

Invited Talk

Olga Beregovaya

What we want, what we need, what we absolutely can't do without – an enterprise user's perspective on machine translation technology and stuff around it

As the gains made by current phrase-based MT systems approach a horizontal asymptote, those of us in industry who benefit from cutting-edge research are giving serious thought to what's next. The goal of this talk is to rekindle the dialog between research and industry circles on the immediate and future needs of the user community and finding mutually beneficial and interesting avenues for collaboration.

Of course, there are too many topics to cover in an hour, but I'll try to touch on some of the most relevant...

These are the areas of utmost interest and importance for the global business community as seen by a user from a major Language Service Provider:

Translation Quality: It's clear that we've nearly hit the wall with what more we can squeeze out of phrase-based MT systems, so where do we go from here? Are factored models the way to go, or should we in industry start giving serious consideration to systems using deep learning techniques and/or deep syntactic/semantic structures?

New real-life applications: User-generated content (UGC) is rapidly becoming a huge opportunity for translation & localization. Whether we're talking about blogs, reviews, or live chat, there are challenges for the MT status quo: normalization, ungrammatical/uncommon syntax, extreme sensitivity to maintaining proper negation or sentiment.

Domain adaptation: How much can we get out of minimal amounts of data? A little more data? Lots of out-of-domain data? How can we seamlessly integrate client/user dictionaries into standard SMT workflows?

Using and interpreting metadata: We're seeing a trend where we often have as much or more metadata than actual text to translate. How can this be leveraged to improve results, and make translators'/post-editors' lives easier?

Quality evaluation, utility prediction: What do we even mean by "utility"? Let's work together on establishing a standard. Can we get robust, reliable QE from limited (or no) bilingual data? (We've got some evidence that the answer is yes...) What can we say about quality *vis à vis* functionality?

Collaboration: We want to help! We're very interested to see research that's directly applicable, and we want to find ways to facilitate academic/industry partnerships. We can work with clients to try to make large amounts of domain-specific (non-parliamentary!) data available. Please reach out to us so we know what you're working on. Better yet, let's work on finding mutually beneficial projects.