



**International
Standard**

ISO/IEC 15938-18

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

**Part 18:
Conformance and reference
software for compression of neural
networks**

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

*Partie 18: Conformité et logiciel de référence pour la compression
des réseaux neuronaux*

**Second edition
2025-02**

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-18:2025



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2025

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	iv
Introduction	v
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 Conformance testing	1
4.1 General	1
4.2 Conformance testing for decoder	2
4.3 Conformance testing for bitstreams	2
4.4 Models and reference bitstreams	2
4.5 Procedure to test decoders	10
4.5.1 General	10
4.5.2 Complete model	10
4.5.3 Incremental model	11
4.5.4 Decoding self-contained NNC bitstreams	11
4.5.5 Decoding NNC bitstreams using out-of-band parameters	11
4.6 Procedure to test bitstreams	11
5 Reference software	11
5.1 General	11
5.2 Software location and license	12
5.3 Software installation	12
5.4 Software architecture	12
5.4.1 General	12
5.4.2 Parameter reduction methods	13
5.4.3 Parameter approximation	14
5.4.4 Reconstruction	14
5.4.5 Encode	14
5.4.6 Decode	14
5.5 Data structures and interfaces	14
5.5.1 model_info: Shared model information	14
5.5.2 approx_data – Data structure	16
5.5.3 nnc_core – Main module	18
5.5.4 nnc_core.nnr_model – Module for handling model related functionalities	22
Annex A (informative) Implementation in Python®	25
Bibliography	26

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 not 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 15938-18:2023), which has been technically revised.

The main changes are:

- additional conformance bitstreams and an extension of the reference software to cover the features added in ISO/IEC 15938-17:2024.

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

This document describes conformance testing and the reference software for ISO/IEC 15938-17. The reference software includes both encoder and decoder functionality.

The reference software is useful in aiding users of a standard for coding neural networks to establish and test conformance and interoperability, and to educate users and demonstrate the capabilities of the standard. For these purposes, the accompanying software is provided as an aid for the study and implementation of ISO/IEC 15938-17.

The purpose of this document is to provide the following:

- A set of reference bitstreams conforming to ISO/IEC 15938-17.
- Description of procedures to test conformance of bitstreams and decoders to ISO/IEC 15938-17.
- Reference decoder software capable of decoding bitstreams that conform to ISO/IEC 15938-17 in a manner that conforms to the decoding process specified in ISO/IEC 15938-17.
- Reference encoder software capable of producing bitstreams that conform to ISO/IEC 15938-17.

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-18:2025

IECNORM.COM : Click to view the full PDF of ISO/IEC 15938-18:2025

Information technology — Multimedia content description interface —

Part 18:

Conformance and reference software for compression of neural networks

1 Scope

This document specifies conformance testing procedures for implementations of ISO/IEC 15938-17 and provides conformance bitstreams. It also provides the reference software for ISO/IEC 15938-17 which is an integral part 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 15938-17:2024, *Information technology — Multimedia content description interface — Part 17: Compression of neural networks for multimedia content description and analysis*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 15938-17 and the following apply.

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

array

ordered list of elements where all elements are of the same type

3.2

dictionary

ordered list of key/value pairs where each key/value pair is a list of two elements with the first element being denoted 'key' and with the second element being denoted 'value'

4 Conformance testing

4.1 General

[Subclauses 4.2](#) through [4.6](#) specify tests for verifying the conformance of bitstreams as well as decoders. These tests make use of test data (bitstreams and related models) available at <https://standards.iso.org/iso-iec/15938/-18/ed-2/en> (described in detail in [4.4](#)) and follow the procedure described in [4.5](#).

4.2 Conformance testing for decoder

The decoder conformance is specified in ISO/IEC 15938-17:2024, Clause 7.

4.3 Conformance testing for bitstreams

The bitstream conformance is specified in ISO/IEC 15938-17:2024, Clause 7.

4.4 Models and reference bitstreams

A set of bitstreams and related neural network models is provided for conformance testing. When neural network models are provided, they are the source models used to generate one or more compressed bitstreams. The code defining these models is provided as part of the reference software distribution available at <https://standards.iso.org/iso-iec/15938/-18/ed-2/en>. Models and bitstreams are provided as "reference_models_bitstreams_ed2_v1.1.zip".

A bitstream may be generated:

- from a commonly used neural network model,
- from synthetic data by encoding it with the reference encoder,
- by creating a conformant synthetic bitstream directly, without using an encoder.

Where applicable, the dataset used to train the model is mentioned for information only. However, the dataset is not needed for conformance testing and thus not provided.

[Table 1](#) summarizes the provided bitstreams for entire models, [Table 2](#) summarizes the provided bitstreams for incrementally coded models. The tables also list the features tested with the bitstreams, and the reference encoder configuration used to generate the bitstream (where applicable). If decoding the bitstream requires out-of-band parameters (e.g. information that is derived from the network topology description), those parameters are also provided. We refer to source models and datasets using the following names:

- ImageNet for the dataset described in Reference [\[5\]](#),
- CIFAR-100 for the dataset described in Reference [\[6\]](#),
- DCase for the dataset described in Reference [\[7\]](#) and the model described in Reference [\[8\]](#),
- PASCAL VOC for the dataset described in Reference [\[12\]](#),
- chest_xray_v3 for the dataset described in Reference [\[13\]](#),
- MobileNet V2 for model described in Reference [\[9\]](#),
- UC12B for the model described in Reference [\[10\]](#),
- VGG16 for the model described in Reference [\[11\]](#), and
- ResNet18 for the model described in Reference [\[14\]](#).

Table 1 — Bitstreams and related models for conformance testing of coding entire models

Bitstream id	Source model	Data-set	Relevant technology ^a	Features tested	Reference encoder configuration
perf_map_sparse_MobileNetV2.nctm	MobileNetV2	Image Net	6.3.4.3	sparsification performance map	qp_density = 2 scan_order = 1 approx_method = "codebook" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
perf_map_prune_DCCase.nctm	DCCase	DCCase	6.3.4.3	pruning performance map	qp_density = 2 scan_order = 1 approx_method = "codebook" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
perf_map_sparse_prune_UC12B.nctm	UC12B	CIFAR-100	6.3.4.3	Sparsification and Pruning Performance map	qp_density = 2 scan_order = 1 approx_method = "codebook" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
perf_map_sparse_VGG16.nctm	VGG16	Image Net	6.3.4.3	Sparsification Performance Map (pruned model)	qp_density = 2 scan_order = 1 approx_method = "codebook" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
prune_tpl_cont_sparse_bm_DCCase.nctm	DCCase	DCCase	6.3.4.5	Prune Topology - sparse bitmask	encode_tpl_only = True partial_data_counter = 0
prune_tpl_cont_prune_bm_VGG16.nctm	VGG16	Image Net	6.3.4.5	Prune Topology - prune bitmask	encode_tpl_only = True partial_data_counter = 0

Table 1 (continued)

Bitstream id	Source model	Data-set	Relevant technology ^a	Features tested	Reference encoder configuration
prune_tpl_cont_comb_bm_VGG16.nctm	VGG16	Image Net	6.3.4.5	Prune Topology - combined bitmask	encode_tpl_only = True partial_data_counter = 0
prune_tpl_cont_prune_dictionary_DCase.nctm	DCase	DCase	6.3.4.5	Prune Topology - prune dictionary	encode_tpl_only = True topology_indexed_reference_flag = False partial_data_counter = 0
prune_tpl_cont_prune_dictionary_idx_ResNet50.nctm	ResNet50	Image Net	6.3.4.5	Prune Topology - prune dictionary (indexed elem id)	encode_tpl_only = True topology_indexed_reference_flag = True partial_data_counter = 0
tpl_reflist_DCase.nctm	DCase	DCase	6.3.4.5, 6.3.3.7	Topology Re-list	encode_tpl_only = True partial_data_counter = 0
partial_data_counter_VGG16_ndu_size_65536.nctm	VGG16	Image Net	6.3.3.1	Partial data counter	max_ndu_nnr_unit_size = 65536
partial_data_counter_VGG16_ndu_size_32768.nctm	VGG16	Image Net	6.3.3.1	Partial data counter	max_ndu_nnr_unit_size = 32768
partial_data_counter_VGG16_ndu_size_16384.nctm	VGG16	Image Net	6.3.3.1	Partial data counter	max_ndu_nnr_unit_size = 16384
partial_data_counter_DCase_ndu_size_2048.nctm	DCase	DCase	6.3.3.1	Partial data counter	max_ndu_nnr_unit_size = 2048
partial_data_counter_DCase_ndu_size_1024.nctm	DCase	DCase	6.3.3.1	Partial data counter	max_ndu_nnr_unit_size = 1024
deepCABAC_ResNet50_1_qp-38_qp_density2.nctm	ResNet50	Image Net	10 9.2.1 / 9.3.1	DeepCABAC entropy coding, uniform quantization	see verify_all.sh
dependent_quantization_ResNet50_2_qp-38_qp_density2.nctm	ResNet50	Image Net	9.2.3 / 9.3.3 6.3.3.7	Dependent scalar quantization	see verify_all.sh
deepCABAC_qp_density_MobileNetV2_3_qp-38_qp_density2.nctm	MobileNetV2	Image Net	9.3	QpDensity	see verify_all.sh
deepCABAC_qp_density_MobileNetV2_4_qp-76_qp_density3.nctm	MobileNetV2	Image Net	9.3	QpDensity	see verify_all.sh
block_scan_order_8x8_cabac_entry_points_ResNet50_5_qp-38_qp_density2.nctm	ResNet50	Image Net	4.12 / 6.4.3.7 / 6.4.3.8 / 7.3.6	Block scan order / cabac entry points	see verify_all.sh

Table 1 (continued)

Bitstream id	Source model	Data-set	Relevant technology ^a	Features tested	Reference encoder configuration
block_scan_order_16x16_cabac_entry_points_ResNet50_6_qp-38_qp_density2.nctm	ResNet50	Image Net	4.12 / 6.4.3.7 / 6.4.3.8 / 7.3.6	Block scan order / cabac entry points	see verify_all.sh
codebook_signaling_MobileNetV2_7_qMse0.00001.nctm	MobileNetV2	Image Net	9.2.3 / 9.3.2	Code-book-based quantization	see verify_all.sh
local_scaling_DCase_8_qp-38_qp_density2.nctm	DCase	DCase	G.1.10 / 8.3.9	Local scaling	see verify_all.sh
batchnorm_folding_MobileNetV2_9_qp-38_qp_density2.nctm	MobileNetV2	Image Net	8.2.1 / 8.3.8	BatchNorm Folding	see verify_all.sh
out_of_band_signaling_ResNet50_10_qp-38_qp_density2.nctm	ResNet50	Image Net	6.3.3.7 / 6.4.3.7	Out-of-band signaling ^b	see verify_all.sh
deepCABAC_8bit_ResNet50_PYT-zoo_11_qp0_qp_density4.nctm	ResNet50	Image Net	9.2.1 / 9.3.1	Uniform quantization with limited precision (8bit)	see verify_all.sh
deepCABAC_8bit_MobileNetV2_PY-Tzoo_12_qp0_qp_density4.nctm	MobileNet V2	Image Net	9.2.1 / 9.3.1	Uniform quantization with limited precision (8bit)	see verify_all.sh
deepCABAC_4bit_VGG16_PYTzoo_13_qp0_qp_density4.nctm	VGG16	Image Net	9.2.1 / 9.3.1	Uniform quantization with limited precision (4bit)	see verify_all.sh
deepCABAC_8bit_UC12B_14_qp0_qp_density4.nctm	UC12B	CIFAR-100	9.2.1 / 9.3.1	Uniform quantization with limited precision (8bit)	see verify_all.sh
deepCABAC_4bit_DCase_15_qp0_qp_density4.nctm	DCase	DCase	9.2.1 / 9.3.1	Uniform quantization with limited precision (4bit)	see verify_all.sh
perf_map_sparse_MobileNetV2_low8.nctm	MobileNetV2	Image Net	6.3.4.3	sparsification performance map (8bit)	qp_density = 2 scan_order = 1 approx_method = "uniform" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0

Table 1 (continued)

Bitstream id	Source model	Data-set	Relevant technology ^a	Features tested	Reference encoder configuration
perf_map_prune_DCase_bw8.nctm	DCase	DCase	6.3.4.3	pruning performance map (8bit)	qp_density = 2 scan_order = 1 approx_method = "uniform" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
perf_map_sparse_prune_UC12B_bw8.nctm	UC12B	CIFAR-100	6.3.4.3	Sparsification and Pruning Performance map (8bit)	qp_density = 2 scan_order = 1 approx_method = "uniform" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
perf_map_sparse_VGG16_bw8.nctm	VGG16	Image Net	6.3.4.3	Sparsification Performance Map (pruned model) (8bit)	qp_density = 2 scan_order = 1 approx_method = "uniform" qp = 35 qp_density = 2 opt_qp = False disable_dq = True lambda_scale = 0.0 cb_size_ratio = 5000 q_mse = 0.00001 param_opt_flag = False cabac_unary_length_minus1 = 9 partial_data_counter = 0
partial_data_counter_DCase_ndu_size_2048_bw8.nctm	DCase	DCase	6.3.3.1	Partial data counter (8bit)	max_ndu_nnr_unit_size = 2048 approx_method = "uniform"
partial_data_counter_DCase_ndu_size_1024_bw8.nctm	DCase	DCase	6.3.3.1	Partial data counter (8bit)	max_ndu_nnr_unit_size = 1024 approx_method = "uniform"
^a Specifies the applicable subclause in ISO/IEC 15938-17:2024.					
^b As specified in out_of_band_signaling_ResNet50_10_qp-38_qp_density2_oob.json					

The bitstreams in [Table 2](#) are provided for testing the decoding of a sequence of bitstreams representing incremental updates of models. They shall be processed in the order indicated by the sequence number. The character '#' in the name of the bitstream indicates the sequence number to be inserted, iterating over all values in the range specified in column 'Seq.' (including the boundaries).

Bitstreams for the tested features Batchnorm folding (BNF), Row skipping (RS), Temporal context modelling/adaptation (TCA) and History dependent significance probability (HDSP) use a common set of parser arguments.

The parser arguments for technologies under test are 'bn_folding': True for enabling BNF, 'row_skip': True for enabling RS, 'temporal_context': True for enabling TCA and 'hdsp': True for enabling HDSP. In a bitstream set only one of these four parser arguments is True, namely that of the technology under test, all others are set to False.

Bitstreams for the feature Predictive Residual Coding (PRE) use a common set of tools from the extended compression tool set in conjunction with PRE: 'row_skip': True, 'temporal_context': True and 'hdsp': True.

Unless otherwise specified in the 'Reference encoder configuration' column of [Table 2](#), all conformance tests indicated by a 'Bitstream id' starting with 'UC10' share a common set of parser arguments: 'use_case': 'UC10', 'batch_size': 64. Likewise for those starting with 'UC14A', 'use_case': 'UC14A', 'qp': -48, 'batch_size': 100, 'model_arch': 'vgg16', 'unidirect_compress': False, 'train_cf': True are used. For PRE and SBT, different qp values are reported as stated in the table, though the rest of the common parser arguments are the same as above. 'UC10' and 'UC14A' refer to the use cases described in the evaluation framework for development of NNC [\[15\]](#).

The common set of parser arguments during all BNF, RS, TCA and HDSP conformance tests is: 'disable_finetuning': True, 'n_epochs': 4, 'approx_method': 'uniform', 'disable_dq': True, 'param_opt': True, 'qp_density': 2, 'qp_diff_offset': 0, 'opt_qp': False, 'lambda_scale': 0.0, 'cb_size_ratio': 5000, 'q_mse': 1e-05, 'nnr_pt_block_enabled': 0, 'filter_sparsity': 0.0, 'filter_spars_gain': 1.0, 'lr': 1e-05, 'model_path': 'PYT_ZOO', 'ini_compress': False, 'max_iterations': None, 'data_split_ratios': [1.0], 'iter_opt': False, 'qp_step': 1, 'bias': 0.125, 'lsa': 0.0, 'ioq': False, 'oob_signaling': False, 'ce': 2, 'sbt_quantization_approach': 'combined', 'sbt_quantization_subapproach': 'random_round', 'export_model_diffs': False, 'sbt_bn_quant': True, 'sbt_mean_quant': False, 'pre': False, 'cache_len': 14, 'running_level': 'model'.

For PRE and SBT, the common set of parser arguments during all conformance tests are: 'disable_dq': True, 'param_opt': True, 'n_epochs': 4, 'qp_density': 2, 'qp_diff_offset': 0, 'lambda_scale': 0.0, 'cb_size_ratio': 5000, 'q_mse': 1e-05, 'nnr_pt_block_enabled': 0, 'target_sparsity': 0, 'filter_sparsity': 0.0, 'filter_spars_gain': 1.0, 'lr': 1e-05, 'model_path': 'PYT_ZOO', 'max_iterations': None, 'data_split_ratios': [1.0], 'bias': 0.125, 'lsa': 0.0, 'ioq': False, 'oob_signaling': False, 'sbt_quantization_approach': 'combined', 'sbt_quantization_subapproach': 'random_round', 'sbt_bn_quant': False, 'sbt_mean_quant': False, 'cache_len': 14, 'running_level': 'model', 'sparse': False, 'target_sparsity': 0.0.

All encoder-relevant parser arguments that are not listed above or differ from the mentioned ones are specified in the "Reference encoder configuration" column of [Table 2](#).

Table 2 — Bitstreams and related models for conformance testing of coding incremental models

Bitstream set	Bitstream id	Seq.	Source model	Dataset	Relevant tech-nol-ogy ^a	Features tested	Reference encoder configuration
UC10 MobileNet_v2 full model	UC10_mobilenet_v2_uniform_qp-34_BNF_epoch_#.nctm	0..3	MobileNet_V2	PAS-CAL-VOC	8.3.8	Batchnorm folding (BNF)	'model_arch': 'mobilenet_v2', 'qp': -34, 'bn_folding': True, 'sparse': False, 'target_sparsity': 0.0
UC10 MobileNet_v2 full model (with sparsification)	UC10_mobilenet_v2_uniform_qp-34_BNF_sparsity_0.9_epoch_#.nctm	0..3	MobileNet_V2	PAS-CAL-VOC	8.3.8 G.1.5 G.1.6	Batchnorm folding (BNF)	'model_arch': 'mobilenet_v2', 'qp': -34, 'bn_folding': True, 'sparse': True, 'target_sparsity': 0.9
UC10 ResNet18 full model	UC10_resnet18_uniform_qp-37_RS_epoch_#.nctm	0..3	ResNet18	PAS-CAL-VOC	10.1.1.2	Row skip-ping (RS)	'model_arch': 'resnet18', 'qp': -37, 'row_skip': True, 'sparse': False, 'target_sparsity': 0.0
UC10 ResNet18 full model (with sparsification)	UC10_resnet18_uniform_qp-37_sparsity_0.9_RS_epoch_#.nctm	0..3	ResNet18	PAS-CAL-VOC	10.1.1.2 G.1.5 G.1.6	Row skip-ping (RS)	'model_arch': 'resnet18', 'qp': -37, 'row_skip': True, 'sparse': True, 'target_sparsity': 0.9

^a Specifies the applicable subclause in ISO/IEC 15938-17:2024.

Table 2 (continued)

Bitstream set	Bitstream id	Seq.	Source model	Dataset	Relevant technology ^a	Features tested	Reference encoder configuration
UC14A VGG16 classifier only center to clients	UC14A_vgg16_uniform_qp-48_TCA_classifier_only_epoch_#_center.nctm	0..3	VGG16	chest_xray_v3	10.1.1.4	Temporal context adaptation (TCA)	'temporal_context': True, 'sparse': False, 'target_sparsity': 0.0
UC14A VGG16 classifier only client 1 to center	UC14A_vgg16_uniform_qp-48_TCA_classifier_only_epoch_#_client1.nctm	0..3	VGG16	chest_xray_v3	10.1.1.4	Temporal context adaptation (TCA)	'temporal_context': True, 'sparse': False, 'target_sparsity': 0.0
UC14A VGG16 classifier only client 2 to center	UC14A_vgg16_uniform_qp-48_TCA_classifier_only_epoch_#_client2.nctm	0..3	VGG16	chest_xray_v3	10.1.1.4	Temporal context adaptation (TCA)	'temporal_context': True, 'sparse': False, 'target_sparsity': 0.0
UC14A VGG16 classifier only center to clients (with sparsification)	UC14A_vgg16_uniform_qp-48_TCA_sparsity_0.8_classifier_only_epoch_#_center.nctm	0..3	VGG16	chest_xray_v3	10.1.1.4 G.1.5 G.1.6	Temporal context adaptation (TCA)	'temporal_context': True, 'sparse': True, 'target_sparsity': 0.8
UC14A VGG16 classifier only client 1 to center (with sparsification)	UC14A_vgg16_uniform_qp-48_TCA_sparsity_0.8_classifier_only_epoch_#_client1.nctm	0..3	VGG16	chest_xray_v3	10.1.1.4 G.1.5 G.1.6	Temporal context adaptation (TCA)	'temporal_context': True, 'sparse': True, 'target_sparsity': 0.8
UC14A VGG16 classifier only client 2 to center (with sparsification)	UC14A_vgg16_uniform_qp-48_TCA_sparsity_0.8_classifier_only_epoch_#_client2.nctm	0..3	VGG16	chest_xray_v3	10.1.1.4 G.1.5 G.1.6	Temporal context adaptation (TCA)	'temporal_context': True, 'sparse': True, 'target_sparsity': 0.8
UC14A VGG16 classifier only center to clients	UC14A_vgg16_uniform_qp-48_HDSP_classifier_only_epoch_#_center.nctm	0..3	VGG16	chest_xray_v3	10.2.1.4 / 10.3.2.2 / 10.3.4.2.2	History dependent significance probability (HDSP)	'hdsp': True, 'sparse': False, 'target_sparsity': 0.0
UC14A VGG16 classifier only client 1 to center	UC14A_vgg16_uniform_qp-48_HDSP_classifier_only_epoch_#_client1.nctm	0..3	VGG16	chest_xray_v3	10.2.1.4 / 10.3.2.2 / 10.3.4.2.2	History dependent significance probability (HDSP)	'hdsp': True, 'sparse': False, 'target_sparsity': 0.0
UC14A VGG16 classifier only client 2 to center	UC14A_vgg16_uniform_qp-48_HDSP_classifier_only_epoch_#_client2.nctm	0..3	VGG16	chest_xray_v3	10.2.1.4 / 10.3.2.2 / 10.3.4.2.2	History dependent significance probability (HDSP)	'hdsp': True, 'sparse': False, 'target_sparsity': 0.0
UC14A VGG16 classifier only center to clients (with sparsification)	UC14A_vgg16_uniform_qp-48_HDSP_sparsity_0.8_classifier_only_epoch_#_center.nctm	0..3	VGG16	chest_xray_v3	10.2.1.4 / 10.3.2.2 / 10.3.4.2.2 G.1.5 G.1.6	History dependent significance probability (HDSP)	'hdsp': True, 'sparse': True, 'target_sparsity': 0.8
UC14A VGG16 classifier only client 1 to center (with sparsification)	UC14A_vgg16_uniform_qp-48_HDSP_sparsity_0.8_classifier_only_epoch_#_client1.nctm	0..3	VGG16	chest_xray_v3	10.2.1.4 / 10.3.2.2 / 10.3.4.2.2 G.1.5 G.1.6	History dependent significance probability (HDSP)	'hdsp': True, 'sparse': True, 'target_sparsity': 0.8

^a Specifies the applicable subclause in ISO/IEC 15938-17:2024.

Table 2 (continued)

Bitstream set	Bitstream id	Seq.	Source model	Dataset	Relevant technology ^a	Features tested	Reference encoder configuration
UC14A VGG16 classifier only client 2 to center (with sparsification)	UC14A_vgg16_uniform_qp-48_HDSP_sparsity_0.8_classifier_only_epoch_#_client2.nctm	0..3	VGG16	chest_xray_v3	10.2.1.4 / 10.3.2.2 / 10.3.4.2.2 G.1.5 G.1.6	History dependent significance probability (HDSP)	'hdsp': True, 'sparse': True, 'target_sparsity': 0.8
UC10 ResNet18 classifier only	UC10_resnet18_uniform_sbt_qp-34_classifier_only_epoch_#.nctm	0..3	ResNet18	PAS-CAL-VOC	9.1.5	Stochastic Binary Ternary (SBT) Quantization	'pre': False, 'approx_method': uniform_sbt, qp: -34, 'train_cf': True
UC10 VGG11 classifier only	UC10_vgg11_uniform_sbt_qp-36_classifier_only_epoch_#.nctm	0..3	VGG11	PAS-CAL-VOC	9.1.5	Stochastic Binary Ternary (SBT) Quantization	'pre': False, 'approx_method': uniform_sbt, qp: -36, 'train_cf': True
UC14A VGG16 full model center to clients	UC14A_vgg16_uniform_sbt_qp-52_full_model_epoch_#.center.nctm	0..3	VGG16	chest_xray_v3	9.1.5	Stochastic Binary Ternary (SBT) Quantization	'pre': False, 'approx_method': uniform_sbt, qp: -52, 'train_cf': False
UC14A VGG16 full model client 1 to center	UC14A_vgg16_uniform_sbt_qp-52_full_model_epoch_#.client1.nctm	0..3	VGG16	chest_xray_v3	9.1.5	Stochastic Binary Ternary (SBT) Quantization	'pre': False, 'approx_method': uniform_sbt, qp: -52, 'train_cf': False
UC14A VGG16 full model client 2 to center	UC14A_vgg16_uniform_sbt_qp-52_full_model_epoch_#.client2.nctm	0..3	VGG16	chest_xray_v3	9.1.5	Stochastic Binary Ternary (SBT) Quantization	'pre': False, 'approx_method': uniform_sbt, qp: -52, 'train_cf': False
UC14A VGG16 classifier only client 1 to center	UC14A_vgg16_uniform_qp-48_PRE_classifier_only_unidir_epoch_#_client1.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform, 'unidirect_compress': True
UC14A VGG16 classifier only client 2 to center	UC14A_vgg16_uniform_qp-48_PRE_classifier_only_unidir_epoch_#_client2.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform, 'unidirect_compress': True
UC14A VGG16 classifier only center to clients	UC14A_vgg16_uniform_qp-48_PRE_classifier_only_bidir_epoch_#.center.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform
UC14A VGG16 classifier only client 1 to center	UC14A_vgg16_uniform_qp-48_PRE_classifier_only_bidir_epoch_#.client1.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform
UC14A VGG16 classifier only client 2 to center	UC14A_vgg16_uniform_qp-48_PRE_classifier_only_bidir_epoch_#.client2.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform

^a Specifies the applicable subclause in ISO/IEC 15938-17:2024.

Table 2 (continued)

Bitstream set	Bitstream id	Seq.	Source model	Dataset	Relevant technology ^a	Features tested	Reference encoder configuration
UC14A VGG16 full model client 1 to center	UC14A_vgg16_uniform_qp-48_PRE_full_model_unidir_epoch_#_client1.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform, 'unidirect_compress': True
UC14A VGG16 full model client 2 to center	UC14A_vgg16_uniform_qp-48_PRE_full_model_unidir_epoch_#_client2.nctm	0..3	VGG16	chest_xray_v3	9.1.6	Predictive Residual Coding (PRE)	'pre': True, 'temporal_context': True, 'row_skip': True, 'hdsp': True, 'approx_method': uniform, 'unidirect_compress': True

^a Specifies the applicable subclause in ISO/IEC 15938-17:2024.

The bitstreams and models are provided at https://standards.iso.org/iso-iec/15938-18/ed-2/en/reference_models_bitstreams_ed2_v1.1.zip.

The parameters of the bitstreams, for which no detailed parameters are listed, can be obtained from the scripts given in column “Reference encoder configuration”, e.g. `verify_all.sh` in “run/conformance_bitstreams” of the reference software. Scripts for generating all incremental bitstreams, as specified in Table 2, and for decoder testing with these bitstreams are provided in the same directory (resp., “`verify_all_incremental.sh`” and “`decode_all_incremental.sh`”). Furthermore, it also provides detailed reference software configurations for each test. If out-of-band parameters are required, they are provided in a file conforming to JSON (as specified in ISO/IEC 21778 [16]), specified in a footnote.

4.5 Procedure to test decoders

4.5.1 General

For testing a decoder, all bitstreams that cover mandatory features as specified in ISO/IEC 15938-17 shall be tested. If support for optional features is claimed, also the bitstreams covering these features shall be tested.

On platforms with only integer or limited integer precision support, only those bitstreams indicated for the supported bit depth (“8bit” or “4bit”) shall be tested.

The procedure for testing features for decoding complete models and decoding a sequence of incremental updates of a model follow the same principles, but differ in details. The two procedures are described in the following subclauses.

4.5.2 Complete model

For each bitstream, the procedure includes the following steps:

- Decode bitstream B with the reference decoder to obtain the decoded model R_r according to 4.5.4 if no out-of-band parameters are specified and according to 4.5.5, otherwise.
- Decode bitstream B with the decoder under test to obtain the decoded model R_t according to 4.5.4 if no out-of-band parameters are specified and according to 4.5.5, otherwise.
- Compare every tensor in R_r against the corresponding tensor in R_t . The test shall be considered as passed if the following conditions hold for every pair of tensors:
 - If an integer tensor has been obtained as a result (decoding method specified in ISO/IEC 15938-17:2024, 7.3.2), the tensors shall be identical.

- If a floating point tensor has been obtained as the result of another decoding method, no corresponding pair of floating point values shall have an absolute difference larger than 0,000001.

4.5.3 Incremental model

For each set of related bitstreams (as indicated in [Table 2](#)), the procedure includes the following steps:

- a) Perform decoding of each of the bitstreams in the set with sequence numbers s (repeating this step for each s from 0 to k , where k is the number of bitstreams in the set) in the specified order:
 - 1) Decode bitstream B^s with the reference decoder to obtain the decoded model R_r^s , according to [4.5.4](#) if no out-of-band parameters are specified and according to [4.5.5](#), otherwise.
 - 2) Decode bitstream B^s with the decoder under test to obtain the decoded model R_t^s according to [4.5.4](#) if no out-of-band parameters are specified and according to [4.5.5](#), otherwise.
 - 3) If s equals 0, initialize the reconstructed models as $R_r = R_r^0$ and $R_t = R_t^0$. Otherwise apply R_r^s and R_t^s as incremental updates to R_r and R_t , respectively (e.g. adding them to the tensor from the previous iteration).
- b) Compare every tensor in R_r against the corresponding tensor in R_t . The test shall be considered as passed if the following conditions hold for every pair of tensors:
 - If an integer tensor has been obtained as a result (decoding method specified in clause 7.3.2 of ISO/IEC 15938-17), the tensors shall be identical.
 - If a floating point tensor has been obtained as the result of another decoding method, no corresponding pair of floating point values shall have an absolute difference larger than 0,000001.

4.5.4 Decoding self-contained NNC bitstreams

For self-contained bitstreams, the decoder shall not use any other inputs than the bitstream to be decoded.

4.5.5 Decoding NNC bitstreams using out-of-band parameters

The testing of decoding bitstreams using out-of-band parameters is agnostic to whether those are application provided or decoded from a particular third party topology format. The decoder shall use only the input bitstream and the parameters provided for this bitstream as specified in [Table 1](#) as inputs.

4.6 Procedure to test bitstreams

In order to test bitstreams created with an encoder other than the reference encoder, the bitstream shall be decoded using the reference decoder. The bitstream is considered conformant, if no errors occur during the decoding process.

All bitstreams that have been created with the script “verify_all.sh” specified in [Table 1](#) can be decoded with the script “decode_all.sh” in “run/conformance_bitstreams” of the reference software. In the same directory, the scripts “verify_all_incremental.sh” and “decode_all_incremental.sh” are provided for the bitstreams in [Table 2](#). The bitstreams are considered conformant, if no errors occur during the decoding process.

5 Reference software

5.1 General

The purpose of this clause is to provide the following:

- Reference decoder software capable of decoding bitstreams that conform to ISO/IEC 15938-17 in a manner that conforms to the decoding process specified in ISO/IEC 15938-17.
- Reference encoder software capable of producing bitstreams that conform to ISO/IEC 15938-17.

Some examples of uses that may be appropriate for the reference decoder software are as follows:

- As an illustration of how to perform the decoding process specified in ISO/IEC 15938-17.
- As the starting basis for the implementation of a decoder that conforms to ISO/IEC 15938-17.
- For testing the conformance of a decoder implementation with the decoding process specified in ISO/IEC 15938-17.
- For testing the conformance of a bitstream to the constraints specified for bitstream conformance in ISO/IEC 15938-17, as the software can detect and report many bitstream conformance violations.

However, the lack of the detection of any conformance violation by the reference decoder software should not be considered as definitive proof that the bitstream conforms to all constraints specified for bitstream conformance in ISO/IEC 15938-17.

Some examples of uses that may be appropriate for the reference encoder software are as follows:

- As an illustration of how to perform an encoding process that produces bitstreams that conform to the constraints specified for bitstream conformance in ISO/IEC 15938-17.
- As the starting basis for the implementation of an encoder that conforms to ISO/IEC 15938-17.
- As a means of generating bitstreams for testing the conformance of a decoder implementation with the decoding process specified in ISO/IEC 15938-17.
- As a means of evaluating and demonstrating examples of the quality that can be achieved by an encoding process that conforms to ISO/IEC 15938-17.

5.2 Software location and license

The reference software is available at <https://standards.iso.org/iso-iec/15938/-18/ed-2/en/NCTM-2.0.zip>.

The software license is provided in the file LICENSE included with the software.

5.3 Software installation

Installation instructions are provided in the file Readme.md included with the software.

5.4 Software architecture

5.4.1 General

[Figure 1](#) gives a general overview of the discussed modules that are commonly present in the end-to-end neural network compression pipeline. Each of the boxes specifies a module representing a type of compression tools. Therefore, the interfaces that define the input/output of these components are specified in order to increase the interoperability between the different components.




Figure 1 — Reference software architecture

The reference software supports out-of-band signalling to allow providing configurations and other parameters to the decoder for testing purposes. Those configurations and parameters are not part of the bitstream and expected to be provided by the application. This may ease testing of supported functionalities. Out-of-band signalling mechanisms are not part of conformance to ISO/IEC 15938-17.

In the following subclauses, float is defined as in 4.4 of ISO/IEC 15938-17, and uint8, uint32, int32 conform to u(8), u(32), i(32) as defined in 6.1.3 of ISO/IEC 15938-17. Mappings of the types to the Python^{®1)} programming language are defined in [Annex A](#).

5.4.2 Parameter reduction methods

This module represents the set of technologies that reduce the set of parameters of the network.

Input (#1): neural network (parameters + architecture) as specified by a particular deep learning framework. Examples of these frameworks are Keras^{®2)}/Tensorflow^{®3)} [2] or PyTorch^{®4)} [1].

Output (#2): neural network (parameters + architecture) as specified by a particular deep learning framework. Examples of these frameworks are Keras/Tensorflow[2] or PyTorch[1].

1) Python is the trademark of a product supplied by the Python Software Foundation. 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) Keras is the trademark of an API by François Chollet. 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) 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.

4) 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.

5.4.3 Parameter approximation

This module applies parameter approximation techniques on the extracted parameter tensors. This may include techniques such as quantization, transformation, prediction, etc.

Input (#3): A dictionary containing all parameter tensors of the network as float tensors and their respective names as keys. The names of each parameter tensor should be consistent with the nomenclature specified by the respective model definition.

Output (#4): A dictionary as specified in [5.5.2](#).

5.4.4 Reconstruction

Each parameter approximation technology provides a method that allows for reconstruction of the approximated parameter tensors resulting from the approximation procedure.

Input (#4): A dictionary in the same format as specified in output (#4) of the approximator.

Output (#6): A dictionary in the same format as specified in input (#3) of the approximator.

5.4.5 Encode

This module performs entropy encoding techniques for input data formats specified in ISO/IEC 15938-17. Auxiliary data shall be passed into the bitstream without entropy encoding in a serialized manner.

Input (#4): A dictionary in the same format as specified in output (#4) of the approximator. The key names and the dictionary structure shall not be encoded into the bitstream.

Output (#5): A bitstream representing the entropy encoded input dictionary as array of uint8 values.

5.4.6 Decode

This module performs the lossless decoding of given bitstreams.

Input (#5): The bitstreams as specified in output (#5).

Output (#4): A dictionary in the same format as specified in output (#4) of the approximator, however, with the same keys and their items are bit-exact to each other.

5.5 Data structures and interfaces

This subclause describes data structures and methods that correspond to interfaces described above.

5.5.1 **model_info**: Shared model information

Information about the model architecture and the corresponding deep learning framework that may be required by the test model is provided as variable *model_info*.

The variable *model_info* is a dictionary containing the key/value pairs specified in [Table 3](#).

Table 3 — Key/value pairs in `model_info`

Key	Value
<code>'parameter_type'</code>	Dictionary containing a key/value pair per parameter tensor of the model. Each key is a string containing the name of the parameter tensor as specified by the model architecture and the corresponding value is one of the strings <code>'conv.weight'</code> , <code>'fc.weight'</code> , <code>'conv.bias'</code> , <code>'fc.bias'</code> , <code>'bn.beta'</code> , <code>'bn.gamma'</code> , <code>'bn.mean'</code> , <code>'bn.var'</code> , or <code>'unspecified'</code> .
<code>'block_identifier'</code>	Dictionary containing a key/value pair per parameter tensor of the model that shall be grouped in a block. Each key is a string containing the name of the parameter tensor as specified by the model architecture and the corresponding value is a string identifying a block in the model architecture the parameter belongs to.
<code>'parameter_dimensions'</code>	Dictionary containing a key/value pair per parameter tensor of the model for which the parameter dimensions are not present in the bitstream. Each key is a string containing the name of the parameter tensor as specified by the model architecture and the corresponding value is a list containing the dimensions of the corresponding parameter tensor.
<code>'parameter_index'</code>	Dictionary containing a key/value pair per parameter tensor of the model. Each key is a string containing the name of the parameter tensor as specified by the model architecture and the corresponding value is an unsigned integer index. Sorting the integer indexes in ascending order defines the order of the parameter tensors.
<code>'topology_storage_format'</code>	An integer value representing syntax element <code>topology_storage_format</code> . A value equal to <code>'None'</code> indicates that no topology unit is contained in the bitstream.
<code>'performance_maps'</code>	Dictionary containing a key/value pair per included performance map. Each key is either <code>'sparsification'</code> or <code>'pruning'</code> depending on the performance map type and each corresponding value is a Dictionary containing the performance map information.
<code>'prune_rep_type'</code>	Integer indicating the type of the prune information. It should match either the identifier <code>NNR_TPL_BMSK</code> or <code>NNR_TPL_DICT</code> .
<code>'sparse_bit_mask'</code>	A list of zero/one values indicating parameter weights that should be zeroed when performing sparsification.
<code>'bit_mask_order_flag'</code>	A Boolean specifying the processing order of bit masks. True indicates row-major order and False indicates column-major order.
<code>'topology_elem_id_list'</code>	A list of strings representing the labels of each parameters in a model. Gets written into the topology reflist.

Each performance map included in `performance_maps` is a dictionary containing the key/value pairs depending on the performance map type of dictionary `model_info['performance_maps']` specified in [Table 4](#).

Table 4 — Key/value pairs in `model_info['performance_maps']`

Values in <code>model_info['performance_maps']</code> contain	Key	Value
<code>'sparsification'</code>	<code>'count_thresholds'</code>	A uint8 specifying the size of <code>'sparsification_threshold'</code> and <code>'non_zero_ratio'</code> .
<code>'sparsification'</code>	<code>'sparsification_threshold'</code>	A list of float specifying the threshold which is applied to the weights of the decoded neural network in order to set the weights to zero.
<code>'sparsification'</code>	<code>'non_zero_ratio'</code>	A list of float specifying the non-zero ratios that are achieved by applying the sparsification thresholds in <code>'sparsification_threshold'</code> to sparsify the weights.
<code>'pruning'</code>	<code>'count_pruning_ratios'</code>	A uint8 specifying the size of the <code>'pruning_ratio'</code> list.
<code>'pruning'</code>	<code>'pruning_ratio'</code>	A list of float specifying the ratios of pruned weights to the total number of weights.

Table 4 (continued)

Values in <i>model_info</i> [<i>performance_maps</i>] contain	Key	Value
'sparsification' or 'pruning'	'nn_accuracy'	A list of float specifying the overall accuracies of the NN for different sparsification thresholds or pruning ratios.
'sparsification' or 'pruning'	'count_classes'	A list of uint8 specifying the number of classes for which separate accuracies are provided for each sparsification threshold or pruning ratio.
'sparsification' or 'pruning'	'class_bitmask'	A list of uint32 specifying a class bitmask for each threshold/ratio that defines the classes included in 'nn_class_accuracy'. A 1-bit represents the presence of the class with a matching index.
'sparsification' or 'pruning'	'nn_class_accuracy'	A list of float specifying the accuracy for a certain class, when a certain sparsification threshold or pruning ratio is applied.

Depending on the values of *model_info*['prune_rep_type'], the dictionary contains one of the key/value pairs specified in [Table 5](#).

Table 5 — Key/value pairs in *model_info*['prune_rep_type']

Value in <i>model_info</i> ['prune_rep_type'] is	Key	Value
NNR_TPL_BMSK	'prune_bit_mask'	A list of zero/one values indicating parameter weights that should be pruned when performing pruning.
NNR_TPL_DICT	'prune_tpl_elem_ids'	A list of strings identifying the parameters for which new dimensions are specified. Should be a unique identifier as long as the topology definition is active.
NNR_TPL_DICT	'prune_tpl_count_dims'	A list of uint32 values indicating the number of dimensions for each parameter.
NNR_TPL_DICT	'prune_tpl_dims'	A list of uint32 values specifying the new dimensions of each parameter.

5.5.2 approx_data – Data structure

Interfaces #3, #4, and #6 use a data structure defined as variable *approx_data*.

The variable *approx_data* is a dictionary containing the key/value pairs specified in [Table 6](#).

Table 6 — Key/value pairs in *approx_data*

Key	Value
' <i>approx_method</i> '	Dictionary containing a key/value pair per approximated parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is specifying the parameter approximation method. Allowed values are ' <i>uniform</i> ' or ' <i>codebook</i> '.
' <i>parameters</i> '	Dictionary containing a key/value pair per parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is a tensor. In the case of a float parameter, it contains float values and in the case of an approximated parameter, it contains int32 values.
' <i>compressed_parameter_types</i> '	Dictionary containing a key/value pair per block identifier of the model. Each key is a string containing the name of the block identifier as specified in <i>model_info</i> and the corresponding value is identical to syntax element <i>compressed_parameter_types</i> .
' <i>decomposition_rank</i> '	Dictionary containing a key/value pair per block identifier of the model for which low-rank decomposition is used. Each key is a string containing the name of the block identifier as specified in <i>model_info</i> and the corresponding value is identical to syntax element <i>decomposition_rank</i> .
' <i>g_number_of_rows</i> '	Dictionary containing a key/value pair per block identifier of the model for which low-rank decomposition is used. Each key is a string containing the name of the block identifier as specified in <i>model_info</i> and the corresponding value is identical to syntax element <i>g_number_of_rows</i> .
' <i>scan_order</i> '	Dictionary containing a key/value pair per weight parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is identical to syntax element <i>scan_order</i> .
' <i>dq_flag</i> '	Dictionary containing a key/value pair per approximated parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is a flag indicating whether dependent quantization is used.
' <i>qp</i> '	Dictionary containing a key/value pair per approximated parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is an int32 value specifying the quantization parameter.
' <i>qp_density</i> '	An int32 value specifying syntax element <i>qp_density</i> .

The parameter tensors of the model are derived from *model_info* and from *approx_data*['*compressed_parameter_types*']. This may result in added or removed parameters relative to *model_info* depending on the values in *approx_data*['*compressed_parameter_types*'].

In addition, the dictionary contains the key/value pairs depending on the values of dictionary *approx_data*['*approx_method*'] specified in [Table 7](#).

Table 7 — Key/value pairs in `approx_data['approx_method']`

Values in <code>approx_data['approx_method']</code> contain (at least once)	Key	Value
<code>'codebook'</code>	<code>'codebooks'</code>	Dictionary containing a key/value pair per codebook-approximated parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is an int32 array specifying the codebook values.
<code>'codebook'</code>	<code>'codebook_zero_offsets'</code>	Dictionary containing a key/value pair per codebook-approximated parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is a int32 value specifying the element of the codebook that corresponds to integer index 0.
<code>'codebook'</code>	<code>'codebooks_egk'</code>	Dictionary containing a key/value pair per codebook-approximated parameter tensor of the model. Each key is a string containing the name of the parameter tensor and the corresponding value is an int32 value specifying syntax element <code>codebook_egk</code> .

The metadata specific to the approximation method shall contain all necessary information required for the later reconstruction process. This is to ensure that the subsequent reconstruction method will be able to entirely reconstruct the int32 parameter tensors of the network back to a float representation. The metadata may be shared across parameters (global), or may also be parameter-specific (local). For the latter, a respective mapping between the parameter name and the metadata should be provided. In order to ensure flexibility across different approximation methods, the choice of the specific format of the metadata is up to the particular method. However, the use of a dictionary structure is recommended in order to ensure consistency.

5.5.3 `nnc_core` – Main module

5.5.3.1 General

This module provides the submodules `approximator`, `coder` and `approx_coder`.

5.5.3.2 `nnc_core.approximator` – Parameter approximation module

This module provides an interface definition for parameter approximation and reconstruction.

```
nnc_core.approximator.approx(approx_info, model_info, approx_data)
```

The variable `model_info` is a dictionary according to [5.5.1](#).

Variable `approx_info` is a dictionary containing the key/value pairs specified in [Table 8](#).

Table 8 — Key/value pairs in `approx_info`

Key	Value
<code>'approx_method'</code>	String specifying the parameter approximation method. Allowed values are <i>'uniform'</i> and <i>'codebook'</i> .
<code>'to_approximate'</code>	A list of strings identifying the parameter types that are approximated.
<code>'dq_flag'</code>	Dictionary containing a key/value pair per parameter tensor of the model that are approximated. Each key is a string containing the name of the parameter tensor and the corresponding value is flag indicating whether dependent quantization is used.
<code>'lambda_scale'</code>	A float value specifying the lambda scale value to be used in dependent quantization.
<code>'cabac_unary_length_minus1'</code>	An int32 value specifying the syntax element <i>cabac_unary_length_minus1</i> .
<code>'integer_aligned_bitdepth'</code>	An uint8 value ≥ 2 and ≤ 32 . Specifies a bit depth for limited precision conformance testing. Typical values: 4 (bit) or 8 (bit).

The dictionary *approx_info* shall contain further keys value pairs depending on the value *approx_info['approx_method']* according to [Table 9](#).

Table 9 — Key/value pairs in `approx_info['approx_method']`

<code>approx_method</code>	Key	Value
<i>'uniform'</i>	<i>'qp'</i>	Dictionary containing a key/value pair per parameter tensor of the model that is approximated. Each key is a string containing the name of the parameter tensor and the corresponding value is an int32 value specifying the quantisation parameter.
<i>'codebook'</i>	<i>'qMse'</i>	A float value specifying the maximum mean squared error between the quantized tensors and the original tensors. If the quantization algorithm fails to achieve this MSE for a tensor using the maximum quantization bits specified by <i>qMax</i> , then the tensor is not quantized.
<i>'codebook'</i>	<i>'cb_size_ratio'</i>	A float value specifying the minimum ratio of tensor to codebook size.
<i>'uniform_sbt'</i>	<i>'sbt_args'</i>	Dictionary containing parameters if the chosen method is <i>uniform_sbt</i> .

If *uniform_sbt* is used, the dictionary shall contain the values according to [Table 10](#).

Table 10 — Key/value pairs when *uniform_sbt* is used

Key	Value
<code>'epoch'</code>	Number of the current epoch.
<code>'n_epoch'</code>	Total number of epochs.
<code>'sbt_mean_quant'</code>	A binary flag to activate mean quantization.
<code>'sbt_subquantization_approach'</code>	A binary flag that determines whether or not the batchnorm quantization shall be activated.
<code>'sbt_bn_quant'</code>	A variable that determines which quantization mode shall be applied. Possible values are: <i>'combined'</i> (switch between binary and ternary), <i>'binary'</i> (only binary quantization is applied), <i>'ternary'</i> (only ternary quantization is applied).
<code>'sbt_quantization_approach'</code>	If the quantization mode is combined, this variable determines how binary/ternary switching will happened. Possible values are <i>'random'</i> and <i>'random_round'</i> .

The variable *approx_data* is a dictionary according to [5.5.2](#).

Returns a dictionary according to [5.5.2](#). The approximation methods are specified in clause 8 of ISO/IEC 15938-17.

```
nnc_core.approximator.rec(approx_data, model_info)
```

The variable *model_info* is a dictionary according to [5.5.1](#).

The variable *approx_data* is a dictionary according to [5.5.2](#).

Quantized values in variable *approx_data* are replaced with reconstructed versions. The employed reconstruction methods are specified in 7.3 of ISO/IEC 15938-17.

```
nnc_core.approximator.fold_bn(model_info, approx_data, approx_info)
```

The variable *model_info* is a dictionary according to [5.5.1](#).

The variable *approx_data* is a dictionary according to [5.5.2](#).

Variable *approx_info* is a dictionary containing the key/value pairs specified in [Table 7](#).

The four batch-norm parameters of each batch-norm layer in *approx_data* are replaced with only two parameters, a scaling and a bias parameter. The two parameters are obtained by merging the batch-norm parameters as specified in 8.2.6 of ISO/IEC 15938-17.

```
nnc_core.approximator.unfold_bn(model_info, approx_data)
```

The variable *model_info* is a dictionary according to [5.5.1](#).

The variable *approx_data* is a dictionary according to [5.5.2](#).

Replaces scaling and bias parameters in *approx_data* with recreated equivalent batch-norm parameters. The recreation process is specified in 8.2.6 of ISO/IEC 15938-17.

```
nnc_core.approximator.set_lsa(model_info, approx_data, lsa_params)
```

The variable *model_info* is a dictionary according to [5.5.1](#).

The variable *approx_data* is a dictionary according to [5.5.2](#).

The variable *lsa_params* is a dictionary containing a key/value pair per parameter tensor of the model for which scaling parameters are specified. Each key is a string containing the name of the parameter tensor the scaling parameters are applied to and the corresponding value is a 1D tensor. In the case of a float parameter, it contains float values and in the case of an approximated parameter, it contains int32 values.

Adds scaling parameters for each tensor specified by *lsa_params* to *approx_data* as specified in 8.2.6 of ISO/IEC 15938-17. The reference software provides an algorithm for deriving values for *lsa_params*, although it is out of scope of the standard.

```
nnc_core.approximator.apply_lsa(model_info, approx_data)
```

The variable *model_info* is a dictionary according to [5.5.1](#).

The variable *approx_data* is a dictionary according to [5.5.2](#).

Applies the scaling parameters to each tensor specified and replaces the parameter tensors in *approx_data* with the scaled versions. Then removes the scaling parameters from *approx_data*. The process is specified in 8.2.6 of ISO/IEC 15938-17.

5.5.3.3 nnc_core.coder – Entropy coding module

This module provides access to lossless entropy encoding and decoding functions.

```
nnc_core.coder.encode(enc_info, model_info, approx_data)
```

The variable *model_info* is a dictionary according to [5.5.1](#).

Variable *enc_info* is a dictionary containing the key/value pairs specified in [Table 11](#).

Table 11 — Key/value pairs in *enc_info*

Key	Value
'cabac_unary_length_minus1'	An int32 value specifying the syntax element cabac_unary_length_minus1.
'param_opt_flag'	A flag specifying whether cabac optimization is used.
'nnr_pt_block_enabled_flag'	If equal to 1, decoding method for NNR compressed payloads of type NNR_PT_BLOCK is enabled.
'partial_data_counter'	A uint8 value specifying the number of partitions for the payload of each ndu nnr unit. Directly sets the initial value of the syntax element <i>partial_data_counter</i> .
'node_id_present_flag'	Equal to 1 indicates that syntax elements device_id, parameter_id, and put_node_depth are present.
'max_ndu_nnr_unit_size'	An int32 value specifying the maximum size (in bytes) of a ndu nnr unit. The payloads of ndu nnr units are partitioned to achieve the desired size. Overrides the value of 'partial_data_counter'.
'device_id'	Uniquely identifies the device that generated the current NDU.
'mps_topology_indexed_reference_flag'	A Boolean value specifying if indexed references are used. Setting the flag to True enables element IDs to be coded using indices. If False, string-based labels are used instead.
'temporal_context_modeling_flag'	A flag specifying whether temporal context modelling is used.
'encode_tpl_only'	A Boolean specifying if model data is written into the bitstream. If True, only topology data is written and no approximation data etc. is written.
'row_skip_enabled_flag'	A flag specifying whether row skip is used.
'out_of_band_signaling'	A Boolean specifying if input parameters (i. e. <i>count_tensor_dimensions</i> , <i>tensor_dimensions</i> , <i>cabac_unary_length_minus1</i> , <i>compressed_parameter_types</i> , <i>decomposition_rank</i> and <i>g_number_of_rows</i>) are signaled out of band. If True, none of these parameters is written to the bitstream but to a dictionary which is stored as JSON file. If False, each parameter, if present, is written to the bitstream.
'parent_node_id_present_flag'	Equal to 1 indicates that syntax element parent_node_id_type is present.
'parent_node_id_type'	Indicates which further syntax elements for uniquely identifying the parent node are present. Allowed values are <i>ICNN_NDU_ID</i> =0, <i>ICNN_NDU_PL_SHA256</i> =1, <i>ICNN_NDU_PL_SHA256</i> =1 or 3 (Reserved).
'parent_device_id'	Is equal to syntax element device_id of the parent NDU.
'general_profile_idc'	Equal to 1 indicates that bitstream conforms to the extended compression tool set. If 0, bitstream conforms to the base feature set (as defined in ISO/IEC 15938-17:2024, sub-clause 6.4.3.2).
'bt_signaling'	If equal to 1, improved context modelling for binary and ternary data is enabled (e.g. in conjunction with SBT)
'mps_parent_signalling_enabled_flag'	Specifies whether parent node information (i.e. element <i>parent_node_id_present_flag</i>) is signalled at NDU level. If equal to 0, <i>parent_node_id_present_flag</i> is not present for any NDU (i.e. the bitstream represents a base neural network). If equal to 1, the bitstream represents a differential update of a base neural network.

The variable *approx_data* is a dictionary according to [5.5.2](#).