

1. TRUE/FALSE - Machine translation (MT) evolved in the manner which is typical of AI systems over the past 50 years, from symbolic rules to statistical models to neural models.

True

3. TRUE/FALSE - For longer sentences, the basic recurrent neural network architecture tends to “forget” earlier parts of the sentence, resulting in rather poor performance. In the late 1990s a research group in Switzerland proposed a solution to this problem: the individual units in a recurrent neural network should have a more complicated structure, with specialized weights that determine what information gets sent on at the next time step and what information can be “forgotten.” These researchers called the more complex units “long short-term memory” (LSTM) units.

True.

5. What does MM think of the claim that machine translation is now close to “human level”? (Please provide at least some of the evidence that she provides in justifying her thinking on this issue.)

She thinks this claim is overstating the abilities of machine translation. It often picks incorrect meanings for words with multiple meanings, such as “bill” and “crisp” in the restaurant example, and it struggles with idioms.

7. TRUE/FALSE - The architecture of an image captioning system differs from that of a machine translation system in that the encoding RNN of the latter is replaced by a CNN in the former.

True

9. TRUE/FALSE - While Microsoft’s CaptionBot says it can “understand the content of any photograph,” the problem is that the opposite is true. Even when their captions are correct, these systems don’t understand photos in the sense that humans understand them. In fact, they tend to be incapable of describing the most interesting aspects of a photo, the way it speaks to us, to our experience, emotions, and knowledge about the world. That is, it misses the meaning of the photo.

True.