


## DeepL Questions

When DeepL, the neural machine translation engine ([www.deepl.com/translator](http://www.deepl.com/translator)), was released at the end of August, it took me a few weeks to talk to someone at DeepL/Linguee about, well, DeepL. (By the way, it's pronounced deep-L.) The company has been a little overwhelmed by the immediate attention it received when, with great fanfare, it released its neural machine translation engine between English, German, French, Spanish, Italian, Polish, and Dutch. This new engine is producing results that in many cases seem to be better—sometimes significantly better—than both Google Translate and Microsoft Bing Translator.

When I finally got to talk to Jaroslaw Kutylowski, DeepL's chief technology officer, there was a lot he couldn't legally and strategically say. What he could say, however, was still very interesting. For instance, when I asked him how long they've been planning to have a machine translation engine, he told me that it really only started to occur to them about a year ago, when neural machine translation first popped up and became everyone's favorite topic of conversation. It does make sense, though, that the company known for Linguee (its multilingual dictionary and corpus tool) is using its somewhat curated corpus in the European Union languages, plus Russian, Chinese, and Japanese, to train a neural machine translation engine. (The entire company is actually called DeepL now, with Linguee being one of its products.) How much the company estimated that curation/editing/dictionary-building plays a role in the relatively high quality was among the questions for which I received no answer.

While there was no time commitment to the immediate roadmap, there is much emphasis on adding new languages, which likely will be those already covered by Linguee but could include others as well. Also on the roadmap is the development of an application programming interface (API). The API is particularly important if DeepL is to be used by professional translators who want to use it not on a



webpage, but integrated into a translation environment tool. And at that point of the discussion I knelt down (actually I didn't, but I would have if it would have made a difference) and asked that, in exchange for payment for the use of this API, if DeepL would commit itself to *not* use the translated data for training purposes. Its current practice is to use the data.

While Jaroslaw didn't completely commit himself to this proposal, he said the likelihood of this commitment was high. This is really, really good news because this is what makes the tool usable for professional translators. Microsoft has remained stubbornly set on its policy of using the data you upload through its API for training purposes. (Some of your clients might not mind this, but many others will, no matter how much Microsoft assures us that it will only happen at a high level and in a temporary manner). Google has understood the issue and assures us that it's not using the data if used through the API, and now, seemingly, DeepL has as well.

One thing I had been very surprised about was the "voice" of the press release for the launch of DeepL. Despite the fact that the company is located in Germany, the announcement was decidedly American in the sense that it was rather uncompromising. (The actual term that came to my mind was "hyperbolic.") Jaroslaw says that you have to be self-confident when you have good reason for it, even if it might not match your surrounding culture. I guess that's a pretty healthy way of looking at things.

To come back to the quality, you can see some numbers in the company's press

release ([www.deepl.com/press.html](http://www.deepl.com/press.html)) that show impressive quality gains. My own—very subjective and limited—testing showed similar results. Particularly when it came to advanced technical and semi-technical texts, the quality of the English>German direction was decidedly better and more natural than Google's and Microsoft's neural output. When it came to relatively high-brow press material, however, the pendulum seemed to swing the other way. But again, the sample I had was obviously very small.

Here is something I found interesting and reassuring. When we first looked at the neural output produced by Microsoft and Google in comparison to their earlier statistical engines, they seemed oh-so elegant and fluid. Suddenly, though, when compared to the results of DeepL, they look terrible again. This reinforces how easily impressed we are with advancements while often forgetting to examine the results on their own merits and realizing that they all have a very long way to go, including DeepL.

For fun, try this interesting little experiment. If you retrieve data from Linguee ([www.linguee.com](http://www.linguee.com)) to be translated by DeepL, you won't get the existing translation already found in Linguee, the one on which DeepL was trained. Instead, you'll receive a completely new one, all carried out by the neural computer "brain." Chances are this would have been different if they had built a statistical machine translation program where the original fragments might in fact have been reassembled, but that's just not how neural machine translation works. ●



**Jost Zetzsche** is the co-author of *Found in Translation: How Language Shapes Our Lives and Transforms the World*, a robust source for replenishing your arsenal of information about how human translation and machine translation each play an important part in the broader world of translation. Contact: [jzetzsche@internationalwriters.com](mailto:jzetzsche@internationalwriters.com).