

intelligence; therefore, it is not strange to talk about →*Neural Machine Translation* as an artificial intelligence application.

Artificial Neural Network	A network of → <i>Artificial Neurons</i> , connected in a specific way to accomplish a specific computational task. Connections have strengths or <i>weights</i> that can be <i>learned</i> and regulate how the state or <i>activation</i> of an artificial neuron is affected by the <i>activations</i> of neighboring neurons (or the value of external inputs), so that the whole network has the desired behaviour.
Artificial Neuron	An artificial neuron (vaguely inspired on a natural neuron) is a very simple computational units. Its state or <i>activation</i> depends on the stimuli received from other artificial neurons (or external inputs) through its connections.
Attention-Based Neural Machine Translation	A kind of → <i>Neural Machine Translation</i> using a <i>learned</i> mechanism called <i>attention</i> to obtain information from selected parts of the source sentence every time a word of the target sentence is produced.
Black-Box approach	According to Wikipedia, "a black box is a device, system or object which can be viewed in terms of its inputs and outputs [...], without any knowledge of its internal workings. For instance, using an MT system in a black-box approach means that one only has access to the output it produces for the source text it is fed. See → <i>Glass-Box Approach</i> .
Controlled Language	A subset of a natural language obtained by restricting its grammar and vocabulary in order to make it easier to process by computers; for example, by removing ambiguities and other natural-language features that make machine translation difficult.
Convolutional encoder	A specific kind of → <i>Encoder</i> using a → <i>Convolutional neural network</i> .
Convolutional neural network	When the same artificial neural network is repeatedly applied to every position (for instance, every word) in the input data and its neighbouring positions (for example, neighbouring words) to obtain a new level of context-informed representation for each position. The resulting representations may be processed again in a similar way to generate another level of <i>deeper</i> representations (hence <i>deep learning</i>) of the original data.

Corpus-based machine translation	Machine translation that uses knowledge <i>learned</i> from very large parallel corpora containing source-language segments and their target-language translations. Examples: → <i>Statistical Machine Translation</i> , → <i>Neural Machine Translation</i> . As opposed to → <i>Rule-Based Machine Translation</i> .
Decoder	In → <i>Neural Machine Translation</i> , an → <i>Artificial Neural Network</i> that computes the likelihood that each word in the vocabulary is the next word in the target segment, by looking at the preceding words in the target segments, at one or more <i>representations</i> of the previous steps it has taken, and, optionally, using <i>attention</i> (see → <i>Attention-based Neural Machine Translation</i>), at the representations produced by the → <i>Encoder</i> at each position of the source segment.
Encoder	In → <i>Neural Machine Translation</i> , an → <i>Artificial Neural Network</i> that reads a source segment and <i>encodes</i> it into a series of → <i>Representations</i> which will later be processed by a → <i>Decoder</i> to generate a target segment.
Entropy	In information theory, a measure of the amount of information that is missing before a message is received, therefore, a measure of <i>uncertainty</i> . When the entropy is measured in <i>bits</i> , it is the minimum number of yes/no questions needed to fully determine (and therefore encode) the contents of the message. Entropy therefore relates to the likelihood of every possible message: when all messages are equally likely, the entropy is maximum. When a single message is the only one that is likely to appear, entropy approaches zero.
Glass-Box approach	A <i>glass box</i> is a device, system or object for which one does not only have access to the output it produces for the inputs it is fed, but one whose internals (such as intermediate results) can be viewed (but usually not altered). For instance, using an MT system in a glass-box approach means that one, in addition to the output it produces for the source text it is fed, has access to internal or intermediate stages of how it has been computed. See → <i>Black-Box Approach</i> .
Interactive Translation Prediction	Also called <i>Interactive Machine Translation</i> , a machine translation system that, instead of producing in one go the translation of a whole source segment, provides professional translators with machine-translated suggestions on how to continue the target segment they are typing.

Language model	A language model (usually a <i>probabilistic</i> language model) assigns a likelihood or score to a sequence of words, usually a sentence. They are commonly used in → <i>Statistical Machine Translation</i> to favour fluent translations.
N-gram	A contiguous sequence of <i>N</i> items (for example, words) as seen in a sample of sequences (for example, a corpus of sentences). Statistics about word <i>N</i> -grams gathered from a text corpus may be used, for instance, to build probabilistic → <i>Language models</i> .
Neural Machine Translation	A kind of → <i>Corpus-Based Machine Translation</i> using → <i>Artificial Neural Networks</i> trained on large parallel corpora. Most Neural Machine Translation systems consist of an → <i>Encoder</i> and a → <i>Decoder</i> , each of which is a specialized artificial neural networks.
Open-Source	A piece of software is <i>open-source</i> (also called <i>free</i> software or <i>free/open-source</i> software) if it is distributed under a license that allows users to run the software for any purpose as well as to study, modify, and distribute it and any modified versions. Typical free/open-source licenses are the Apache license, the MIT license, the BSD 3-clause license, the GPL license, and two Creative Commons licenses (CC-BY and CC-BY-SA).
Phrase-Based Machine Translation	[also <i>Phrase-Based Statistical Machine Translation, PBSMT</i>] → <i>Statistical machine translation</i> using <i>phrase</i> pairs, that is, correspondences between sequences of one or more contiguous words, such as (<i>statistical machine translation, traduction automatique statistique</i>) as building blocks.
Recurrent encoder	A specific kind of → <i>Encoder</i> using a → <i>Recurrent neural network</i> .
Recurrent neural network	An → <i>Artificial neural network</i> that at each cycle or time step <i>t</i> receives, among its inputs, part of the outputs it produced in the previous cycle or time step <i>t-1</i> . Recurrent neural networks are capable of processing a sequence of inputs of any length and produce an output at each step that depends on what has been processed until now.
Representation	In → <i>Artificial Neural Networks</i> , the vector made up of the activations of a selected group (usually called a <i>layer</i>) of → <i>Artificial neurons</i> , written as a → <i>Vector</i> . Representations of

similar items of knowledge (words, sentences, concepts) are usually mathematically related.

Rule-Based Machine Translation	Machine translation that works by applying grammar rules and dictionaries written by experts, as opposed to → <i>Corpus-Based Machine Translation</i> .
Statistical Machine Translation	A type of → <i>Corpus-Based Machine Translation</i> that generates translations based on probabilistic models obtained by performing statistics on parallel corpora. Statistical machine translation systems assign an approximate probability or likelihood to many possible translations of the source segment and then pick up the most likely. See also → <i>Phrase-Based Machine Translation</i> .f
Tokenization	Segmentation of source text in <i>tokens</i> , usually (but not always) words or similar units, before sending to a machine translation system: <i>We won't go to Ben's party</i> → We + wo + n't + go + to + Ben + 's + party.
Vector	An ordered, fixed-size list of real numbers such as "(-0.73, 0.31, 0.76, -0.02)".
Word Embedding	A representation of a word as a → <i>Vector</i> . Used in → <i>Neural Machine Translation</i> . Word embeddings are usually learned from large monolingual texts. For similar words, embeddings are also mathematically similar.