicates how many left-shifting operations are to be performed.

shift_index[k][i] indicates the axis index of the k th topology element to be left-shifted.

shift_value[k][i] indicates the amount of left-shift on the axis with index `index[k][i]`.

6.4.3.8 NNR aggregate unit header semantics

nnr_aggregate_unit_type specifies the type of the aggregate NNR unit.

The NNR aggregate unit types are specified in [Table 10](#).

Table 10 — NNR aggregate unit types

nnr_aggregate_unit_type	Identifier	NNR Aggregate Unit Type	Description
0	NNR_AGG_GEN	Generic NNR aggregate unit	A set of NNR units
1	NNR_AGG_SLF	Self-contained NNR aggregate unit	When extracted and then concatenated with an NNR_STR and NNR_MPS, an NNR_AGG_SLF shall be decodable without any need of additional information and a full or partial NN model shall be successfully reconstructable with it.
2..127	NNR_RSVD	Reserved	ISO/IEC-reserved range
128..255	NNR_UNSP	Unspecified	Unspecified range

The values in the range NNR_RSVD are reserved for uses in future versions of this or related specifications. Encoders shall not use these values. Decoders conforming to this version of the specification may ignore NNR units using these values. The values in the range NNR_UNSP are not specified, their use is outside the scope of this specification. Decoders conforming to this version of the specification may ignore NNR units using these values.

entry_points_present_flag specifies whether individual NNR unit entry points are present.

num_of_nnr_units_minus2 + 2 specifies the number of NNR units present in the NNR aggregate unit's payload.

nnr_unit_type[i] specifies the NNR unit type of the NNR unit with index i. This value shall be the same as the NNR unit type of the NNR unit at index i.

nnr_unit_entry_point[i] specifies the byte offset from the start of the NNR aggregate unit to the start of the NNR unit in NNR aggregate unit's payload and at index i. This value shall not be equal or greater than the total byte size of the NNR aggregate unit. nnr_unit_entry_point values can be used for fast and random access to NNR units inside the NNR aggregate unit payload.

quant_bitdepth[i] specify the max bit depth of quantized coefficients for each tensor in the NNR aggregate unit.

ctu_scan_order[i] specify the CTU-wise scan order for each tensor in the NNR aggregate unit. Value 0 indicates that the CTU-wise scan order is raster scan order at horizontal direction, value 1 indicates that the CTU-wise scan order is raster scan order at vertical direction.

6.4.4 NNR unit payload semantics

6.4.4.1 General

The following clauses define the semantics of NNR units.

6.4.4.2 NNR start unit payload semantics

Start unit payload (reserved for future use).

6.4.4.3 NNR model parameter set unit payload semantics

topology_carriage_flag specifies whether the NNR bitstream carries the topology internally or externally. When set to 1, it specifies that topology is carried within one or more NNR units of type "NNR_TPL". If 0, it specifies that topology is provided externally (i.e. out-of-band with respect to the NNR bitstream).

mpps_sparsification_flag specifies whether sparsification is applied to the model in the NNR compressed data units that utilize this model parameter set.

mps_pruning_flag specifies whether pruning is applied to the model in the NNR compressed data units that utilize this model parameter set.

mps_unification_flag specifies whether unification is applied to the model in the NNR compressed data units that utilize this model parameter set.

mps_decomposition_performance_map_flag equal to 1 specifies that tensor decomposition was applied to at least one layer of the model and a corresponding performance map is transmitted.

mps_quantization_method_flags specifies the quantization method(s) used for the model in the NNR compressed data units that utilize this model parameter set. If multiple models are specified, they are combined by OR. The methods are defined in [Table 11](#).

Table 11 — Quantization method identifiers

Quantization method	Quantization method ID	Value
Scalar uniform	NNR_QSU	0x01
Codebook	NNR_QCB	0x02
Reserved		0x04-0x07

mps_topology_indexed_reference_flag specifies whether topology elements are referenced by unique index. When set to 1, topology elements are represented by their indexes in the topology data defined by the topology payload of type NNR_TPL_REFLIST. If this flag is set to 0, then topology_data of NNR topology unit shall contain the topology information.

base_model_id_present_flag specifies whether a base_model_id is provided. base_model_id_present_flag shall be set to 1, if nnr_pre_flag is 1.

base_model_id provides an identifier for referencing the base model to which a bitstream is related, e.g, to which a weight update is applied. It is up to the application how to use the id.

validation_set_performance_present_flag specifies whether the validation_set_performance is present. When set to 1, the validation_set_performance is present.

metric_type_performance_map_valid_flag specifies that the performance_metric_type is valid performance metric for performance maps. When set to 1, the values reported in performance_maps correspond to the definition provided by performance_metric_type.

performance_metric_type specifies which metric has been used to represent the performance of weight updates. It is a null terminated string that can be defined by the application (see [Annex F](#)).

mps_parent_signalling_enabled_flag specifies whether parent node information (i.e. element parent_node_id_present_flag) is signalled at NDU level. If mps_parent_signalling_enabled_flag is equal to 0, parent_node_id_present_flag is not present for any NDU (i.e. the bitstream represents a base neural network). If mps_parent_signalling_enabled_flag is equal to 1, the bitstream represents a differential update of a base neural network.

nnr_pre_flag specifies whether the bitstream contains encoded residuals of weight updates with respect to previous weight updates or contains the encoded weight updates. If set to 1, the bitstream contains encoded residuals of weight updates. If set to 0, the bitstream contains encoded weight updates. The nnr_pre_flag can be 1 only after the decoder has processed at least one update of the same model. Compressed data units of type NNR_PT_BLOCK are always represented as weight updates. Residuals of weights shall not be used in combination with dependent scalar quantization. If nnr_pre_flag is not present, it is inferred to be 0.

mps_qp_density specifies density information of syntax element mps_quantization_parameter in the NNR compressed data units that utilize this model parameter sets.

mps_quantization_parameter specifies the quantization parameter for scalar uniform quantization of parameters of each layer of the neural network for arithmetic coding in the NNR compressed data units that utilize this model parameter set.

sparsification_performance_map() specifies a mapping between different sparsification thresholds and resulting NN inference accuracies. The resulting accuracies are provided separately for different aspects or characteristics of the output of the NN. For a classifier NN, each sparsification threshold is mapped to separate accuracies for each class, in addition to an overall accuracy which considers all classes. Classes are ordered based on the neural network output order, i.e. the order specified during training.

spm_count_thresholds specifies the number of sparsification thresholds. This number shall be non-zero.

sparsification_threshold specifies a list of thresholds where each threshold is applied to the weights of the decoded neural network in order to set the weights to zero. I.e. the weights whose values are less than the threshold are set to zero.

non_zero_ratio specifies a list of non-zero ratio values where each value is the non-zero ratio that is achieved by applying the `sparsification_threshold` to sparsify the weights.

spm_nn_accuracy specifies a list of accuracy values where each value is the overall accuracy of the NN (e.g. classification accuracy by considering all classes) when sparsification using the corresponding threshold in `sparsification_threshold` is applied.

spm_count_classes specifies a list of number of classes where each such number is the number of classes for which separate accuracies are provided for each sparsification thresholds.

spm_class_bitmask specifies a subset of classes for which the accuracies are signalled, when a certain sparsification threshold is applied. The order of bits indicates the indexes of classes, with the most significant bit representing the presence of the smallest indexed class.

spm_nn_class_accuracy specifies a list of lists of class accuracies, where each value is accuracy for a certain class, when a certain sparsification threshold is applied.

pruning_performance_map() specifies a mapping between different pruning ratios and resulting NN inference accuracies. The resulting accuracies are provided separately for different aspects or characteristics of the output of the NN. For a classifier NN, each pruning ratio is mapped to separate accuracies for each class, in addition to an overall accuracy which considers all classes. Classes are ordered based on the neural network output order, i.e. the order specified during training.

ppm_count_pruning_ratios specifies the number of pruning ratios. This number shall be non-zero.

pruning_ratio specifies the pruning ratio.

ppm_nn_accuracy specifies a list of accuracy values where each value is the overall accuracy of the NN (e.g. classification accuracy by considering all classes) when pruning using the corresponding ratio in `pruning_ratio` is applied.

ppm_class_bitmask specifies a subset of classes for which corresponding accuracies are signalled, when a certain pruning ratio is applied. The order of bits indicates the indexes of classes, with the most significant bit representing the presence of the smallest indexed class.

ppm_count_classes specifies a list of number of classes where each such number is the number of classes for which separate accuracies are provided for each pruning ratio.

ppm_nn_class_accuracy specifies a list of lists of class accuracies, where each value is accuracy for a certain class, when a certain pruning ratio is applied.

unification_performance_map() specifies a mapping between different unification thresholds and resulting NN inference accuracies. The resulting accuracies are provided separately for different aspects or characteristics of the output of the NN. For a classifier NN, each unification threshold is mapped to separate accuracies for each class, in addition to an overall accuracy which considers all classes. Classes are ordered based on the neural network output order, i.e. the order specified during training.

upm_count_thresholds specifies the number of unification thresholds. This number shall be non-zero.

count_resaped_tensor_dimensions specifies a counter of how many dimensions are specified for reshaped tensor. For example, for a weight tensor reshaped to 3-dimensional tensor, count_resaped_tensor_dimensions is 3.

reshaped_tensor_dimensions specifies an array or list of dimension values. For example, for a convolutional layer reshaped to 3-dimensional tensor, dim is an array or list of length 3.

count_super_block_dimensions specifies a counter of how many dimensions are specified. For example, for a 3-dimensional superblock, count_super_block_dimensions is 3.

super_block_dimensions specifies an array or list of dimension values. For example, for a 3-dimensional superblock, dim is an array or list of length 3, i.e. [64, 64, kernel_size].

count_block_dimensions specifies a counter of how many dimensions are specified. For example, for a 3-dimensional block, count_block_dimensions is 3.

block_dimensions specifies an array or list of dimension values. For example, for a 3-dimensional block, dim is an array or list of length 3, i.e. [2, 2, 2].

unification_threshold specifies the threshold which is applied to tensor block in order to unify the absolute value of weights in this tensor block.

upm_nn_accuracy specifies the overall accuracy of the NN (e.g. classification accuracy by considering all classes).

upm_count_classes specifies number of classes for which separate accuracies are provided for each unification thresholds.

upm_class_bitmask specifies a subset of classes for which corresponding accuracies are signalled, when a certain unification threshold is applied. The order of bits indicates the indexes of classes, with the most significant bit representing the presence of the smallest indexed class.

upm_nn_class_accuracy specifies the accuracy for a certain class, when a certain unification threshold is applied.

decomposition_performance_map() specifies a mapping between different mean square error (MSE) thresholds between the decomposed tensors and their original version and resulting NN inference accuracies. The resulting accuracies are provided separately for different aspects or characteristics of the output of the NN. For a classifier NN, each MSE threshold is mapped to separate accuracies for each class, in addition to an overall accuracy which considers all classes. Classes are ordered based on the neural network output order, i.e. the order specified during training.

dpm_count_thresholds specifies the number of decomposition MSE thresholds. This number shall be non-zero.

mse_threshold specifies an array of MSE thresholds which are applied to derive the ranks of the different tensors of weights.

dpm_nn_accuracy specifies the overall accuracy of the NN (e.g. classification accuracy by considering all classes).

nn_reduction_ratio[i] specifies the ratio between the total number of parameters after tensor decomposition of the whole model and the number of parameters in the original model.

dpm_count_classes specifies number of classes for which separate accuracies are provided for each decomposition thresholds.

dpm_nn_class_accuracy specifies an array of accuracies for a certain class, when a certain decomposition threshold is applied.

validation_set_performance specifies a performance indicator obtained on a local validation set to communicate the performance of a weight update after it is dequantized and employed into the base model.

6.4.4.4 NNR layer parameter set unit payload semantics

lps_sparsification_flag specifies whether sparsification was applied to the model in the NNR compressed data units that utilizes this layer parameter set.

lps_pruning_flag specifies whether pruning was applied to the model in the NNR compressed data units that utilizes this layer parameter set.

lps_unification_flag specifies whether unification was applied to the model in the NNR compressed data units that utilizes this layer parameter set.

lps_quantization_method_flags specifies the quantization method used for the data contained in the NNR compressed data units to which this layer parameter set refers. If multiple models are specified, they are combined by OR. The methods are defined in [Table 12](#).

Table 12 — Quantization method identifiers

Quantization method	Quantization method ID	Value
Scalar uniform	NNR_QSU	0x01
Codebook	NNR_QCB	0x02
Reserved		0x04-0x07

lps_qp_density specifies density information of syntax element `lps_quantization_parameter` in the NNR compressed data units that utilize this model parameter set.

lps_quantization_parameter specifies the quantization parameter for scalar uniform quantization of parameters of each layer of the neural network for arithmetic coding in the NNR compressed data units that utilize this model parameter set.

The variable `QpDensity` is derived as follows:

- If an active NNR layer parameter set is present, the variable `QpDensity` is set to `lps_qp_density`.
- Otherwise, the variable `QpDensity` is set to `mps_qp_density`.

The variable `QuantizationParameter` is derived as follows:

- If an active NNR layer parameter set is present, the variable `QuantizationParameter` is set to `lps_quantization_parameter`.
- Otherwise, the variable `QuantizationParameter` is set to `mps_quantization_parameter`.

sparsification_performance_map() is as defined in [subclause 6.4.4.3](#).

When `lps_sparsification_flag` of a certain layer is equal to 1 and `mps_sparsification_flag` is equal to 0, then the information in `sparsification_performance_map()` of the layer parameter set is valid when performing sparsification only on that layer. More than one layer can have `lps_sparsification_flag` equal to 1 in their layer parameter set.

When both `mps_sparsification_flag` and `lps_sparsification_flag` are equal to 1, the following shall apply:

- If sparsification is applied on the whole model (i.e. all layers), then the information in `sparsification_performance_map()` of the model parameter set is valid.
- If sparsification is applied on only one layer, and for that layer `lps_sparsification_flag` is equal to 1, then the information in `sparsification_performance_map()` of the layer parameter set of that layer is valid.

pruning_performance_map() is as defined in [subclause 6.4.4.3](#).

When `lps_pruning_flag` of a certain layer is equal to 1 and `mps_pruning_flag` is equal to 0, then the information in `pruning_performance_map()` of the layer parameter set is valid when performing pruning only on that layer. More than one layer can have `lps_pruning_flag` equal to 1 in their layer parameter set.

When both `mps_pruning_flag` and `lps_pruning_flag` are equal to 1, the following shall apply:

- If pruning is applied on the whole model (i.e. all layers), then the information in `pruning_performance_map()` of the model parameter set is valid.
- If pruning is applied on only one layer, and for that layer `lps_pruning_flag` is equal to 1, then the information in `pruning_performance_map()` of the layer parameter set of that layer is valid.

unification_performance_map() is as defined in [subclause 6.4.4.3](#).

When `lps_unification_flag` of a certain layer is equal to 1 and `mps_unification_flag` is equal to 0, then the information in `unification_performance_map()` of the layer parameter set is valid when performing unification only on that layer. More than one layer can have `lps_unification_flag` equal to 1 in their layer parameter set.

When both `mps_unification_flag` and `lps_unification_flag` are equal to 1, the following shall apply:

- If unification is applied on the whole model (i.e. all layers), then the information in `unification_performance_map()` of the model parameter set is valid.
- If unification is applied on only one layer, and for that layer `lps_unification_flag` is equal to 1, then the information in `unification_performance_map()` of the layer parameter set of that layer is valid.

6.4.4.5 NNR topology unit payload semantics

`topology_storage_format` value is as signalled in the corresponding NNR topology unit header of the same NNR unit of type `NNR_TPL`.

topology_data is a byte sequence of length determined by the NNR unit size describing the neural network topology, in the format specified by `topology_storage_format`.

If `topology_storage_format` is set to `NNR_TPL_UNREC`, definition and identification of the storage format of `topology_data` is out of scope of this document.

NOTE If `topology_storage_format` is set to `NNR_TPL_UNREC`, the (header) structure of `topology_data` can be used to identify the format.

nnr_rep_type specifies whether pruning information is represented as a bitmask or as a dictionary of references of topology elements. The permitted values are specified in [Table 13](#).

Table 13 — Pruning information representation types

nnr_rep_type value	Identifier	Description
0x00	NNR_TPL_BMSK	Topology related information signalled as bitmask
0x01	NNR_TPL_DICT	Topology related information signalled as dictionary of topology elements
0x02-0x03		Reserved

prune_flag when set to 1 indicates that pruning step is used during parameter reduction and pruning related topology information is present in the payload.

order_flag when set to 1 indicates that the bitmask should be processed row-major order; and column-major otherwise.

sparse_flag when set to 1 indicates that sparsification step is used during parameter reduction and related topology information is present in the payload.

count_ids specifies the number of element ids that are updated. When present, its value shall be greater than zero.

element_id specifies the unique id that is used to reference a topology element

element_id_index specifies the unique index of the topology element which is present in the `nnr_topology_unit_payload()` where `topology_storage_format` is equal to `NNR_TPL_REFLIST`.

count_dims specifies the number of dimensions. When present, its value shall be greater than zero.

dim specifies array of dimensions that contain the new dimensions for the specified element. When present, its value shall be greater than zero.

bit_mask_value when set to 1 indicates that this specific neuron's weight is pruned if `pruning_flag` is set to 1 or is sparsified (the weight value is 0) if `sparse_flag` is set to 1.

count_bits specifies the number of bits present in the bit mask information. When present, its value shall be greater than zero.

6.4.4.6 NNR quantization unit payload semantics

quantization_data is a byte sequence of length determined by the NNR unit size describing the neural network quantization information, in the format specified by `quantization_storage_format`.

If `quantization_storage_format` is set to `NNR_QNT_UNREC`, definition and identification of the storage format of `quantization_data` is out of scope of this document.

NOTE If `quantization_storage_format` is set to `NNR_QNT_UNREC`, the (header) structure of `quantization_data` can be used to identify the format.

6.4.4.7 NNR compressed data unit payload semantics

raw_float32_parameter is a float parameter tensor.

6.4.4.8 NNR aggregate unit payload semantics

NNR aggregate unit payload carries multiple NNR units. `num_of_nnr_units_minus2 + 2` parameter in NNR aggregate unit header shall specify how many NNR units are present in the NNR aggregate unit's payload.

7 Decoding process

7.1 General

A decoder that complies with this document shall take an NNR bitstream, as specified in [subclause 6.3](#), as input and

- generate decompressed data which complies with an NNR decompressed data format (as defined in [Table 15](#)) or
- generate ASCII or compressed data outputs as indicated by using the `NNR_TPL` and `NNR_QNT` NNR unit payloads (as described in [subclause 6.3.3](#))

For the decoding process, the following conditions shall apply:

- Any information that is required for decoding an NNR unit of the NNR bitstream should be signalled as part of the NNR bitstream. If such information is not part of the NNR bitstream, then it shall be provided to the decoding process by other means (e.g. out-of-band topology information or parameters required for decoding but not signalled or carried in the NNR bitstream).
- The decoding process shall be initiated with an NNR unit of type `NNR_STR`. With the reception of the `NNR_STR` unit, the decoder shall reset its internal states and get ready to receive an NNR bitstream. The presence and cardinality of preceding NNR units shall be as specified in the relevant clauses and annexes of this document.

NOTE For example, a decoder can be further initialized via an NNR unit of type NNR_MPS in order set global neural network model parameters.

A decoder that complies with this document shall output data structures which comply with the decompressed NNR data formats as soon as it decompresses them. This allows low delay between inputting NNR compressed data units and accessing decompressed data structures from its output. How to establish the relationship between the input NNR units and NNR decompressed output data is out of scope of this document and left to implementation.

7.2 NNR decompressed data formats

Depending on the compression methods used to create a particular bitstream, the NNR decoder is expected to output different decompressed data formats as a result of decoding an NNR data unit. Table 15 specifies the data types for NNR decompressed data formats, and Table 15 specifies the bit depth of the supported data formats that result after decompressing NNR compressed data units. If general_profile_idc has the value 0, nnr_decompressed_data_format shall only take the value 0 or 1.

Table 14 — Data types for NNR decompressed data formats

Parameter identifier	Parameter description
TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model
TENSOR_FLOAT	Tensor of float values used for representing tensor-shaped float parameters of the model

Table 15 — NNR decompressed data formats

nnr_decompressed_data_format	Data type	Format description	Bit depth
0	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	32
1	TENSOR_FLOAT	Tensor of float values used for representing tensor-shaped float parameters of the model	32
2	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	2
3	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	3
4	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	4
5	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	8
6	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	16
7	TENSOR_INT	Tensor of integer values used for representing tensor-shaped signed integer parameters of the model	64
8	TENSOR_FLOAT	Tensor of float values used for representing tensor-shaped float parameters of the model	16
9	TENSOR_FLOAT	Tensor of float values used for representing tensor-shaped float parameters of the model	64

The value range for TENSOR_INT with bit depth n is determined as $-2^{n-1}..2^{n-1}-1$.

7.3 Decoding methods

7.3.1 General

This subclause specifies the decoding methods of this document. Depending on the value of `nnr_compressed_data_unit_payload_type`, one of the subclauses as specified in [Table 16](#) is invoked.

Table 16 — NNR compressed data payload types

Payload identifier	Description	nnr_compressed_data_unit_payload_type	Sub-clause
NNR_PT_INT	integer parameter tensor	0	7.3.2
NNR_PT_FLOAT	float parameter tensor	1	7.3.3
NNR_PT_RAW_FLOAT	uncompressed float parameter tensor	2	7.3.4
NNR_PT_BLOCK	float parameter tensors including a (optionally decomposed) weight tensor and, optionally, local scaling parameters, biases, and batch norm parameters that form a block in the model architecture	3	7.3.5

If the payload identifier is `NNR_PT_INT`, `NNR_PT_FLOAT`, or `NNR_PT_FLOAT_RAW` and if multiple topology elements are combined (as signalled in the NNR compressed data unit header via `nnr_multiple_topology_elements_present_flag`), then NNR decompressed tensors shall be further split into multiple tensors after the decoding process as follows:

- Tensor `RecParam` is split into multiple tensors by invoking `TensorSplit( RecParam, split_index, concatenation_axis_index )`.
- The output of function `TensorSplit` is the list of split output tensors associated with topology elements as specified by array `topology_elem_id_list`.
- Output tensors are further processed by swapping their axis as signalled in `topology_tensor_dimension_mapping()` by invoking `AxisSwap()`.

7.3.2 Decoding method for NNR compressed payloads of type NNR_PT_INT

Input to this process are:

- One or more NNR compressed data units which are marked to be decompressed together by `partial_data_counter` and `nnr_compressed_data_unit_payload_type` fields are set as `NNR_PT_INT`.

Output of this process is a variable `RecParam` of type `TENSOR_INT` as specified in [Table 14](#). The dimensions of `RecParam` are equal to `ShiftArrayIndex( TensorDimensions, first_tensor_dimension_shift )`. Decoding of a bitstream conforming to method `NNR_PT_INT` shall only produce values for `RecParam` that can be represented in two's complement representation:

- if `general_profile_idc` equals 0, as 32bit integer value,
- if `general_profile_idc` equals 1, and `nnr_decompressed_data_format_present_flag` equals 0, as 32bit integer value,
- if `general_profile_idc` equals 1, and `nnr_decompressed_data_format_present_flag` equals 1, as integer with the bit depth defined in [Table 15](#) according to the value of `nnr_decompressed_data_format`.

The arithmetic coding engine and context models are initialized as specified in [subclause 10.3.2](#).

A syntax structure `shift_parameter_ids( cabac_unary_length_minus1, -1 )` according to [subclause 10.2.1.6](#) is decoded from the bitstream and the initialization process for probability estimation parameters as specified in [subclause 10.3.2.2](#) is invoked.

A syntax structure `quant_tensor( TensorDimensions, cabac_unary_length_minus1, 0, -1 )` according to [subclause 10.2.1.4](#) is decoded from the bitstream and `RecParam` is set equal to `QuantParam`.

A syntax structure `terminate_cabac()` according to [subclause 10.2.1.2](#) is decoded from the bitstream.

The return value of `DimensionShift( RecParam, TensorDimensions, first_tensor_dimension_shift )` is assigned to `RecParam`.

If `nnr_pre_flag` equals 1, `RecParam` contains residuals of weight updates with respect to the previous weight update. To retrieve the current weight update, the previous weight update is fetched from a local cache `WeightUpdateCache` which always contains the previous update and is added to the `RecParam`. The return value is assigned to `RecParam` and is stored locally into `WeightUpdateCache`. `WeightUpdateCache` is a list of `TENSOR_INT` objects with length 1. If `nnr_pre_flag` equals 0, `RecParam` already contains weight updates and is stored locally into `WeightUpdateCache`.

7.3.3 Decoding method for NNR compressed payloads of type NNR_PT_FLOAT

Input to this process are:

- One or more NNR compressed data units which are marked to be decompressed together by `partial_data_counter` and their `nnr_compressed_data_unit_payload_type` fields are set as `NNR_PT_FLOAT`.

Output of this process is a variable `RecParam` of type `TENSOR_FLOAT` as specified in [Table 14](#):

- if `general_profile_idc` equals 0, as 32bit floating point value,
- if `general_profile_idc` equals 1, and `nnr_decompressed_data_format_present_flag` equals 0, as 32bit floating point value,
- if `general_profile_idc` equals 1, and `nnr_decompressed_data_format_present_flag` equals 1, as floating point value with the bit depth defined in [Table 15](#) according to the value of `nnr_decompressed_data_format`.

The dimensions of `RecParam` are equal to `ShiftArrayIndex( TensorDimensions, first_tensor_dimension_shift )`.

The arithmetic coding engine and context models are initialized as specified in [subclause 10.3.2](#).

[Subclause 7.3.6](#) is invoked with `TensorDimensions`, 0, and `(codebook_present_flag ? 0 : -1)` as inputs, and the output is assigned to `RecParam`.

A syntax structure `terminate_cabac()` according to [subclause 10.2.1.2](#) is decoded from the bitstream.

Decoding of a bitstream conforming to method `NNR_PT_FLOAT` shall only produce values for `RecParam` that can be represented as float value without loss of precision.

The return value of `DimensionShift( RecParam, TensorDimensions, first_tensor_dimension_shift )` is assigned to `RecParam`.

If `nnr_pre_flag` equals 1, `RecParam` contains residuals of weight updates with respect to the previous weight update. To retrieve the current weight update, the previous weight update is fetched from a local cache `WeightUpdateCache` which always contains the previous update and is added to the `RecParam`. The return value is assigned to `RecParam` and is stored locally into `WeightUpdateCache`. `WeightUpdateCache` is a list of `TENSOR_FLOAT` objects with length 1. If `nnr_pre_flag` equals 0, `RecParam` already contains weight updates and is stored locally into `WeightUpdateCache`.

7.3.4 Decoding method for NNR compressed payloads of type NNR_PT_RAW_FLOAT

Output of this process is a variable `RecParam` of type `TENSOR_FLOAT` as specified in [Table 14](#) The dimensions of `RecParam` are equal to `ShiftArrayIndex( TensorDimensions, first_tensor_dimension_shift )`.

RecParam is set equal to raw_float32_parameter.

The return value of DimensionShift(RecParam, TensorDimensions, first_tensor_dimension_shift) is assigned to RecParam.

7.3.5 Decoding method for NNR compressed payloads of type NNR_PT_BLOCK

Inputs to this process are:

- One or more NNR compressed data units which are marked to be decompressed together by partial_data_counter and their nnr_compressed_data_unit_payload_type fields are set as NNR_PT_BLOCK.

Output of this process are one or more variables of type TENSOR_FLOAT as specified in [Table 14](#) depending on the value of compressed_parameter_types as follows:

If (compressed_parameter_types & NNR_CPT_DC) == 0: RecWeight

If (compressed_parameter_types & NNR_CPT_DC) != 0: RecWeightG, RecWeightH

If (compressed_parameter_types & NNR_CPT_LS) != 0: RecLS

If (compressed_parameter_types & NNR_CPT_BN) != 0: RecBeta, RecGamma, RecMean, RecVar

If (compressed_parameter_types & NNR_CPT_BI) != 0: RecBias

The resulting variables are represented with following bit depth:

- if general_profile_idc equals 0, as 32bit floating point value,
- if general_profile_idc equals 1, and nnr_decompressed_data_format_present_flag equals 0, as 32bit floating point value,
- if general_profile_idc equals 1, and nnr_decompressed_data_format_present_flag equals 1, as floating point value with the bit depth defined in [Table 15](#) according to the value of nnr_decompressed_data_format.

If present, the dimensions of RecWeight are equal to ShiftArrayIndex(TensorDimensions, first_tensor_dimension_shift).

If present, the dimensions of RecWeightG are equal to TensorDimensionsG.

If present, the dimensions of RecWeightH are equal to TensorDimensionsH.

If present, the variables RecLS, RecBeta, RecGamma, RecMean, RecVar, and RecBias are 1D and their length is equal to the first dimension of TensorDimensions.

The arithmetic coding engine and context models are initialized as specified in [subclause 10.3.2](#).

If (compressed_parameter_types & NNR_CPT_LS) != 0, [subclause 7.3.6](#) is invoked with the dimensions of RecLS, -1, and -1 as inputs, and the output is assigned to RecLS.

If (compressed_parameter_types & NNR_CPT_BI) != 0, [subclause 7.3.6](#) is invoked with the dimensions of RecBias, -1, and -1 as inputs, and the output is assigned to RecBias.

If (compressed_parameter_types & NNR_CPT_BN) != 0, [subclause 7.3.6](#) is invoked with the dimensions of RecBeta, -1, and -1 as inputs, and the output is assigned to RecBeta.

If (compressed_parameter_types & NNR_CPT_BN) != 0, [subclause 7.3.6](#) is invoked with the dimensions of RecGamma, -1, and -1 as inputs, and the output is assigned to RecGamma.

If (compressed_parameter_types & NNR_CPT_BN) != 0, [subclause 7.3.6](#) is invoked with the dimensions of RecMean, -1, and -1 as inputs, and the output is assigned to RecMean.

If (compressed_parameter_types & NNR_CPT_BN) != 0, [subclause 7.3.6](#) is invoked with the dimensions of RecVar, -1, and -1 as inputs, and the output is assigned to RecVar.

If (compressed_parameter_types & NNR_CPT_DC) == 0, [subclause 7.3.6](#) is invoked with the dimensions of RecWeight, 0, and (codebook_present_flag ? 0 : -1) as inputs, and the output is assigned to RecWeight.

If (compressed_parameter_types & NNR_CPT_DC) != 0, the following applies:

[Subclause 7.3.6](#) is invoked with TensorDimensionsG, 0, and (codebook_present_flag ? 0 : -1) as inputs, and the output is assigned to RecWeightG.

[Subclause 7.3.6](#) is invoked with TensorDimensionsH, (TensorDimensionsG[0] + (4 << scan_order) - 1) >> (2 + scan_order) - 1, and (codebook_present_flag ? 1 : -1) as inputs, and the output is assigned to RecWeightH.

NOTE From the decoded RecWeightG and RecWeightH, the variable RecWeight can be derived as follows:

RecWeight = TensorReshape (RecWeightG * RecWeightH, TensorDimensions)

The return value of DimensionShift(RecWeight, TensorDimensions, first_tensor_dimension_shift) is assigned to RecWeight.

A syntax structure terminate_cabac() according to [subclause 10.2.1.2](#) is decoded from the bitstream.

If (compressed_parameter_types & NNR_CPT_DC) == 0, the return value of DimensionShift(RecWeight, TensorDimensions, first_tensor_dimension_shift) is assigned to RecWeight.

7.3.6 Decoding process for an integer weight tensor

Inputs to this process are:

- A variable tensorDims specifying the dimensions of the tensor to be decoded.
- A variable entryPointOffset indicating whether entry points are present for decoding and, if entry points are present, an entry point offset.
- A variable codebookId indicating whether a codebook is applied and, if a codebook is applied, which codebook shall be used.

Output of this process is a variable recParam of type TENSOR_FLOAT as specified in [Table 14](#) with dimensions equal to tensorDims.

A syntax structure quant_param(QpDensity) according to [subclause 10.2.1.3](#) is decoded from the bitstream.

A syntax structure shift_parameter_ids(cabac_unary_length_minus1, codebookId) according to [subclause 10.2.1.6](#) is decoded from the bitstream and the initialization process for probability estimation parameters as specified in [subclause 10.3.2.2](#) is invoked.

A syntax structure quant_tensor(tensorDims, cabac_unary_length_minus1, entryPointOffset, codebookId) according to [subclause 10.2.1.4](#) is decoded from the bitstream and recParam is set as follows:

```

if( codebookId == -1 )
    recParam = QuantParam
else {
    for( i = 0; i < Prod( tensorDims ); i++ ) {
        idx = TensorIndex( tensorDims, i )
        if( codebookId == 0 )
            recParam[idx] = Codebook[ QuantParam[idx] + CbZeroOffset ]
        else
            recParam[idx] = CodebookDC[ QuantParam[idx] + CbZeroOffsetDC ]
    }
}

```

A variable stepSize is derived as follows:

```
mul = (1 << QpDensity) + ( (qp_value + QuantizationParameter) & ( ( 1 << QpDensity ) - 1 ) )
shift = (qp_value + QuantizationParameter) >> QpDensity
stepSize = mul * 2shift - QpDensity
```

Variable recParam is updated as follows:

```
recParam = recParam * stepSize
```

NOTE Following from the above calculations, recParam can always be represented as binary fraction.

8 Parameter reduction

8.1 General

This clause specifies the reconstruction of parameters for specific parameter reduction tools. Additional information about the parameter reduction tools not essential for decoding is provided in Annex [G.1](#).

8.2 Methods

8.2.1 Batchnorm folding

If batchnorm folding is applied, the batchnorm function is represented as

$$BN(X) = \alpha \circ W * X + \delta$$

where $\alpha = \frac{\gamma}{\sqrt{\sigma^2 + \epsilon}}$ and where $\delta = \frac{(b - \mu) \circ \gamma}{\sqrt{\sigma^2 + \epsilon}} + \beta$. X is the input, $BN(X)$ is the output, W is a weight tensor of the convolutional or fully-connected layer (represented as 2D matrix), b is a bias parameter, and where the remaining parameters are batch-normalization parameters. Note that b , μ , σ^2 , γ , and β have the same shape as X and that X is shaped as a transposed vector. Parameter ϵ is a scalar close to zero.

Parameter α can be present in NNR compressed payloads of type NNR_PT_BLOCK as output variable RecLS and it is quantized using either uniform quantization or dependent scalar quantization.

Parameter δ can be present in NNR compressed payloads of type NNR_PT_BLOCK as output variable RecBias and it is quantized using either uniform quantization or dependent scalar quantization.

Note that the four batchnorm parameters RecBeta, RecGamma, RecMean, and RecVar can be recreated from RecLS and RecBias according to the following equations:

$$\text{RecBeta} = \text{RecBias}$$

$$\text{RecGamma} = \text{RecLS}$$

$$\text{RecMean} = 0$$

$$\text{RecVar} = 1 - \epsilon, \text{ where parameter } \epsilon \text{ is a scalar close to zero.}$$

In case the four batchnorm parameters have been recreated, RecBias is set to 0 and RecLS is set to 1. If RecBias and RecLS are not required, they can simply be ignored.

In the context of incremental compression for (federated) training of neural networks, δ (RecBias), α (RecLS), RecMean and RecVar are required to continue the training process. In order to implement batchnorm folding in an incremental scenario, first a local copy of the batchnorm folded network shall be stored, including the calculated δ_0 and α_0 parameters per batchnorm module. Then, one round of training

is executed with the unfolded model. After training, the updated model is batchnorm folded to create the respective δ_1 and α_1 parameters of the updated model.

For compressing the model update, only $\Delta\delta = \delta_1 - \delta_0$ and $\Delta\alpha = \alpha_1 - \alpha_0$ are included in the quantization and coding process whereas μ and σ^2 are skipped. To continue training with the model update,

$$\beta_u = \beta + \Delta\delta$$

$$\gamma_u = \gamma + \Delta\alpha.$$

The mean and variance parameters μ and σ^2 are reset to 0 and 1. Note that during training, the batchnorm module usually keeps running estimates of μ and σ^2 with a momentum. This information is lost due to incremental batchnorm folding.

8.3 Syntax and semantics

8.3.1 Sparsification using compressibility loss

The presence and semantics of syntax elements are specified in [Table 17](#).

Table 17 — Syntax and semantics for sparsification using compressibility loss

Syntax element	condition	semantics
tensor_dimensions	present	Dimension and shape of original tensors

8.3.2 Sparsification using micro-structured pruning

The presence and semantics of syntax elements are specified in [Table 18](#).

Table 18 — Syntax and semantics for sparsification using micro-structured pruning

Syntax element	condition	semantics
count_tensor_dimension	present	counter of how many dimensions of reshaped weight tensor
reshaped_tensor_dimensions[]	present	dimensions of reshaped weight tensor
count_super_block_dimension	present	counter of how many dimensions of superblock
super_block_dimensions[]	present	dimensions of superblock
count_block_dimension	present	counter of how many dimensions of block
block_dimensions[]	present	dimensions of block

8.3.3 Combined pruning and sparsification

The presence and semantics of syntax elements are specified in [Table 19](#).

Table 19 — Syntax and semantics for combined pruning and sparsification

Syntax element and functions	condition	semantics
nnr_rep_type	present	The flag to indicate what type of output is produced
prune_flag	present	The flag to indicate pruning is applied
order_flag	present	The flag to indicate the order of processing of information in row-major or column-major
sparse_flag	present	The flag to indicate sparsification is applied
count_ids	(prune_flag == 1) && (nnr_rep_type == NNR_TPL_DICT)	The number of elements that are pruned
element_id[]	(prune_flag == 1) && (nnr_rep_type == NNR_TPL_DICT)	The IDs of the elements that are pruned
count_dims[]	(prune_flag == 1) && (nnr_rep_type == NNR_TPL_DICT)	The number of dimensions of each pruned element
dim[][]	(prune_flag == 1) && (nnr_rep_type == NNR_TPL_DICT)	The new dimensions of the pruned elements
bit_mask()	sparse_flag == 1	A bitmask to indicate which matrix elements are preserved during sparsification. A bit value of 1 shall indicate that the corresponding element is preserved and a bit value of 0 shall indicate that the corresponding element is sparsified
	(prune_flag == 1) && (nnr_rep_type == NNR_TPL_BMSK)	A bitmask to indicate which matrix elements or output channels are preserved during pruning. A bit value of 1 shall indicate that the corresponding element is preserved and a bit value of 0 shall indicate that the corresponding element is pruned

8.3.4 Unstructured statistics-adaptive sparsification

The semantics of syntax elements are specified in [Table 20](#).

Table 20 — Syntax and semantics for structured sparsification

Syntax element	condition	semantics
qp_value	present	integer
QpDensity	present	integer

8.3.5 Structured sparsification (global and local approach)

The semantics of syntax elements are specified in [Table 21](#).

Table 21 — Syntax and semantics for structured sparsification

Syntax element	condition	semantics
row_skip_enabled_flag	present	indicates whether row skipping is used for the sparsified element
row_skip_list	present, if row_skip_enabled_flag==1	

8.3.6 Weight unification

The presence and semantics of syntax elements are specified in [Table 22](#).

Table 22 — Syntax and semantics for weight unification

Syntax element	condition	semantics
count_tensor_dimension	present	counter of how many dimensions of reshaped weight tensor
reshaped_tensor_dimensions[]	present	dimensions of reshaped weight tensor
count_super_block_dimension	present	counter of how many dimensions of superblock
super_block_dimensions[]	present	dimensions of superblock
count_block_dimension	present	counter of how many dimensions of block
block_dimensions[]	present	dimensions of block

8.3.7 Low rank/low displacement rank for convolutional and fully connected layers

The presence and semantics of syntax elements are specified in [Table 23](#).

Table 23 — Syntax and semantics for low rank/low displacement rank

Syntax element	condition	semantics
compressed_parameter_types	(compressed_parameter_types && NNR_CPT_DC) != 0	One bit indicating whether decomposition is present
decomposition_rank	present	rank
g_number_of_rows	present	rows of G
tensor_dimensions	present	dimensions of original tensor

8.3.8 Batchnorm folding

The presence and semantics of syntax elements are specified in [Table 24](#).

Table 24 — Syntax and semantics for batchnorm folding

Syntax element/Variable	condition	semantics
compressed_parameter_types	(compressed_parameter_types && NNR_CPT_BN) != 0	One bit indicating whether batchnorm parameters are present
QpDensity	present	unsigned integer
QuantizationParameter	present	integer
qp_value	present	integer
dq_flag	present	flag

8.3.9 Local scaling adaptation (LSA)

The presence and semantics of syntax elements are specified in [Table 25](#).

Table 25 — Syntax and semantics for local scaling

Syntax element/Variable	condition	semantics
compressed_parameter_types	(compressed_parameter_types && NNR_CPT_LS) != 0	One bit indicating whether a local scaling parameter is present
QpDensity	present	unsigned integer
QuantizationParameter	present	integer
qp_value	present	integer
dq_flag	present	flag

9 Parameter quantization

9.1 General

This clause specifies the reconstruction of parameters for specific parameter quantization tools. Additional information about the parameter quantization tools not essential for decoding is provided in Annex [G.2](#).

9.2 Methods

9.2.1 Uniform quantization method

Uniform quantization is applied to the parameter tensors using a fixed step size represented by parameters `mps_qp_density` (or `lps_qp_density`, if present) and `qp_value` according to the specification in [subclause 7.3.3](#) and a flag, denoted as `dq_flag`, equal to zero. The reconstructed values in the decoded tensor are integer multiples of the step size.

9.2.2 Codebook-based method

The parameter tensors are represented as a codebook and tensors of indices, the latter having the same shape as the original tensors. The size of the codebook is chosen at the encoder and is transmitted as a metadata parameter. The indices have integer values, they will be further entropy coded. The codebook is composed of integer values that are strictly monotonically increasing.

The reconstructed integer tensors are the values of codebook elements referred to by their index value and the reconstructed tensors are derived by multiplying the reconstructed integer tensors with a step size that is derived from parameters `mps_qp_density` (or `lps_qp_density`, if present) and `qp_value`.

9.2.3 Dependent scalar quantization method

Dependent scalar quantization is applied to the parameter tensors using a fixed stepsize represented by parameters `mps_qp_density` (or `lps_qp_density`, if present) and `qp_value` according to the specification in [subclause 7.3.3](#) and a state transition table of size 8, whenever a flag, denoted as `dq_flag`, is equal to one. The reconstructed values in the decoded tensor are integer multiples of the step size.

9.2.4 Predictive residual encoding (PRE)

PRE aims at minimizing the amount of data transferred between subsequent updates. PRE obtains this minimization goal by encoding and communicating the residual of the weight updates with respect to a base signal instead of the actual weight updates whenever using the residual is beneficial. Residuals are computed as the subtraction of the actual weight updates and the base signal on MPS level. Each decoder instance adopts its own previous weight update as the base signal for computing the current residual. The residuals and the weight updates are then entropy coded separately and encoder decides based on the bitstream sizes which one is to be communicated.

When PRE is activated, residual calculation is done for all quantized parameters except for quantized parameters in `NNR_PT_BLOCK`.

9.3 Syntax and semantics

9.3.1 Uniform quantization method

The presence and semantics of syntax elements are specified in [Table 26](#).

Table 26 — Syntax and semantics for uniform quantization method

Syntax element/Variable	condition	semantics
QpDensity	present	unsigned integer
QuantizationParameter	present	integer
qp_value	present	integer
dq_flag	dq_flag == 0	flag

9.3.2 Codebook-based method

The presence and semantics of syntax elements are specified in [Table 27](#).

Table 27 — Syntax and semantics for codebook-based method

Syntax element/Variable	condition	semantics
QpDensity	present	unsigned integer
QuantizationParameter	present	integer
qp_value	present	integer
codebook_egk	present	unsigned integer
codebook_size	present	unsigned integer
codebook_centre_offset	present	integer
codebook_zero_value	present	integer
codebook_delta_left	present	unsigned integer, multiple instances thereof
codebook_delta_right	present	unsigned integer, multiple instances thereof

9.3.3 Dependent scalar quantization method

The presence and semantics of syntax elements are specified in [Table 28](#).

Table 28 — Syntax and semantics for dependent scalar quantization method

Syntax element	condition	semantics
QpDensity	present	unsigned integer
QuantizationParameter	present	integer
qp_value	present	integer
dq_flag	dq_flag == 1	flag

10 Entropy coding

10.1 Methods

10.1.1 DeepCABAC

10.1.1.1 Binarization

The encoding method scans the parameter tensor in a manner as defined by function `TensorIndex()`. Each quantized parameter level is encoded according to the following procedure employing an integer parameter `maxNumNoRemMinus1`:

In the first step, a binary syntax element `sig_flag` is encoded for the quantized parameter level, which specifies whether the corresponding level is equal to zero. If the `sig_flag` is equal to one, a further binary syntax element `sign_flag` is encoded. The bin indicates if the current parameter level is positive or negative. Next, a unary sequence of bins is encoded, followed by a fixed length sequence as follows:

A variable k is initialized with zero and X is initialized with $1 \leq k$. A syntax element $\text{abs_level_greater_x/x2}$ is encoded, which indicates, that the absolute value of the quantized parameter level is greater than x . If $\text{abs_level_greater_x/x2}$ is equal to 1 and if x is greater than maxNumNoRemMinus1 , the variable k is increased by 1. Afterwards, $1 \leq k$ is added to x and a further $\text{abs_level_greater_x/x2}$ is encoded. This procedure is continued until an $\text{abs_level_greater_x/x2}$ is equal to 0. Now, it is clear that X takes one of the values $(x, x - 1, \dots, X - (1 \leq k) + 1)$. A code of length k is encoded, which points to the values in the list which is absolute quantized parameter level.

10.1.1.2 Row skipping

If enabled by flag $\text{row_skip_flag_enabled_flag}$, the row skipping technique signals one flag $\text{row_skip_list}[i]$ for each value i along the first axis of the parameter tensor. If the flag $\text{row_skip_list}[i]$ is 1, all elements of the parameter tensor for which the index for the first axis equals i are set to zero. If the flag $\text{row_skip_list}[i]$ is 0, all elements of the parameter tensor for which the index for the first axis equals i are encoded individually.

10.1.1.3 Context modeling

Context modeling corresponds to associating the three type of flags sig_flag , sign_flag , and $\text{abs_level_greater_x/x2}$ with context models. In this way, flags with similar statistical behavior should be associated with the same context model so that the probability estimator (inside of the context model) can adapt to the underlying statistics.

The context modeling is as follows:

Twenty-four context models are distinguished for the sig_flag , depending on the state value and whether the neighbouring quantized parameter level to the left is zero, smaller, or larger than zero.

If dq_flag is 0, only the first three context models are used.

Three other context models are distinguished for the sign_flag depending on whether the neighbouring quantized parameter level to the left is zero, smaller, or larger than zero.

For the $\text{abs_level_greater_x/x2}$ flags, each x uses either one or two separate context models. If $x \leq \text{maxNumNoRemMinus1}$, two context models are distinguished depending on the sign_flag . If $x > \text{maxNumNoRemMinus1}$, only one context model is used.

10.1.1.4 Temporal context modeling

If enabled by flag $\text{temporal_context_modeling_flag}$, additional context model sets for flags sig_flag , sign_flag and $\text{abs_level_greater_x}$ are available. The derivation of ctxIdx is then also based on the value of a quantized co-located parameter level in the previously encoded parameter update tensor, which can be uniquely identified using the device_id , parameter_id and put_node_depth elements defined in [subclause 6.3.3.7](#). If the co-located parameter level is not available or equal to zero, the context modeling according to [subclause 10.1.1.3](#) is applied. Otherwise, if the co-located parameter level is not equal to zero, the temporal context modeling of the presented approach is as follows:

Sixteen context models are distinguished for the sig_flag , depending on the state value and whether the absolute value of the quantized co-located parameter level is greater than one or not.

If dq_flag is 0, only the first two additional context models are used.

Two more context models are distinguished for the sign_flag depending on whether the quantized co-located parameter level is smaller or greater than zero.

For the $\text{abs_level_greater_x}$ flags, each x uses two separate context models. These two context models are distinguished depending on whether the absolute value of the quantized co-located parameter level is greater or equal to $x-1$ or not.

10.2 Syntax and semantics

10.2.1 DeepCABAC syntax

10.2.1.1 General

This subclause specifies the entropy coding syntax as used by the decoding process of [Clause 7](#).

10.2.1.2 DeepCABAC termination syntax

terminate_cabac() {	Descriptor
terminating_one_bit	at(v)
while(!byte_aligned())	
nesting_zero_bit	f(1)
}	

terminating_one_bit specifies a terminating bit equal to 1.

nesting_zero_bit is one bit set to 0.

10.2.1.3 Quantization parameter syntax

quant_param(qpDensity) {	Descriptor
qp_value	iae(6 + qpDensity)
}	

qp_value is the quantization parameter.

10.2.1.4 Quantized tensor syntax

quant_tensor(dimensions, maxNumNoRemMinus1, entryPointOffset, codebookId) {	Descriptor
if(general_profile_idc == 1 && codebookId > -1){	
if(codebookId == 0) {	
codebookSize = CbSize	
cbZeroOffset = CbZeroOffset	
}	
else {	
codebookSize = CbSizeDC	
cbZeroOffset = CbZeroOffsetDC	
}	
}	
else {	
codebookSize = 0	
cbZeroOffset = 0	
}	
tensor2DHeight = dimensions[0]	
tensor2DWidth = Prod(dimensions) / tensor2DHeight	
if(general_profile_idc == 1 && tensor2DWidth > 1 && !(codebookId > -1 && codebookSize == 1)) {	
if(parent_node_id_present_flag == 1) {	

hist_dep_sig_prob_enabled_flag	uae(1)
}	
row_skip_enabled_flag	uae(1)
if(row_skip_enabled_flag)	
for(i = 0; i < tensor2DHeight; i++)	
row_skip_list[i]	ae(v)
}	
stateId = 0	
bitPointer = get_bit_pointer()	
lastOffset = 0	
for(i = 0; i < Prod(dimensions); i++) {	
idx = TensorIndex(dimensions, i, scan_order)	
if(entryPointOffset != -1 && GetEntryPointIdx(dimensions, i, scan_order) != -1 && scan_order > 0) {	
lvlCurrRange = 256	
j = entryPointOffset + GetEntryPointIdx(dimensions, i, scan_order)	
lvlOffset = cabac_offset_list[j]	
if(dq_flag)	
stateId = dq_state_list[j]	
set_bit_pointer(bitPointer + lastOffset + BitOffsetList[j])	
lastOffset += BitOffsetList[j]	
init_prob_est_param()	
}	
QuantParam[idx] = 0	
if(general_profile_idc != 1 tensor2DWidth <= 1 !row_skip_enabled_flag !row_skip_list[idx[0]])	
int_param(idx, maxNumNoRemMinus1, stateId, codebookSize, cbZeroOffset)	10.2.1.5
if(dq_flag) {	
nextSt = StateTransTab[stateId][QuantParam[idx] & 1]	
if(QuantParam[idx] != 0) {	
QuantParam[idx] = QuantParam[idx] << 1	
if(QuantParam[idx] < 0)	
QuantParam[idx] += stateId & 1	
else	
QuantParam[idx] += - (stateId & 1)	
}	
stateId = nextSt	
}	
}	
}	

hist_dep_sig_prob_enabled_flag specifies whether history dependent significance probability modeling is enabled. A **hist_dep_sig_prob_enabled_flag** equal to 1 specifies that history dependent significance

probability modeling is enabled. When not present, the value of `hist_dep_sig_prob_enabled_flag` is inferred to be zero.

row_skip_enabled_flag specifies whether row skipping is enabled. A `row_skip_enabled_flag` equal to 1 indicates that row skipping is enabled. If `row_skip_enabled_flag` is not present in the bitstream, it is inferred to be zero.

row_skip_list specifies a list of flags where the *i*-th flag `row_skip_list[ i ]` indicates whether all tensor elements of `QuantParam` for which the index for the first dimension equals *i* are zero. If `row_skip_list[ i ]` is equal to 1, all tensor elements of `QuantParam` for which the index for the first dimension equals *i* are zero. `QuantParam` is an array holding the decoded integer values of the tensor.

`init_prob_est_param()` invokes the initialization process specified in [subclause 10.3.2.2](#).

The 2D integer array `StateTransTab[][]` specifies the state transition table for dependent scalar quantization and is as follows:

`StateTransTab[][] = { {0, 2}, {7, 5}, {1, 3}, {6, 4}, {2, 0}, {5, 7}, {3, 1}, {4, 6} }`

10.2.1.5 Quantized parameter syntax

<code>int_param(i, maxNumNoRemMinus1, stateId, codebookSize, cbZeroOffset) {</code>	Descriptor
<code>if(codebookSize != 1 general_profile_idc != 1) {</code>	
<code> sig_flag</code>	<code>ae(v)</code>
<code> if(sig_flag) {</code>	
<code> QuantParam[i]++</code>	
<code> if (general_profile_idc == 1 && codebookSize > 0 && (cbZeroOffset</code>	
<code> == 0 cbZeroOffset == codebookSize-1)) {</code>	
<code> signVal = cbZeroOffset != 0 ? 1 : 0</code>	
<code> }</code>	
<code> } else {</code>	
<code> sign_flag</code>	<code>ae(v)</code>
<code> signVal = sign_flag</code>	
<code> }</code>	
<code> j = -1</code>	
<code> if(general_profile_idc == 1 && codebookSize > 0)</code>	
<code> maxAbsVal = signVal ? cbZeroOffset : (codebookSize -</code>	
<code> cbZeroOffset - 1)</code>	
<code> } else</code>	
<code> maxAbsVal = -1</code>	
<code> if(maxAbsVal > 1 maxAbsVal == -1) {</code>	
<code> do {</code>	
<code> j++</code>	
<code> abs_level_greater_x[j]</code>	<code>ae(v)</code>
<code> QuantParam[i] += abs_level_greater_x[j]</code>	
<code> } while(abs_level_greater_x[j] == 1 && j <</code>	
<code> maxNumNoRemMinus1 && (maxAbsVal > QuantParam[i] </code>	
<code> maxAbsVal = -1))</code>	
<code> if(abs_level_greater_x[j] == 1 && (maxAbsVal ></code>	
<code> QuantParam[i] maxAbsVal = -1)) {</code>	
<code> RemBits = 0</code>	

j = -1	
do {	
j++	
abs_level_greater_x2 [j]	ae(v)
if(abs_level_greater_x2[j] && (maxAbsVal > QuantParam[i] maxAbsVal = -1)) {	
QuantParam[i] += 1 << RemBits	
RemBits++	
}	
} while(abs_level_greater_x2[j] && j < 30 && (maxAbsVal ≥ (QuantParam[i] + (1 << RemBits)) maxAbsVal == -1))	
abs_remainder	uae(RemBits)
QuantParam[i] += abs_remainder	
}	
}	
QuantParam[i] = signVal ? -QuantParam[i] : QuantParam[i]	
}	
}	
}	

sig_flag specifies whether the quantized weight QuantParam[i] is nonzero. A sig_flag equal to 0 indicates that QuantParam[i] is zero.

sign_flag specifies whether the quantized weight QuantParam[i] is positive or negative. A sign_flag equal to 1 indicates that QuantParam[i] is negative.

abs_level_greater_x[j] indicates whether the absolute level of QuantParam[i] is greater j + 1.

abs_level_greater_x2[j] comprises the unary part of the exponential Golomb remainder.

abs_remainder indicates a fixed length remainder.

The variable curParaId is set equal to the parameter identifier of the currently decoded parameter. QuantParam[i]. When no parameter with a parameter identifier equal to curParaId has been decoded before, the variable AnySigBeforeFlag[curParaId] is set equal to 0. The variable AnySigBeforeFlag[curParaId] is modified as follows:

AnySigBeforeFlag[curParaId] = AnySigBeforeFlag[curParaId] || (QuantParam[i] != 0)

10.2.1.6 Shift parameter indices syntax

shift_parameter_ids(maxNumNoRemMinus1, codebookId) {	Descriptor
if(general_profile_idc == 1 && codebookId > -1) {	
if(codebookId == 0) {	
codebookSize = CbSize	
cbZeroOffset = CbZeroOffset	
}	
else {	
codebookSize = CbSizeDC	

cbZeroOffset = CbZeroOffsetDC	
}	
}	
else {	
codebookSize = 0	
cbZeroOffset = 0	
}	
for(i = 0; i < (dq_flag ? 24 : 3; i++) {	
if(codebookSize != 1) {	
shift_idx(i, ShiftParameterIdsSigFlag)	10.2.1.7
}	
else {	
ShiftParameterIdsSigFlag[i] = 0	
}	
}	
if(temporal_context_modeling_flag){	
for(i = 24; i < (dq_flag ? 40 : 26); i++) {	
if(codebookSize != 1) {	
shift_idx(i, ShiftParameterIdsSignFlag)	
}	
else {	
ShiftParameterIdsSigFlag[i] = 0	
}	
}	
if(hist_dep_sig_prob_enabled_flag) {	
for(i = 40; i < (dq_flag ? 48 : 41); i++) {	
if(codebookSize != 1) {	
shift_idx(i, ShiftParameterIdsSigFlag)	
}	
else {	
ShiftParameterIdsSigFlag[i] = 0	
}	
}	
for(i = 0; i < (temporal_context_modeling_flag ? 5:3); i++) {	
if(!(codebookSize > 0 && (cbZeroOffset == 0 cbZeroOffset == codebookSize -1)) && codebookSize != 1) {	
shift_idx(i, ShiftParameterIdsSignFlag)	
}	
else {	
ShiftParameterIdsSignFlag[i] = 0	
}	
}	

maxAbsVal = codebookSize > 0 ? max(cbZeroOffset, codebookSize -1 - cbZeroOffset) : -1	
for(i = 0; i < (: 2)*(maxNumNoRemMinus1+1); i++) {	
if(maxAbsVal == -1 i < (maxAbsVal-1)*2) {	
shift_idx(i, ShiftParameterIdsAbsGrX)	
}	
else {	
ShiftParameterIdsAbsGrX[i] = 0	
}	
}	
if(temporal_context_modeling_flag) {	
for(i = 2*(maxNumNoRemMinus1+1); i < 4*(maxNumNoRemMinus1+1); i++) {	
if(maxAbsVal == -1 i < (maxAbsVal-1)*2 + (	
2*(maxNumNoRemMinus1+1))) {	
shift_idx(i, ShiftParameterIdsAbsGrX)	
}	
else {	
ShiftParameterIdsAbsGrX[i] = 0	
}	
}	
}	
currX2Level= maxNumNoRemMinus1	
for(i = 0; i < 31; i++) {	
currX2Level = currX2Level + (1 << i)	
if(maxAbsVal == -1 currX2Level < maxAbsVal) {	
shift_idx(i, ShiftParameterIdsAbsGrX2)	
}	
else {	
ShiftParameterIdsAbsGrX2[i] = 0	
}	
}	
}	

10.2.1.7 Shift parameter syntax

	Descriptor
shift_idx(ctxId, shiftParameterIds) {	
shiftParameterIds[ctxId] = 0	
shift_idx_minus_1_present_flag	ae(v)
if(shift_idx_minus_1_present_flag) {	
shift_idx_minus_1	uae(3)
shiftParameterIds[ctxId] += shift_idx_minus_1 + 1	
}	
}	

shift_idx_minus_1_present_flag specifies whether the shift parameter index shiftParameterIds[ctxId] is present. A shift_idx_minus_1_present_flag equal to zero indicates that shiftParameterIds[ctxId] is zero.

shift_idx_minus_1 specifies the absolute value of the shift parameter index `shiftParameterIds[ ctxId ]` minus one. The shift parameter index is `shiftParameterIds[ ctxId ] = shift_idx_minus_1 + 1`

10.3 Entropy decoding process

10.3.1 General

Inputs to this process are a request for a value of a syntax element and values of prior parsed syntax elements.

Output of this process is the value of the syntax element.

The parsing of syntax elements proceeds as follows:

For each requested value of a syntax element a binarization is derived as specified in [subclause 10.3.3](#).

The binarization for the syntax element and the sequence of parsed bins determines the decoding process flow as described in [subclause 10.3.4](#).

10.3.2 Initialization process

10.3.2.1 General

Outputs of this process are initialized DeepCABAC internal variables.

The context variables of the arithmetic decoding engine are initialized as follows:

The initialization process for context variables is invoked as specified in [subclause 10.3.2.3](#).

The decoding engine registers `IvlCurrRange` and `IvlOffset` both in 16 bit register precision are initialized by invoking the initialization process for the arithmetic decoding engine as specified in [subclause 10.3.2.4](#).

10.3.2.2 Initialization process for probability estimation parameters

Outputs of this process are the initialized probability estimation parameters `shift0`, `shift1`, `pStateIdx0`, and `pStateIdx1` for each context model of syntax elements `sig_flag`, `sign_flag`, `abs_level_greater_x`, and `abs_level_greater_x2`.

The 2D array `CtxParameterList [ ] [ ]` is initialized as follows:

`CtxParameterList [ ] [ ] = { { 1, 4, 0, 0 }, { 1, 4, -41, -654 }, { 1, 4, 95, 1519 }, { 0, 5, 0, 0 }, { 2, 6, 30, 482 }, { 2, 6, 95, 1519 }, { 2, 6, -21, -337 }, { 3, 5, 0, 0 }, { 3, 5, 30, 482 } }`

When `hist_dep_sig_prob_enabled_flag` is equal to 1, the following applies for `i` in the range of 40 to `(dq_flag ? 47 : 40)`, inclusive:

- The context parameters associated with the i -th context model of the `sig_flag` are set as follows with `setId` equal to `ShiftParameterIdsSigFlag[ i ]`:
 - `shift0` is set to `CtxParameterList[ setId ][ 0 ]`
 - `shift1` is set to `CtxParameterList[ setId ][ 1 ]`,
 - `pStateIdx0` is set to `CtxParameterList[ setId ][ 2 ]`
 - `pStateIdx1` is set to `CtxParameterList[ setId ][ 3 ]`

The list `idcsSig` is derived as follows:

- When `hist_dep_sig_prob_enabled_flag` is equal to 1, the values of the range `40..(dq_flag ? 47 : 40)` are added to `idcsSig`.

- When `dq_flag` is equal to 1, the values of the range $0..(\text{temporal_context_modeling_flag} ? 39 : 23)$ are added to `idcsSig`.
- When `dq_flag` is equal to 0 and `temporal_context_modeling_flag` is equal to 1, the values $\{0, 1, 2, 24, 25\}$ are added to `idcsSig`.

The initialization process for a probability estimation parameter range as specified in subclause [10.3.2.2](#) is invoked with `idcs` equal to `idcsSig`, `elementName` equal to `sig_flag`, and `shiftParameterIds` equal to `ShiftParameterIdsSigFlag`.

The initialization process for a probability estimation parameter range as specified in subclause [10.3.2.2](#) is invoked with `idcs` including the values of the range of $0..(\text{temporal_context_modeling_flag} ? 4 : 2)$, `elementName` equal to `sign_flag`, and `shiftParameterIds` equal to `ShiftParameterIdsSignFlag`.

The initialization process for a probability estimation parameter range as specified in subclause [10.3.2.2](#) is invoked with `idcs` including the values of the range of $0..(\text{temporal_context_modeling_flag} ? 4 : 2) * (\text{cabac_unary_length_minus1} + 1) - 1$, `elementName` equal to `abs_level_greater_x`, and `shiftParameterIds` equal to `ShiftParameterIdsAbsGrX`.

The initialization process for a probability estimation parameter range as specified in subclause [10.3.2.2](#) is invoked with `idcs` including the values of the range of $0..30$, `elementName` equal to `abs_level_greater_x2`, and `shiftParameterIds` equal to `ShiftParameterIdsAbsGrX2`. To initialize a probability estimation parameter range, the input are the list `idcs`, the syntax element name `elementName`, and the shift parameters `shiftParameterIds`.

The following applies for each entry *i* of the list `idcs`:

- The context parameters associated with the *i*-th context model of the syntax element `elementName` are set as follows with `setId` equal to `shiftParameterIds[i]`:
 - `shift0` is set to `CtxParameterList[setId][0]`
 - `shift1` is set to `CtxParameterList[setId][1]`
 - `pStateIdx0` is set to `CtxParameterList[setId][2]`
 - `pStateIdx1` is set to `CtxParameterList[setId][3]`

10.3.2.3 Initialization process for context variables

Outputs of this process are the initialized DeepCABAC context variables distinguished by the associated syntax element and by `ctxIdx`.

For each context variable, the two variables `pStateIdx0` and `pStateIdx1` are both set to 0, a variable `shift0` is set to 1 and a variable `shift1` is set to 4.

10.3.2.4 Initialization process for the arithmetic decoding engine

The status of the arithmetic decoding engine is represented by the variables `IvlCurrRange` and `IvlOffset`. In the initialization procedure of the arithmetic decoding process, `IvlCurrRange` is set equal to 510 and `IvlOffset` is set equal to the value returned from `read_bits(9)` interpreted as a 9 bit binary representation of an unsigned integer with the most significant bit written first.

The bitstream shall not contain data that result in a value of `IvlOffset` being equal to 510 or 511.

10.3.3 Binarization process

10.3.3.1 General

Input to this process is a request for a syntax element.

Output of this process is the binarization of the syntax element.

All syntax elements use fixed-length (FL) binarization process.

The specification of the fixed-length (FL) binarization process is given in [subclause 10.3.3.2](#).

10.3.3.2 Fixed-length binarization process

Input to this process is a request for a fixed-length (FL) binarization.

Output of this process is the FL binarization associating each value symbolVal with a corresponding bin string.

FL binarization is constructed by using the fixedLength-bit unsigned integer bin string of the symbol value symbolVal, where $\text{fixedLength} = \text{Ceil}(\text{Log}_2(\text{cMax} + 1))$. The indexing of bins for the FL binarization is such that the binIdx = 0 relates to the most significant bit with increasing values of binIdx towards the least significant bit.

10.3.4 Decoding process flow

10.3.4.1 General

Inputs to this process are all bin strings of the binarization of the requested syntax element as specified in [subclause 10.3.3](#).

Output of this process is the value of the syntax element.

This process specifies how each bin of a bin string is parsed for each syntax element. After parsing each bin, the resulting bin string is compared to all bin strings of the binarization of the syntax element and the following applies:

- If the bin string is equal to one of the bin strings, the corresponding value of the syntax element is the output.
- Otherwise (the bin string is not equal to one of the bin strings), the next bit is parsed.

While parsing each bin, the variable binIdx is incremented by 1 starting with binIdx being set equal to 0 for the first bin.

The parsing of each bin is specified by the following two ordered steps:

1. The derivation process for ctxIdx and bypassFlag as specified in [subclause 10.3.4.2](#) is invoked with binIdx as input and ctxIdx and bypassFlag as outputs.
2. The arithmetic decoding process as specified in [subclause 10.3.4.3.2](#) is invoked with ctxIdx and bypassFlag as inputs and the value of the bin as output.

10.3.4.2 Derivation process for ctxIdx and bypassFlag

10.3.4.2.1 General

Input to this process is the position of the current bin within the bin string, binIdx.

Outputs of this process ctxIdx and bypassFlag.

The values of ctxIdx and bypassFlag are derived as follows based on the entries for binIdx of the corresponding syntax element in [Table 29](#):

- If the entry in [Table 29](#) is not equal to "bypass" or "na", the values of binIdx are decoded by invoking the DecodeDecision process as specified in [subclause 10.3.4.3.2](#) and the following applies:
 - The variable ctxInc is specified by the corresponding entry in [Table 29](#) and when more than one value is listed in [Table 29](#) for a binIdx, the assignment process for ctxInc for that binIdx is further specified in the subclauses given in parenthesis.
 - bypassFlag is set equal to 0.
- Otherwise, if the entry in [Table 29](#) is equal to "bypass", the values of binIdx are decoded by invoking the DecodeBypass process as specified in [subclause 10.3.4.3.4](#) and the following applies:
 - ctxIdx is set equal to 0.
 - bypassFlag is set equal to 1.
- Otherwise (the entry in [Table 29](#) is equal to "na"), the values of binIdx do not occur for the corresponding syntax element.

Table 29 — Assignment of ctxInc to syntax elements with context coded bins

Syntax element	binIdx					
	0	1	2	3	4	>= 5
row_skip_list[j]	0	na	na	na	na	na
sig_flag	0.47 (subclause 10.3.4.2.2)	na	na	na	na	na
sign_flag	0.4 (subclause 10.3.4.2.3)	na	na	na	na	na
abs_level_greater_x[j]	$4*j+(0..1)$ (subclause 10.3.4.2.4)	na	na	na	na	na
abs_level_greater_x2[j]	j	na	na	na	na	na
shift_idx_minus_1_present_flag	0	na	na	na	na	na

10.3.4.2.2 Derivation process of ctxInc for the syntax element sig_flag

Inputs to this process are the sig_flag decoded before the current sig_flag, if present, the state value stateId, the associated sign_flag, if present, and, if present, the co-located parameter level (coLocParam) from an incremental update contained in a reference NDU decoded before the current incremental update. A reference NDU is identified as follows: Syntax elements parameter_id and device_id of the reference NDU are equal to the syntax elements parameter_id and device_id (of the current NDU), respectively. Syntax element put_node_depth of the reference NDU is equal to the syntax element put_node_depth minus 1 (of the current NDU). Whenever the put_node_depth (of the current NDU) is smaller or equal to one, the reference incremental update is not available.

If nnr_pre_flag equals 1 and if nnr_compressed_data_unit_payload_type!=NNR_PT_BLOCK, coLocParam is obtained from the residuals of the incremental updates contained in the reference NDU. Otherwise, coLocParam is obtained from the weight updates of the incremental updates contained in the reference NDU.

If no sig_flag was decoded before the current sig_flag, or if the current sig_flag is the first sig_flag after an invocation of the initialization process for probability estimation parameters of [subclause 10.3.2.2](#), the sig_flag decoded before the current sig_flag is inferred to be 0. If no sign_flag associated with the previously decoded sig_flag was decoded, or if the current sig_flag is the first sig_flag after an invocation of the initialization process for probability estimation parameters of [subclause 10.3.2.2](#), the sign_flag associated with the previously decoded sig_flag is inferred to be 0. If no co-located parameter level from an incremental update decoded before the current incremental update is available or if temporal_context_modeling_flag is

equal to zero, it is inferred to be 0. A co-located parameter level means the parameter level in the same tensor at the same position in the reference incremental update.

Output of this process is the variable `ctxInc`.

The variable `curParaId` is set equal to parameter identifier of the currently decoded parameter.

The variable `ctxInc` is derived as follows:

- If `AnySigBeforeFlag[ curParaId ]` is equal 1 and `hist_dep_sig_prob_enabled_flag` is equal to 1, the following applies:
 - `ctxInc` is set to `stateId + 40`
- Otherwise (`AnySigBeforeFlag[ curParaId ]` is equal to 0 or `hist_dep_sig_prob_enabled_flag` is equal to 0), the following applies:
 - If `coLocParam` is equal to 0 the following applies:
 - If the previously decoded `sig_flag` is equal to 0, `ctxInc` is set to `stateId*3`.
 - Otherwise, if `sign_flag` associated with the previously decoded `sig_flag` is equal to 0, `ctxInc` is set to `stateId*3+1`.
 - Otherwise, `ctxInc` is set to `stateId*3+2`.
 - If `coLocParam` is not equal to 0 the following applies:
 - If `coLocParam` is greater than 1 or less than -1, `ctxInc` is set to `stateId*2+24`.
 - Otherwise, `ctxInc` is set to `stateId*2+25`.

10.3.4.2.3 Derivation process of `ctxInc` for the syntax element `sign_flag`

Inputs to this process are the `sig_flag` decoded before the current `sig_flag`, if present, the associated `sign_flag`, if present, and, if present, the co-located parameter level (`coLocParam`) from an incremental update contained in a reference NDU decoded before the current incremental update. A reference NDU is identified as follows: Syntax elements `parameter_id` and `device_id` of the reference NDU are equal to the syntax elements `parameter_id` and `device_id` (of the current NDU), respectively. Syntax element `put_node_depth` of the reference NDU is equal to the syntax element `put_node_depth` minus 1 (of the current NDU). Whenever the `put_node_depth` (of the current NDU) is smaller or equal to one, the reference incremental update is not available.

If `nnr_pre_flag` equals 1 and if `nnr_compressed_data_unit_payload_type` != `NNR_PT_BLOCK`, `coLocParam` is obtained from the residuals of the incremental updates contained in the reference NDU. Otherwise, `coLocParam` is obtained from the weight updates of the incremental updates contained in the reference NDU.

If no `sig_flag` was decoded before the current `sig_flag`, or if the current `sig_flag` is the first `sig_flag` after an invocation of the initialization process for probability estimation parameters of [subclause 10.3.2.2](#), the `sig_flag` decoded before the current `sig_flag` is inferred to be 0. If no `sign_flag` associated with the previously decoded `sig_flag` was decoded or if the current `sign_flag` is the first `sign_flag` after an invocation of the initialization process for probability estimation parameters of [subclause 10.3.2.2](#), the `sign_flag` associated with the previously decoded `sig_flag`, it is inferred to be 0. If no co-located parameter level from an incremental update decoded before the current incremental update is available or if `temporal_context_modeling_flag` is equal to zero, it is inferred to be 0. A co-located parameter level means the parameter level in the same tensor at the same position in the reference incremental update.

Output of this process is the variable `ctxInc`.

The variable `ctxInc` is derived as follows:

- If `coLocParam` is equal to 0 the following applies:
 - If the previously decoded `sig_flag` is equal to 0, `ctxInc` is set to 0.
 - Otherwise, if `sign_flag` associated with the previously decoded `sig_flag` is equal to 0, `ctxInc` is set to 1.
 - Otherwise, `ctxInc` is set to 2.
- If `coLocParam` is not equal to 0 the following applies:
 - If `coLocParam` is less than 0, `ctxInc` is set to 3.
 - Otherwise, `ctxInc` is set to 4.

10.3.4.2.4 Derivation process of `ctxInc` for the syntax element `abs_level_greater_x[j]`

Inputs to this process are the `sign_flag` decoded before the current syntax element `abs_level_greater_x[j]` and, if present, the co-located parameter level (`coLocParam`) from an incremental update contained in a reference NDU decoded before the current incremental update. A reference NDU is identified as follows: Syntax elements `parameter_id` and `device_id` of the reference NDU are equal to the syntax elements `parameter_id` and `device_id` (of the current NDU), respectively. Syntax element `put_node_depth` of the reference NDU is equal to the syntax element `put_node_depth` minus 1 (of the current NDU). Whenever the `put_node_depth` (of the current NDU) is smaller or equal to one, the reference incremental update is not available.

If `nnr_pre_flag` equals 1 and if `nnr_compressed_data_unit_payload_type` != `NNR_PT_BLOCK`, `coLocParam` is obtained from the residuals of the incremental updates contained in the reference NDU. Otherwise, `coLocParam` is obtained from the weight updates of the incremental updates contained in the reference NDU.

If no co-located parameter level from an incremental update decoded before the current incremental update is available or if `temporal_context_modeling_flag` is equal to zero, it is inferred to be 0. A co-located parameter level means the parameter level in the same tensor at the same position in the reference incremental update.

Output of this process is the variable `ctxInc`.

The variable `ctxInc` is derived as follows:

- If `coLocParam` is equal to zero the following applies:
 - If `sign_flag` is equal to 0, `ctxInc` is set to $2*j$.
 - Otherwise, `ctxInc` is set to $2*j+1$.
- If `coLocParam` is not equal to zero the following applies:
 - If `coLocParam` is greater or equal to j or is lower or equal to $-j$, `ctxInc` is set to $2*j+2*maxNumNoRemMinus1$
 - Otherwise, `ctxInc` is set to $2*j + 2* macNumNoRemMinus1 +1$.

10.3.4.3 Arithmetic decoding process

10.3.4.3.1 General

Inputs to this process are `ctxIdx` and `bypassFlag`, as derived in [subclause 10.3.4.2](#).

Output of this process is the value of the bin.

For decoding the value of a bin, the `ctxIdx` and the `bypassFlag` are passed to the arithmetic decoding process `DecodeBin( ctxIdx, bypassFlag )`, which is specified as follows:

- If `bypassFlag` is equal to 1, `DecodeBypass( )` as specified in [subclause 10.3.4.3.4](#) is invoked.
- Otherwise, `DecodeDecision( ctxIdx )` as specified in [subclause 10.3.4.3.2](#) is invoked.

NOTE Arithmetic coding is based on the principle of recursive interval subdivision. Given a probability estimation $p(0)$ and $p(1) = 1 - p(0)$ of a binary decision $(0, 1)$, an initially given code sub-interval with the range `IvlCurrRange` would be subdivided into two sub-intervals having range $p(0) * \text{IvlCurrRange}$ and $\text{IvlCurrRange} - p(0) * \text{IvlCurrRange}$, respectively. Depending on the decision, which has been observed, the corresponding sub-interval will be chosen as the new code interval, and a binary code string pointing into that interval would represent the sequence of observed binary decisions. It is useful to distinguish between the most probable symbol (MPSym) and the least probable symbol (LPSym), so that binary decisions are identified as either MPSym or LPSym, rather than 0 or 1. Given this terminology, each context is specified by the probability p_{LPSym} of the LPSym and the value of MPSym (`valMps`), which is either 0 or 1. The arithmetic core engine in this document has three distinct properties:

- The probability estimation is performed by means of two exponential decay estimators, where the average of the probability estimates is used for determining sub-intervals.
- The range `IvlCurrRange` representing the state of the coding engine is quantized to a small set $\{Q_1, \dots, Q_8\}$ of pre-set quantization values prior to the calculation of the new interval range. Storing a table containing all 32x8 pre-computed product values of $Q_i * p_{\text{LPSym}}(p\text{StateIdx})$ allows a multiplication-free approximation of the product $\text{IvlCurrRange} * p_{\text{LPSym}}(p\text{StateIdx})$.
- For syntax elements or parts thereof for which an approximately uniform probability distribution is assumed to be given a separate simplified encoding and decoding bypass process is used.

10.3.4.3.2 Arithmetic decoding process for a binary decision

10.3.4.3.2.1 General

Input to this process is the variables `ctxIdx`.

Output of this process is the decoded value `binVal`.

For decoding a single decision (`DecodeDecision`), the following applies:

1. The value of the variable `ivlLpsRange` is derived as follows:
 - Given the current value of `IvlCurrRange`, the variable `qRangeIdx` is derived as follows:

$$qRangeIdx = \text{IvlCurrRange} \& 0xe0$$
 - Given `qRangeIdx`, `pStateIdx0` and `pStateIdx1` associated with `ctxIdx` and the current syntax element, `valMps` and `ivlLpsRange` are derived as follows:

$$\text{valMps} = 16 * pStateIdx0 + pStateIdx1 \geq 0$$

$$\text{rlpsTable} = [128, 112, 97, 84, 74, 65, 57, 50, 45, 39, 34, 30, 27, 23, 20, 18, 15, 14, 12, 11, 10, 9, 7, 7, 5, 5, 4, 4, 3, 3, 2, 2, 142, 125, 108, 93, 82, 72, 63, 56, 50, 43, 38, 33, 30, 26, 22, 20, 17, 16, 13, 12, 11, 10, 8, 8, 6, 6, 5, 5, 3, 3, 2, 2, 156, 137, 119, 103, 90, 79, 70, 61, 55, 48, 42, 37, 33, 28, 24, 22, 19, 17, 15, 13, 12, 11, 9, 9, 6, 6, 5, 5, 4, 4, 2, 2, 171, 150, 130, 112, 99, 87, 76, 67, 60, 52, 46, 40, 36, 31, 27, 24, 21, 19, 16, 15, 13, 12, 10, 10, 7, 7, 6, 6, 4, 4, 3, 3, 185, 162, 141, 121, 107, 94, 82, 73, 65, 56, 50, 43, 39, 34, 29, 26, 22, 21, 17, 16, 14, 13, 11, 11, 8, 8, 6, 6, 4, 4, 3, 3, 199, 175, 152, 131, 115, 101, 89, 78, 70, 61, 54, 47, 42, 36, 31, 28, 24, 22, 19, 17, 15, 14, 12, 12, 8, 8, 7, 7, 5, 5, 3, 3, 213, 187, 163, 140, 123, 108, 95, 84, 75, 65, 58, 50, 45, 39, 33, 30, 26, 24, 20, 18, 16, 15, 13, 13,$$

9, 9, 7, 7, 5, 5, 3, 3, 228, 200, 174, 150, 132, 116, 102, 90, 80, 70, 62, 54, 48, 42, 36, 32,
28, 26, 22, 20, 18, 16, 14, 14, 10, 10, 8, 8, 6, 6, 4, 4]
ivlLpsRange = rlpTable[(abs((16 * pStateIdx0 + pStateIdx1) >> 7)) + qRangeIdx]

2. The variable IvlCurrRange is set equal to IvlCurrRange – ivlLpsRange and the following applies:
 - If IvlOffset is greater than or equal to IvlCurrRange, the variable binVal is set equal to 1 – valMps, IvlOffset is decremented by IvlCurrRange, and IvlCurrRange is set equal to ivlLpsRange.
 - Otherwise, the variable binVal is set equal to valMps.

Given the value of binVal, pStateIdx0, pStateIdx1, shift0, and shift 1 associated with ctxIdx and the current syntax element, the state transition is performed as specified in [subclause 10.3.4.3.2.2](#). Depending on the current value of IvlCurrRange, renormalization is performed as specified in [subclause 10.3.4.3.3](#).

10.3.4.3.2.2 State transition process

Inputs to this process are the current pStateIdx0 and pStateIdx1, the variables shift0 and shift1 and the decoded value binVal.

Outputs of this process are the updated pStateIdx0 and pStateIdx1 of the context variable associated with ctxIdx.

Depending on the decoded value binVal, the update of the two variables pStateIdx0 and pStateIdx1 associated with ctxIdx and with the syntax element is derived as follows:

transitionTable = [2512, 2288, 2064, 1840, 1616, 1392, 1168, 944, 720, 560, 464, 368, 272, 208, 144, 80, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 0]

sign = 2 * binVal – 1

pStateIdx0 += sign * (transitionTable[16 + (sign * pStateIdx0 >> 3)] >> (4 + shift0))

pStateIdx1 += sign * (transitionTable[16 + (sign * pStateIdx1 >> 7)] >> shift1)

10.3.4.3.3 Renormalization process in the arithmetic decoding engine

The current value of IvlCurrRange is first compared to 256 and further steps are specified as follows:

- If IvlCurrRange is greater than or equal to 256, no renormalization is needed and the process is finished;
- Otherwise (IvlCurrRange is less than 256), the renormalization loop is entered. Within this loop, the value of IvlCurrRange is doubled, i.e. left-shifted by 1 and a single bit is shifted into IvlOffset by using read_bits(1).

The bitstream shall not contain data that result in a value of IvlOffset being greater than or equal to IvlCurrRange upon completion of this process.

10.3.4.3.4 Bypass decoding process for binary decisions

Output of this process is the decoded value binVal.

First, the value of IvlOffset is doubled, i.e. left-shifted by 1 and a single bit is shifted into IvlOffset by using read_bits(1). Then, the value of IvlOffset is compared to the value of IvlCurrRange and further steps are specified as follows:

- If IvlOffset is greater than or equal to IvlCurrRange, the variable binVal is set equal to 1 and IvlOffset is decremented by IvlCurrRange.
- Otherwise (IvlOffset is less than IvlCurrRange), the variable binVal is set equal to 0.

The bitstream shall not contain data that result in a value of `IvlOffset` being greater than or equal to `IvlCurrRange` upon completion of this process.

10.3.4.3.5 Decoding process for binary decisions before termination

Output of this process is the decoded value `binVal`.

The decoding process is specified as follows:

First, the value of `IvlCurrRange` is decremented by 2. Then, the value of `IvlOffset` is compared to the value of `IvlCurrRange` and further steps are specified as follows:

- If `IvlOffset` is greater than or equal to `IvlCurrRange`, the variable `binVal` is set equal to 1, no renormalization is carried out, and DeepCABAC decoding is terminated. The last bit inserted in register `IvlOffset` is equal to 1.
- Otherwise (`IvlOffset` is less than `IvlCurrRange`), the variable `binVal` is set equal to 0 and renormalization is performed as specified in [subclause 10.3.4.3.3](#).

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-17:2024

Annex A (normative)

Implementation for NNEF

A.1 General

Neural Network Exchange Format is a data format for exchanging information about (trained) neural networks as specified in NNEF-v1.0.3. This annex specifies how compressed representation of neural networks, as specified in this document, are utilized in NNEF context.

NNEF as specified in NNEF-v1.0.3 is taken as reference for Neural Network Exchange Format (NNEF). Any bitstream elements in NNEF format shall be implemented in accordance with NNEF-v1.0.3.

A.2 Identifiers

The identifiers defined in [Table A.1](#) shall be used to indicate NNEF information included in the bitstream.

Table A.1 — Identifiers for NNEF

element	value	Identifier	Description
topology_storage_format	1	NNR_TPL_NNEF	Neural network topology information is stored in NNEF as specified in NNEF-v1.0.3
quantization_storage_format	1	NNR_QNT_NNEF	Neural network (optional) quantization information is stored in NNEF as specified in NNEF-v1.0.3

A.3 Definitions for use in NNR bitstream

NNEF topology is a textual information describing the structure of the neural network according to the syntax as specified in [subclause 3.2](#) of NNEF-v1.0.3. This information is stored as a textual file under NNEF-specified file directory tree structure.

Other NNEF specific data structures and acronyms are used as specified in NNEF-v1.0.3.

A.4 Carriage of NNEF data in NNR bitstream

If NNEF topology information is provided in-band of the NNR bitstream, then the following constraints shall apply:

- NNEF topology information shall be carried in the NNR topology unit's topology_data syntax element, represented as null-terminated string, which shall be encoded as UTF-8 characters in accordance with ISO/IEC 10646.
- Additionally, NNEF topology elements may be carried as a reference list in a consecutive NNR topology unit where topology_storage_format is equal to NNR_TPL_REFLIST. Listed topology elements shall be NNEF variable labels as specified in NNEF-v1.0.3 and as null-terminated strings encoded as UTF-8 characters. All topology elements in the list shall also be present in the NNEF topology information.
- NNR topology units shall precede any NNR compressed data unit.
- When an NNR topology unit carries NNEF topology information inside the topology_data, topology_storage_format value of NNR topology unit header shall be set to NNR_TPL_NNEF.

If NNEF optional quantization information as specified in [subclause 5.1](#) of NNEF-v1.0.3 is provided in-band of the NNR bitstream, then the following constraints shall apply:

- Optional NNEF quantization information shall be carried in the NNR quantization unit's quantization_data syntax element, represented as null-terminated string, which shall be encoded as UTF-8 characters in accordance with ISO/IEC 10646.
- NNR quantization unit shall precede any NNR compressed data unit.
- quantization_storage_format value of NNR quantization unit header shall be set to NNR_QNT_NNEF.

The following constraints shall apply to NNR compressed data units:

topology_elem_id in NNR compressed data unit header and topology_elem_id_list values inside the topology_elements_ids_list() of NNR topology unit payload shall contain the related NNEF variable label as specified in NNEF-v1.0.3. This enables mapping of a uniquely identifiable NNEF data structure to an NNR compressed data unit.

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-17:2024

Annex B (informative)

Implementation for ONNX®

B.1 General

Open Neural Network Exchange^[1] is a data format for exchanging information about (trained) neural networks. This annex specifies how compressed representation of neural networks, as specified in this document, are utilized in ONNX context.

In this version of specification,^[1] is taken as reference for Open Neural Network Exchange (ONNX).

B.2 Identifiers

The identifiers defined in [Table B.1](#) are used to indicate ONNX information included in the bitstream.

Table B.1 — Identifiers for ONNX

element	value	Identifier	Description
topology_storage_format	2	NNR_TPL_ONNX	Neural network topology information is stored in ONNX format as specified in Reference [1] .
quantization_storage_format	2	NNR_QNT_ONNX	Neural network (optional) quantization information is stored in ONNX format as specified in Reference [1] .

B.3 Definitions for use in NNR bitstream

ONNX topology is described in a GraphProto using NodeProto protobuf messages that contain information about the structure of the neural network. This information is stored under ONNX-specified GraphProto at model.graph.node.

ONNX weights are optionally stored in TensorProtos as specified in Reference [\[1\]](#) under model.graph.initializer.

B.4 Carriage of ONNX data in NNR bitstream

If ONNX topology information is provided in-band of the NNR bitstream, then the following constraints should apply:

- ONNX topology information should be carried in the NNR topology unit's topology_data syntax element, represented as null-terminated string encoded as UTF-8 characters as specified in ISO/IEC 10646.
- Topology_storage_format value of NNR topology unit header should be set to NNR_TPL_ONNX.
- If ONNX topology information is not carried as part of NNR bitstream, then an NNR topology unit with topology_storage_format equal to NNR_TPL_ONNX should carry a zero-sized and null-terminated string encoded as UTF-8 characters in its topology_data.
- Additionally, ONNX topology elements may be carried as a reference list in a consecutive NNR topology unit where topology_storage_format is equal to NNR_TPL_REFLIST. All topology elements in the list should be present in the ONNX topology information, if ONNX topology information is carried as part of the NNR bitstream (i.e. inside the topology_data as a non-zero sized and null terminated string encoded as UTF-8 characters).