

---

**cONNXr**  
*Release 0.0.0*

**Mar 11, 2021**



<b>1</b>	<b>Out of the box example</b>	<b>3</b>
1.1	Documentation . . . . .	3
1.2	Contributing . . . . .	11
1.3	Testing . . . . .	14
1.4	Onnx Operator Status . . . . .	15



---

**Note:** This project is in a very early stage and we are looking for contributors.

---

A onnx runtime written in pure C99 with zero dependencies focused on embedded devices. Run inference on your machine learning models no matter which framework you train it with and no matter the device that you use. This is the perfect way to go in old hardware that doesn't support fancy modern C or C++.



---

## Out of the box example

---

Estimate the number written in an image using the MNIST model. The input image is stored in *input\_0.pb* and the MNIST model is stored in onnx format in *model.onnx*.

First compile the project.

```
make all
```

And run inference on the input using the model.

```
build/connxr test/mnist/model.onnx test/mnist/test_data_set_0/input_0.pb
```

## 1.1 Documentation

### 1.1.1 Introduction

#### What is ONNX?

If you don't know about `onnx` you might want to read about it before since its the building block of this project. They have a nice [website](#) and great repositories with a lot of documentation to read about. Everything is open source, and really big companies in the industry are behind it (AMD, ARM, AWS, Nvidia, IBM) just to name a few. Here you can find a mix of official and non official related repositories:

- <https://github.com/onnx/onnx>
- <https://github.com/onnx/onnx-r>
- <https://github.com/onnx/models>
- <https://github.com/owulveryck/onnx-go>
- <https://github.com/microsoft/onnxruntime>

In short, `onnx` provides a **O**pen **N**eural **N**etwork **E**xchange format. This format, describes a huge set of operators, that can be mixed to create every type of machine learning model that you ever heard of, from a simple neural network to complex deep convolutional networks. Some examples of operators are: matrix multiplications, convolutions, adding, maxpool, sin, cosine. They provide a standardised set of operators [here](#). So we can say that `onnx` provides a layer of abstraction to ML models, which makes all frameworks compatible between them. Exporters are provided for a huge variety of frameworks (PyTorch, TensorFlow, Keras, Scikit-Learn) so if you want to convert a model from Keras to TensorFlow, you just have to use Keras exporter to export `Keras->ONNX` and then use the importer to import `ONNX-TensorFlow`.

In the following image, you can find an example on how a `onnx` model looks like. Its just a bunch of nodes that are connected between them to form a graph. Each node has an operator that takes some inputs with some specific dimensions and some attributes and calculates some outputs. This is how the inference is calculated, just forward propagating the input along every node until the last one is reached.





So this `.onnx` format can describe machine learning models through a set of nodes that contain specific operators connected among them. On top of that, they also offer a set of the so called “runtimes” [see link](#). A runtime allows to run inference on a model, and they offer different ones in a wide variety of languages and hardware. Unfortunately, all the runtimes rely on modern versions of C/C++ with many abstractions and dependancies, which might be a no go for some specific projects. Here is where we come in.

## What is cONNXr

Well, now that you know about `onnx`, our project is just a runtime that runs inference on `onnx` models. The `c` means that is implemented in C language and the `r` means that its a runtime for `ONNX`. The only difference between this runtime and the others, is that this one is written in pure C99 without any dependancy. This means that it should be able to compile with almost any compiler, no matter how old it is. Our goal is to enable embedded devices that doesn't have much resources or fancy features (like GPUs or any type of hardware accelerator) to run inference. No GPUs, no multithreading, no dependancies, just pure C code with the lowest possible footprint. Train your model in whatever ML framework you want, export it to `.onnx` and deploy it wherever you want.

You might also find this project useful if you work with some bare metal hardware with dedicated accelerators. If this is the case, you might find useful to reuse the architecture and replace the specific operators (functions) by your own ones.

Or perhaps you have a different use case. If this is the case, we would love to hear about it.

### 1.1.2 File structure

The most relevant files and folders are the following ones:

- `include/src`: Source and header C files.
- `include/operators/onnx/`: Autogenerated header files for all operators and versions. There is one header file per operator, and within each file there is one function declaration for each data type. For example, for the operator `conv` version 11 there is one header file, that contains 3 function declarations (double, float and float16). Note that these functions are defined as extern and not implemented.
- `src/operators/implementation`: Here is where all the operators implementations are located. The functions that are defined here must use the prototype defined in `include/operators/onnx` mentioned above. See some examples.
- `src/operators/resolve`: This is autogenerated code, that resolves the mapping between an operator + data type and the function that should be executed.
- `src/operators/operator_set.c`: Autogenerated file that stores the relationship between an operator name (i.e. `Conv`) and the function that resolves its data type. Operator version is also taken into account.
- `scripts/onnx_generator`: Set of Python scripts that generates the header and C files existing in `include/operators/onnx`, `src/operators/resolve` and `src/operators/operator_set.c`.
- `src/test` and `include/test`: Test code divided into two levels: operator level and model level.
- `test`: Test vectors on operator and model level. Bunch of `.onnx` models and `.pb` files (input + expected output)

### 1.1.3 Operators interface

All `ONNX` operators comply with the following interface. This struct contains all the data that an operator needs to run, like its inputs and attributes. In `operator_executer` it contains the pointer to the function that runs that

specific node. For example, if node  $i$  contains a Conv, that `operator_executer` will point to the `conv` function defined in `src/operators/implementation`.

You can check more information about `Onnx__NodeProto` and `Onnx__TensorProto` structs in `onnx.pb-c.h`, which is a file that is generated from the ONNX common interface.

```
struct node_context {
    Onnx__NodeProto    *onnx_node;
    Onnx__TensorProto **inputs;
    Onnx__TensorProto **outputs;
    operator_executer resolved_op;
};
```

## 1.1.4 Types and Structures

We can divide the types and structures that this repo uses into two:

- ONNX structures: These structures are exactly (or almost) the same as the ones that ONNX defines in its `onnx.proto`. See `onnx.pb-c.h`. We list the most relevant ones below.
- Custom structures: These structures are defined *ad hoc* for this project.

### ONNX structures

#### Onnx\_\_GraphProto

Defines the graph of a model, made of different nodes `Onnx__NodeProto`. See `onnx.pb-c.h`.

```
struct __Onnx__GraphProto
{
    ProtobufCMessage base;
    size_t n_node;
    Onnx__NodeProto **node;
    char *name;
    size_t n_initializer;
    Onnx__TensorProto **initializer;
    size_t n_sparse_initializer;
    Onnx__SparseTensorProto **sparse_initializer;
    char *doc_string;
    size_t n_input;
    Onnx__ValueInfoProto **input;
    size_t n_output;
    Onnx__ValueInfoProto **output;
    size_t n_value_info;
    Onnx__ValueInfoProto **value_info;
    size_t n_quantization_annotation;
    Onnx__TensorAnnotation **quantization_annotation;
};
```

#### Onnx\_\_NodeProto

Information for a given node (a node has inputs and attributes and provides an output using a given operator). See `onnx.pb-c.h`.

```

struct  _Onnx__NodeProto
{
  ProtobufCMessage base;
  size_t n_input;
  char **input;
  size_t n_output;
  char **output;
  char *name;
  char *op_type;
  char *domain;
  size_t n_attribute;
  Onnx__AttributeProto **attribute;
  char *doc_string;
};

```

## Onnx\_\_TensorProto

One of the most important types that you will see, is the `Onnx__TensorProto`. It just defines a vector, an array, a matrix, or whatever you want to call it. It is quite convenient to use, because it is quite generic. You can store different types of values, with different sizes. As an example, let's say that we want to store a 3 dimension vector. In that case `n_dims=3` and `dims[0]`, `dims[1]`, `dims[2]` will store some values. Let's store some float values, so `data_type=ONNX__TENSOR_PROTO__DATA_TYPE__FLOAT`. In this case `n_float_data=dims[0]*dims[1]*dims[2]` and `float_data` will contain all the values in a single dimension array. The tensor has also a name.

```

struct  _Onnx__TensorProto
{
  ProtobufCMessage base;
  size_t n_dims;
  int64_t *dims;

  protobuf_c_boolean has_data_type;
  int32_t data_type;
  Onnx__TensorProto__Segment *segment;

  size_t n_float_data;
  float *float_data;

  size_t n_int32_data;
  int32_t *int32_data;

  size_t n_string_data;
  ProtobufCBinaryData *string_data;

  size_t n_int64_data;
  int64_t *int64_data;

  char *name;
  char *doc_string;

  protobuf_c_boolean has_raw_data;
  ProtobufCBinaryData raw_data;

  size_t n_external_data;
  Onnx__StringStringEntryProto **external_data;
};

```

(continues on next page)

(continued from previous page)

```

protobuf_c_boolean has_data_location;
Onnx__TensorProto__DataLocation data_location;

size_t n_double_data;
double *double_data;

size_t n_uint64_data;
uint64_t *uint64_data;
};

```

## Custom structures

Since we are at a very early stage, structures are constantly evolving. Have a look to the code in order to know more. At some point we will also document them here.

### 1.1.5 Protocol Buffers

In order to convert from the ONNX interface defined in `onnx.proto`, the `protoc` library is used. `onnx` uses protocol buffers to serialize the models data. Note that `protobuf-c` is used to generate the `pb/onnx.pb-c.c` and `pb/onnx.pb-c.h`. Files are already provided, but you can generate it like this:

```
protoc --c_out=. onnx.proto
```

### 1.1.6 Autogenerated code

In order to avoid boilerplate code, some `.c` and `.h` files are autogenerated from Python. You can find these Python scripts and more information under `scripts/onnx_generator` folder. These Python scripts take into account a given ONNX version and generates all the “infrastructure” from it.

### 1.1.7 Versioning

ONNX is constantly evolving and new operators and versions are regularly being released. This project aims to keep up with that using the autogenerated code mentioned above.

### 1.1.8 Operators

All `onnx` operators are defined in the [official doc](#). Our goal is to implement as many as possible, but if you have a model with a custom operator, that also fine.

All operators share a common interface. This structure is available as an input parameter to all operators, and it contains all the data you need, like the input tensors and attributes and a pointer to the output.

```

struct node_context{
    Onnx__NodeProto    *onnx_node;
    Onnx__TensorProto **inputs;
    Onnx__TensorProto **outputs;
    operator_executer resolved_op;
};

```

Lets say that you want to implement the Add operator, that just adds two numbers or tensors. You can access to its inputs with the following code. Once you have that, its pretty straightforward to access the data within the tensor (see `Onnx__TensorProto` struct)

```
Onnx__TensorProto *A = searchInputByName(ctx, 0);
Onnx__TensorProto *B = searchInputByName(ctx, 1);
```

On the other hand you will need to store the result in a variable, so that other nodes can reuse that output. Just use the following function and populate the content.

```
Onnx__TensorProto *C = searchOutputByName(ctx, 0);
```

If the operator you are implementing has some attributes, you can also easily get them with. Just replace `auto_pad` by your attribute name.

```
Onnx__AttributeProto *auto_pad = searchAttributeNyName (
    ctx->onnx_node->n_attribute,
    ctx->onnx_node->attribute,
    "auto_pad");
```

Its also important to note that each operator is usually defined for more than one data type like float, double or uint8\_t. In order to address this, we have decided to define one function per data type, trying of course to share as many code as possible between different data types implementations.

### 1.1.9 Tracing

We utilize macros to enable fine tracing of components. Trace macros can act as asserts and therefore may abort execution if an erroneous state is detected.

**If not explicitly enabled, all tracing components are stripped before compilation!**

Trace prints look like this:

```
[MACRO LEVEL] FILE:LINE SCOPE MSG
```

- MACRO describes which macro produced the print
- LEVEL describes the minimum `TRACE_LEVEL` needed to generate this print
- FILE is the file in which the macro was called
- LINE is the line number in which the macro was called
- SCOPE is an optional block description
- MSG is the actual logged message

To enable tracing you need to specify a default `TRACE_LEVEL` starting from 0.

```
make clean all TRACE_LEVEL=0
```

|TRACE\_LEVEL|Meaning |-----:|----- | undefined|disable tracing | 0|no prints, except for errors | 1|execution flow related prints (function entry/exists, decisions) | 2|data flow related prints (input/output meta data) | 3+|internal state of algorithms

The `TRACING_LEVEL` can be overwritten at file and function granularity without recompilation. When tracing is enabled, each macro checks the `CONNXR_TRACE_FILE` and `CONNXR_TRACE_FUNCTION` env variable for overrides.

Overrides are specified in following format:

```
CONNXR_TRACE_FILE=<filename>:<trace_level>{;<filename>:<trace_level>}
CONNXR_TRACE_FUNCTION=<function>:<trace_level>{;<function>:<trace_level>}
```

So you may execute

```
make clean all TRACE_LEVEL=0
CONNXR_TRACE_FUNCTION=operator__onnx__maxpool__12__T_tensor_float:3 make test_
↪operators
```

to only generate a detailed trace for a specific operator (maxpool in this case).

## 1.2 Contributing

This project is in very early state so we are looking for contributors. Feel free to open a PR or an issue with questions or suggestions. We would be also happy to guide you if you have an idea. You can contribute in many ways, but the most common is to implement a new operator. The `onnx` specification defines more than 150 operators, that are supported for different types (float, int, double, ...) and for some of them there are different versions. Some common ones are implemented with limited functionality like convolutions, matrix multiplications and so on, but there are lots of them remaining.

You can also contribute by improving an existing operator or fixing some bugs. Feel free to have a look to the opened issues, where we also have some simple issues for newcomers.

### 1.2.1 Add new operator

If you want to add a new operator, we provide a simple and generic interface that you can use, with all the information in place ready to be used. The interface is the following. You can check [here](#) more information about each struct. In `onnx_node`, `inputs` and `outputs` you will have all the information that you need such as the inputs of the operators with its names, the attributes and where to write the output. Refer to [ONNX Operators](#) to check the inputs/outputs/attributes of the operator you are implementing. Note that some might be optional.

```
struct node_context{
    Onnx__NodeProto    *onnx_node;
    Onnx__TensorProto **inputs;
    Onnx__TensorProto **outputs;
    operator_executer resolved_op;
};
```

First of all decide the operator that you want to implement. You can see the [official onnx operators list](#). Lets say that you want to implement the `Abs` operator, fair enough.

Since each operator has different versions, you would need to choose the version that you are implementing. The differences can be rather tiny, but its not the same `Abs-13` operator than `Abs-6`. Note that the number is the onnx version that introduced that operator. We recommend implementing the latest one available at the moment.

Once you know the operator and version, there is one extra “dimension” to take into account. Most of the operators work with different data types, so you must choose which one to implement. Going back to the `Conv` example, we would suggest to start implementing the `float` version. Note that `Conv` is defined for double, float and float16. The code would be of course quite similar, but since we are in C we can’t just use the same functions for all types.

Lets see how you can implement a new operator in few simple steps.

## 1. Generate the files

You don't have to create any files, just do a small modification in the `Makefile` and run a script. The files that you need will be generated automatically:

- If you want to implement a new operator, go to the makefile and add a new line with the operator that you want to generate, i.e. `ONNX_INCLUDE+='^Add$$'`.
- Once you have that, and assuming that you have the latest onnx Python version installed, run `make onnx_generator`.
- The previous step will generate all files that are needed. It will also update the `operator_sets.c` file and create a resolver for that operator. You don't really need to care about this.
- Now you can populate the `.c` files with your implementation. Note that there is one implementation per data type (i.e. float, double,...)

Lets see some of the files that Add operator has:

- `operator__onnx__add__7__T_tensor_double.c`: The 7 indicates the operator version and double indicates the type. Add operator will have other files such as `_int32` or `_int64`. The implementations would be different, but they can share most of the code. You have of course to write in these file the actual implementation of the operator.
- `resolve_operator__onnx__add__7.c`: You don't need to touch this file. It just maps (resolves) the function that an operator needs based on the input time.

## 2. Implement the operator

Now that you have generated all the files you are ready to start implementing the operator. Lets say that you want to implement `operator__onnx__add__7__T_tensor_float.c`, which is the Add operator for opset version 11 and type float.

This operator is adding two values or tensors so first of all you need these values. According to [the specification](#) this operator takes two inputs A and B with no attributes. There are two different ways that you can access them: by index or by name.

### Inputs

You can **access by index** as follows. Note that `ctx->onnx_node->input[n]` doesn't contain the tensor, but just a name of the tensor.

```
// Access by index
Onnx__TensorProto *A = ctx->inputs[0];
Onnx__TensorProto *B = ctx->inputs[1];
```

You can also **access by name** with `searchInputByName()`. This function takes an index with the name of the tensor that wants to be accessed.

```
// Access by name
Onnx__TensorProto *A = searchInputByName(ctx, 0);
Onnx__TensorProto *B = searchInputByName(ctx, 1);
```

Both ways are perfectly valid, but the **access by index** is preferred since it doesn't need to search. But note that in some cases it can't be used. Imagine an operator with two inputs, but the second one is optional. Using `ctx->inputs[1]` will only work if the second input is provided, and will fail when it doesn't. For this cases, use `searchInputByName()` that will return `NULL`.



**Note:** There is some ongoing work in here, so might change the way it works.

At this point you have A and B ready to be added. Both variables belong to `Onnx__TensorProto` so feel free to have a look [here](#). You can access the elements of the tensor like this, assuming that the type stored in it is `float`.

```
for (int i = 0; i < A->n_float_data) {
    //A->float_data[i]
}
```

## Attributes

If the operator you are implementing has some attributes, you can also easily get them with. Just replace `auto_pad` by your attribute name.

```
Onnx__AttributeProto *auto_pad = searchAttributeByName(
    ctx->onnx_node->n_attribute,
    ctx->onnx_node->attribute,
    "auto_pad");
```

Same than before applies here. You can access the attributes **by name** or **by index**

Access by index:

- If there is only 1 attribute, and that attribute is mandatory, this way can be used. Not the case of `LeakyRelu` because that attribute could be empty (and the default value will be taken).
- If there are more than 1 attribute but all of them are mandatory, we can also use this way. If some attributes are mandatory this can't be done.

```
// Access by index
Onnx__AttributeProto *a_alpha = ctx->onnx_node->attribute[0]->f;
```

And this other way can be used in the rest of the cases.

```
// Access by name
Onnx__AttributeProto *a_alpha = searchAttributeByName(ctx->onnx_node->n_attribute, ctx->
    onnx_node->attribute, "alpha");
if (a_alpha) {
    alpha = a_alpha->f;
}
```

## Outputs

Last but not least, you will need to store the result in a variable, so that other nodes can reuse that output. Just use the following function and populate the content.

You can do it also **by index**. This is the way to go if there if only one output that is mandatory.

```
Onnx__TensorProto *C = ctx->outputs[0];
```

And in more complex cases where you can optional outputs or more than one, you can use the following.

```
Onnx__TensorProto *C = searchOutputByName(ctx, 0);
```

### 3. Test the operator

Once the operator is implemented, you are ready to test it. Luckily, `onnx` provides a set of `test vectors` for all operators, and we have taken care of integrating them, so you just need to go to `tests.c` file and uncomment the test case of your operator. For example, if you have implemented the `Shrink` operator, you should uncomment `test_shrink_hard` and `test_shrink_hard`. Its important to note that this test cases are not testing all data types (float, double,...). Most of these test cases run on float, but not always.

Currently `onnx` only provides test cases for the latest operator, so if you want to test and old version of an operator, you will have to do some manual work. However, there is some ongoing work in [here](#) to address this issue.

You can have a look to [this PR](#) that contains an example of how an operator can be implemented.

#### 1.2.2 Other contributions

You can also contribute in other ways, like improving an existing operator, fixing bugs or writing documentation.

## 1.3 Testing

### 1.3.1 Testing

Everything is tested using Python `unittest`, that runs the C code under the hood. Note that `onnx` provides data for testing a backend implementation. Have a look to <https://github.com/onnx/onnx/blob/master/docs/OnnxBackendTest.md> and <https://github.com/onnx/onnx/tree/master/onnx/backend/test/data/node>. In this link you will find sets of inputs/outputs + a `model.onnx` that allows to test both an individual operator or a whole model, hence we divide two different types of tests:

- **Operator Tests:** Tests an individual operator. Lets see a simple example. We want to test the operator `Sum` over a one dimension array. So we provide a set of two inputs  $X1 = [1 \ 1]$  and  $X2 = [2 \ 2]$  and an expected output  $O = [3 \ 3]$ . We also provides a model with just one two inputs an a node which operator is `Sum`. All these inputs/outputs
- **Model Tests:** Runs inference using a model. As an example we can use MNIST model (a digit recognition library). The input is an image with a written digit and the output is an int  $[0-9]$ . This type of tests run several operators.

#### 1.3.2 Operator Tests

You can run all operator tests with the following command. This command just uses `tests/test_operators.py` script under the hood.

```
make test_operators
```

You can run a single test as follows:

```
python3 tests/test_operators.py TestOperators.test_relu
```

#### 1.3.3 Model Tests

Different end to end models are implemented and tested, and you can run them all with the following command (make sure you can compiled before `make all`). Note that this uses `tests/test_models.py` under the hood.

```
make test_models
```

You can also run a single model with the following command:

```
python3 tests/test_models.py TestModels.test_mnist
```

### 1.3.4 Debugging tools

If you are implementing a new model but you can get the end to end test cases to pass, you might want to use some of the tools that we have. Using the CLI, you can provide the option `--dump-file`. This will dump all the intermediate outputs to a file, that you can use later on to compare against the expected values.

The `dump.txt` file will look something like this.

```
name=output_name
shape=shape1,shape2,...
tensor=value1,value2,value3,...
name=output_name
shape=shape1,shape2,...
tensor=value1,value2,value3,...
...
```

So in this file you have all the data that cONNXr has calculated for every node. Now you can compare this against the expected outputs of that model. For this purpose, we have `scripts/assert_nodes.py`. This script gets the output tensor that cONNXr calculated, and compares it against the one calculated by the official ONNX backend. This is a great tool if you want to find the node that is causing problems. This script will point to the node that is not matching the expected values.

Note that this feature is under development.

Example how to dump data

```
build/connxr test/mobilenetv2-1.0/mobilenetv2-1.0.onnx test/mobilenetv2-1.0/test_data_
↪set_0/input_0.pb --dump-file
```

Example how to use the Python script:

TODO

```
python3 scripts/assert_nodes.py
```

## 1.4 Onnx Operator Status

### 1.4.1 Onnx Operator Status

This will be generated using the tests as reference.

- ai.onnx
  - not\_implemented Abs
  - not\_implemented Acos
  - not\_implemented Acosh
  - implemented Add

- not\_implemented And
- implemented ArgMax
- not\_implemented ArgMin
- not\_implemented Asin
- not\_implemented Asinh
- not\_implemented Atan
- not\_implemented Atanh
- not\_implemented AveragePool
- implemented BatchNormalization
- not\_implemented BitShift
- implemented Cast
- not\_implemented Ceil
- implemented Clip
- not\_implemented Compress
- not\_implemented Concat
- not\_implemented ConcatFromSequence
- not\_implemented Constant
- not\_implemented ConstantOfShape
- implemented Conv
- implemented ConvInteger
- not\_implemented ConvTranspose
- not\_implemented Cos
- not\_implemented Cosh
- not\_implemented CumSum
- not\_implemented DepthToSpace
- not\_implemented DequantizeLinear
- not\_implemented Det
- not\_implemented Div
- not\_implemented Dropout
- not\_implemented Einsum
- implemented Elu
- not\_implemented Equal
- not\_implemented Erf
- not\_implemented Exp
- not\_implemented Expand
- not\_implemented EyeLike

- not\_implemented Flatten
- not\_implemented Floor
- not\_implemented GRU
- not\_implemented Gather
- not\_implemented GatherElements
- not\_implemented GatherND
- not\_implemented Gemm
- implemented GlobalAveragePool
- not\_implemented GlobalLpPool
- not\_implemented GlobalMaxPool
- not\_implemented Greater
- not\_implemented HardSigmoid
- not\_implemented Hardmax
- implemented Identity
- not\_implemented If
- not\_implemented InstanceNormalization
- not\_implemented Inverse
- not\_implemented IsInf
- not\_implemented IsNaN
- not\_implemented LRN
- not\_implemented LSTM
- implemented LeakyRelu
- not\_implemented Less
- not\_implemented Log
- not\_implemented LogSoftmax
- not\_implemented Loop
- not\_implemented LpNormalization
- not\_implemented LpPool
- implemented MatMul
- implemented MatMulInteger
- not\_implemented Max
- implemented MaxPool
- not\_implemented MaxRoiPool
- not\_implemented MaxUnpool
- not\_implemented Mean
- not\_implemented Min

- not\_implemented Mod
- implemented Mul
- not\_implemented Multinomial
- not\_implemented Neg
- not\_implemented NonMaxSuppression
- not\_implemented NonZero
- not\_implemented Not
- not\_implemented OneHot
- not\_implemented Or
- not\_implemented PRelu
- not\_implemented Pad
- not\_implemented Pow
- not\_implemented QLinearConv
- not\_implemented QLinearMatMul
- implemented QuantizeLinear
- not\_implemented RNN
- not\_implemented RandomNormal
- not\_implemented RandomNormalLike
- not\_implemented RandomUniform
- not\_implemented RandomUniformLike
- not\_implemented Reciprocal
- not\_implemented ReduceL1
- not\_implemented ReduceL2
- not\_implemented ReduceLogSum
- not\_implemented ReduceLogSumExp
- not\_implemented ReduceMax
- not\_implemented ReduceMean
- not\_implemented ReduceMin
- not\_implemented ReduceProd
- not\_implemented ReduceSum
- not\_implemented ReduceSumSquare
- implemented Relu
- implemented Reshape
- not\_implemented Resize
- not\_implemented ReverseSequence
- not\_implemented RoiAlign

- not\_implemented Round
- not\_implemented Scan
- not\_implemented Scatter
- not\_implemented ScatterElements
- not\_implemented ScatterND
- not\_implemented Selu
- not\_implemented SequenceAt
- not\_implemented SequenceConstruct
- not\_implemented SequenceEmpty
- not\_implemented SequenceErase
- not\_implemented SequenceInsert
- not\_implemented SequenceLength
- not\_implemented Shape
- not\_implemented Shrink
- stubbed Sigmoid
- not\_implemented Sign
- not\_implemented Sin
- not\_implemented Sinh
- not\_implemented Size
- not\_implemented Slice
- implemented Softmax
- not\_implemented Softplus
- not\_implemented Softsign
- not\_implemented SpaceToDepth
- not\_implemented Split
- not\_implemented SplitToSequence
- not\_implemented Sqrt
- not\_implemented Squeeze
- not\_implemented StringNormalizer
- not\_implemented Sub
- not\_implemented Sum
- not\_implemented Tan
- not\_implemented Tanh
- not\_implemented TfIdfVectorizer
- not\_implemented ThresholdedRelu
- not\_implemented Tile

- not\_implemented TopK
- not\_implemented Transpose
- not\_implemented UnfoldToDepth
- not\_implemented Unique
- not\_implemented Unsqueeze
- not\_implemented Upsample
- not\_implemented Where
- not\_implemented Xor
- Functions
  - \* not\_implemented Celu
  - \* not\_implemented DynamicQuantizeLinear
  - \* not\_implemented GreaterOrEqual
  - \* not\_implemented LessOrEqual
  - \* not\_implemented MeanSquaredDistance
  - \* not\_implemented MeanVarianceNormalization
  - \* not\_implemented NegativeLogLikelihoodLoss
  - \* not\_implemented Range
  - \* not\_implemented SoftmaxCrossEntropyLoss
- ai.onnx.training
  - not\_implemented Adagrad
  - not\_implemented Gradient
  - not\_implemented GraphCall
  - not\_implemented Momentum
- ai.onnx.ml
  - not\_implemented ArrayFeatureExtractor
  - not\_implemented Binarizer
  - not\_implemented CastMap
  - not\_implemented CategoryMapper
  - not\_implemented DictVectorizer
  - not\_implemented FeatureVectorizer
  - not\_implemented Imputer
  - not\_implemented LabelEncoder
  - not\_implemented LinearClassifier
  - not\_implemented LinearRegressor
  - not\_implemented Normalizer
  - not\_implemented OneHotEncoder



- not\_implemented SVMClassifier
- not\_implemented SVMRegressor
- not\_implemented Scaler
- not\_implemented TreeEnsembleClassifier
- not\_implemented TreeEnsembleRegressor
- stubbed ZipMap