

The future of localization



AI Business
eBook Series



*How to employ language and culture as
part of your digital transformation toolkit*

In collaboration with:

LIONBRIDGE

2

At the core of business expansion

Localization is concerned with adapting products, communications, services, and more, so that they fit with customer expectations in a particular location.

In a globalized world, localization enables companies to tap into international markets to increase their revenue, drawing on economies of scale in the process. Localization is therefore at the core of any business strategy that goes beyond a local product. Even a company that is linked to a specific region through its branding – such as Apple’s promise of products that are “Designed by Apple in California” – must make localization efforts in order to appeal to customers across borders.

Of course, translation remains a core part of localization. Some of this work – as will be discussed in this eBook – is carried out by human linguists. Other aspects are entrusted to cutting-edge machine translation (MT) tools. Still others are carried out using a hybrid approach, called augmented translation (see page 5), involving both.

But true localization goes beyond language. Successfully transplanting a service or product from one market to another requires understanding of the subtleties of each location. “We basically recreate an experience in another language,” Inés Rubio, senior manager for language services at NI (National Instruments), told AI Business. “It’s not about the language; it’s about the culture.”

In this eBook, sponsored by translation and localization specialist Lionbridge, we look at some of the challenges facing international businesses, and the benefits of localization done right. We interview the experts, and feature perspectives from organizations including Zynga, Royal Caribbean, Volvo, and the Wikimedia Foundation. We hope that this publication will equip you with the knowledge required to build truly global products and services.

Luke Dormehl | Associate Editor | AI Business

eBook in collaboration with:

LIONBRIDGE

The future of localization | www.lionbridge.com



3

Introduction to modern localization services
Language and culture

4

Market overview
Localization goes mainstream

5

Putting AI to work
What is augmented translation?

7

Case study
How Lionbridge helped Royal Caribbean overhaul its localization efforts

8

Working for a global audience
How Zynga uses localization to capture the hearts of millions of gamers

10

Thinking locally
Using language services to improve Wikipedia

12

Advances in machine translation
Inflection points

14

Case study
How Lionbridge helped Volvo Cars create a unified online marketplace

15

What’s next?
The future of language technology

17

The academic perspective
A question of parity

Language and culture

“It’s not about the language; it’s about the culture. So when you talk about, for example, Spanish, where are we targeting? Is it global Spanish? Is it European Spanish? How do we make it relatable to the target audience? It’s definitely not just about translating words. It goes much, much [deeper than that.]”

Translation is important – but it is even more important to understand the target audience

Most commonly, localization refers to linguistic translation of assets such as websites, instruction manuals, and software applications, so that they can be read or otherwise understood.

“The purpose of localization is really making the product or offering local, and a big part of that is speaking the language of the end consumer,” Kajetan Malinowski, product leader and strategist at Lionbridge, told AI Business. But language is just a part of the story – and some of the world’s most lucrative markets require cultural expertise.

Beyond translation

Have you ever heard the phrase, “it’s not what you said, but how you said it?” There is a reason why localization is not simply called translation. Localization means taking into account a broad range of aspects of local communication. Linguistics is part of it, but so are regional differences. Units of measurement are one simple

example of a standard that varies from place to place. Other elements that vary between location are regulatory or technical aspects, and even ‘tone of voice’ or register.

“We basically recreate an experience in another language,” Inés Rubio, senior manager for language services at NI (National Instruments), told AI Business. “It’s not about the language; it’s about the culture. So when you talk about, for example, Spanish, where are we targeting? Is it global Spanish? Is it European Spanish? How do we make it relatable to the target audience? It’s definitely not just about translating words. It goes much, much [deeper than that.]”

Localization could also involve translating graphics, such as interface icons, in order to meet the target audiences’ preferences. Such icons do not necessarily have universal or unequivocal meanings. For example, a light bulb to indicate an idea is commonplace in markets such as the US and UK, but does not have the same meaning elsewhere. A ‘thumbs up’ emoji has positive connotations in the US and the UK, but is offensive in Latin America, the Middle East, and West Africa. The implications of colors vary as well. Orange in India is equated with love and courage, but in Middle Eastern cultures is more commonly associated with loss and mourning.

Get localization right, and the benefits are enormous. Get it wrong, or fail to understand the importance of local customs, and the results can be extremely damaging. It is one of the most important tools any modern business should wield in the 2020s.



4

Localization goes mainstream

“When you’re just starting, localization is completely off your mind,” she said. “It’s something to think about at the last minute. It’s not considered upstream at all. As you learn from mistakes, and as you realize how much more costly it is to make those changes afterwards, the company and the whole team [comes to appreciate it].”



Looking at the size of the localization market – and the drivers behind its growth

The birth of modern localization – often abbreviated as ‘l10n’ where the number 10 refers to the number of letters between l and n – dates back to the 1980s. This was, not coincidentally, the decade in which computer software started to become a fixture in offices and homes all over the world. To capitalize on this growing market opportunity, companies had to embrace localization, whether it was supporting new character sets, or adding or altering functions for certain markets. While this was initially done in-house, the scale and importance of such work rapidly led to the development of a fully-fledged localization industry.

Bigger than ever

Over the past decade, the language services market – comprising instruction, translation, and broader localization efforts – has nearly doubled in size. A recent report from Sclator, a leading publication for the language services industry, claimed that the addressable translation, localization, and interpreting market represented \$23.8bn in 2020. Technology, gaming, media, and life sciences were the fastest-growing

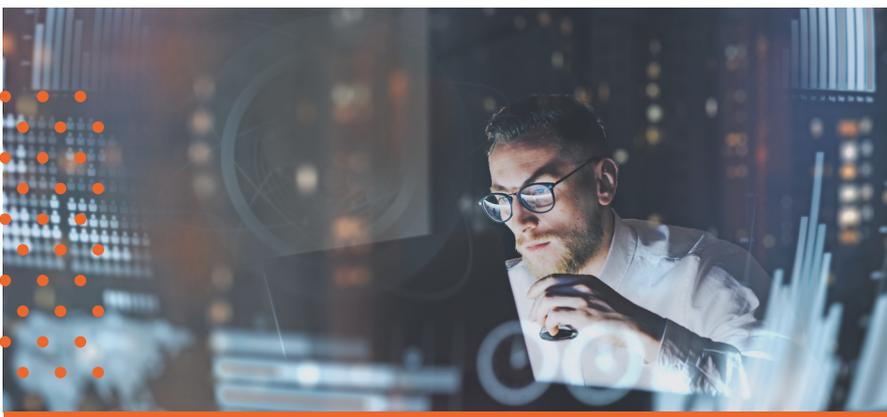
customer verticals, and the industry as a whole was expected to grow by up to 10% in 2021.

According to Statista, Europe represents around half of the global language services market, despite being home to less than 10% of the global population. That is due to the variety of languages spoken in the region. The United States, meanwhile, remains the leader by market size.

The disparity isn’t only observed on a country-by-country basis. While localization is a mature discipline in some sectors, in other (particularly newer) industries and companies it is still cementing its importance.

Inés Rubio, senior manager for language services at NI, suggested that the commitment shown to localization can be seen as a proxy for maturity. “When you’re just starting, localization is completely off your mind,” she said. “It’s something to think about at the last minute. It’s not considered upstream at all. As you learn from mistakes, and as you realize how much more costly it is to make those changes afterwards, the company and the whole team [comes to appreciate it]. Then they start bringing those things into the process.”

Thanks to technological advances involving localization, the tools that power this industry are getting better all the time. New innovations, such as neural machine translation (NMT), have only emerged in the past few years – but are already having a significant impact. These breakthroughs will continue to transform the world of localization for both customers and end-users alike.





What is augmented translation?

Discussing the relationship between humans and machines with Lionbridge's Will Rowlands-Rees and Jamie Punishill

“Machine translation is one of the foundational applications of AI,” Will Rowlands-Rees, chief product officer at Lionbridge, told AI Business. “The first machine translation occurred in the early 1950s. Even back then, they were predicting there would be no more human translators necessary. They might have been slightly off on those predictions.”

Nearly 75 years later, humans are still an important part of translation efforts. This is true even in cases where machine translation is involved. While free, publicly available translation engines can be effective in certain scenarios, they are rarely appropriate for business use.

“We’ve all got funny stories about pasting in some copy and the stuff that comes out the other side, right?” Jaime Punishill, chief marketing officer at Lionbridge, told AI Business. “Those are the untrained public engines. It is a little better in a commercial environment, but you still get weird linguistic issues between

two languages. Trying to translate English to Korean isn’t easy, much less German to Korean. The machine can do funny things.”

The importance of good translation

Reading a translated news article might be acceptable if all you want is the gist of the story. But there are other situations in which “good enough” isn’t, frankly, good enough. A prospectus for an investment product, a medical research label, or a legal text are just a few examples in which translation must be carried out perfectly.

“If I’m producing a manual for a leading sports car manufacturer, if you give the wrong instructions, there’s health and safety to consider,” Rowlands-Rees said. “Or think about clinical research or any other area that could be considered [a matter of] life and death. From a company perspective, everything is also going to have brand implications. You want to make sure it’s translated correctly, because if it isn’t, that reflects badly on the brand.” Traditionally, this process has been

called post-editing, in which machine translation tools are used to carry out initial translation. These raw MT segments are then handed over to human linguists who correct any errors made by computers. While this approach theoretically saves time, it also has its limitations, like the fact that human translators are only brought in at the end of the process.

Enter augmented translation

Augmented translation, a term originally coined by CSA Research, utilizes a range of technologies to allow human linguists to achieve high quality translation faster, along with improving consistency and accuracy. Augmented translation improves the translation experience by better pairing humans and technology. Many of these approaches have previously been explored but, by linking them together, it is possible to create a more comprehensive framework. It automates many tasks deemed low-value, but which take up large amounts of time, while allowing human linguists to utilize their expertise to the fullest. ➤

6

► Augmented translation covers seven main areas:

Translation memory

Adaptive neural machine translation

Quality estimation

Automated content enrichment (ACE)

Terminology management

Lights-out project management

Translation management system (TMS)

All these work side by side. For instance, translation memory can help with repeat translation tasks. “We might get the same product manual to translate every year,” Punishill said. “Most of that product manual is the same as it was the year before, yet in the previous, analog world, the company was paying to translate it twice – or more. Translation memories are designed to look at the source text and say if there’s a match.”

If there’s an exact match, the text can be sent through without requiring any further modification. If, on the other hand, some sections differ, these can be flagged for a human translator to update, giving them the opportunity to check the surrounding words and sentences to ensure that the addition doesn’t affect the context of the text.

Another aspect of augmented translation is being able to automatically determine the best human and machine resources to use for a specific translation job. That might mean handing off work to a human translator or, alternatively, knowing which machine translation engine will yield the best results.

Playing to the strengths

“One of the things you discover in this area is that every machine is good at different things,” Punishill continued. “One is good at English to German; another one’s good at Spanish to Korean; one’s good with

pharmaceutical content; one’s better with marketing content. You have to mix and match MT engines. You need a mix of capabilities to power your translation system.

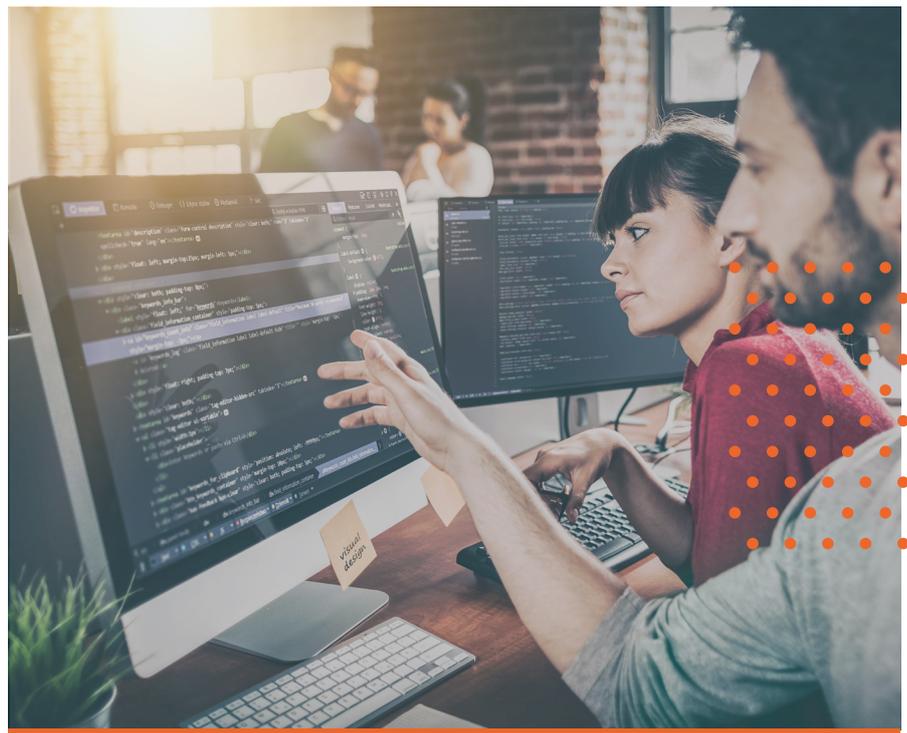
“At Lionbridge, we help companies set those up, and we run [the systems] for them. That involves all the things you have to do, from training the engines to maintaining the brand and voice guidelines, terminology dictionaries, and all the linguistic assets that keep these things running – and running really well. That’s not just translating well, but actually being able to translate in, for example, Apple’s voice, or in Microsoft’s voice.”

Rowlands-Rees noted: “We’re able to tell the user, ‘we think this document is best if it’s human-translated and here’s why,’ or ‘we want to route this to MT and here’s why we think MT is going to be good.’ We’ve understood some of the reasons why human translation is better, or MT is better, depending on the situation. We’ve programmatically understood that in submissions, and we can then route the content most appropriately,

which is efficient for us and, most importantly, for our customers.”

Augmented translation can support translation efforts in other ways, too – such as the ability to use predictive analytics to gauge which pieces of text will translate well, allowing improvement both up and downstream in the content creation efforts. In addition, it can carry out comprehensive assessments of areas like Search Engine Optimization (SEO) in other languages. This can greatly help with discoverability online.

Augmented translation represents the very cutting edge of machine translation. Using it can greatly improve the quality of output. By bringing together technologies like translation and quality estimation (providing a statistical estimation of how high quality a machine translation is likely to achieve) with the latest adaptive neural machine translation tools, augmented translation is a valuable weapon in the arsenal of any localization team. You can expect for this to become a more commonplace term in the years to come.



How Lionbridge helped Royal Caribbean overhaul its localization efforts

Leading cruise ship operator was looking to better appeal to global audiences

A cruise liner that sails around the world is, by definition, international. Each year, around 5.3 million travelers from across the globe line up to set sail to various exotic destinations with Royal Caribbean. To ensure they both get the experience they're hoping for, and that Royal Caribbean's brand integrity is maintained, high quality translation is needed for the company's website – royalcaribbean.com. Not only does this allow users to book their vacation, but also to explore Royal Caribbean's range of ships, activities, and destinations.

The problem to solve

For a prestigious brand like Royal Caribbean, an unwieldy and ad-hoc content translation workflow can have a negative impact on both the accuracy and communication abilities of its content. When the company decided to update its content platform, it sought out established experts in localization to help. Royal Caribbean elected to use the services of Lionbridge, a company with more than 25 years of experience in language

services. The goal was to drive both traffic and sales with a centralized site for content, improve translation workflows, and deliver a superior user experience.

The solution employed

The solution called into service consisted of two tools: Lionbridge's Connector™ and Adobe Experience Manager™. This provided both the efficient workflows and multilingual expertise that was required to ensure that not only was Royal Caribbean's content properly updated, but that it was also optimized for local search. The materials were available in the languages necessary, in an acceptable time frame, at a reasonable cost, and a high level of proficiency. In just four months, upward of 80% of the existing Royal Caribbean content – adding up to

approximately 1,500 pages – was translated into seven languages.

“Using the Lionbridge Connector for Adobe Experience Manager to automate the Web's translation process has allowed us to translate thousands of web pages with ease,” said Ariadna Castro, product content optimization manager for Royal Caribbean Cruise Lines. “The technology Lionbridge brings to the table helped us to implement an efficient translation workflow. The account team has been excellent. They are always ready to jump to support our goals.”

Localization projects should, first and foremost, provide a better user experience for customers. However, they must also translate into tangible results that showcase why it is such a beneficial investment for companies. In the case of Royal Caribbean, the effects were noticeable almost immediately. Within the first few years of adoption, the company reported a 35% increase in organic clicks, an even greater 37% increase in organic search impressions, and a very impressive 3x improvement in global clicks.

With a more personalized experience for customers from all over the world, everyone involved reaped the rewards of a localization overhaul well done.

1,500
pages, translated into 7
languages in 4 months by
Lionbridge



8

How Zynga uses localization to capture the hearts of millions of gamers

“We try not to use the term localization; we try to use the term ‘culturalization,’” said Bernard Kim, Zynga’s president of publishing. “So not only do we shift the way that the product reads out with regards to the language, but we try to adapt the game dynamically for the local market with regards to gameplay and engaging features

President of publishing Bernard Kim and head of strategic partnerships Vineet Rajosi Sharma share notes on developing games with worldwide appeal

Games are universal and appeal across borders. In order to thrive in the current gaming landscape, the approach of large game developers has to follow suit. Established in 2007, Zynga is one of the world’s leading developers of social games that are played by millions of people around the world, every single day. One of its best-known titles, FarmVille, launched on Facebook in June 2009 and accumulated 10 million daily active users within weeks. More than a decade on, having largely pivoted to mobile, Zynga continues to attract tens of millions of monthly players in more than 150 countries, with titles like Zynga Poker, Words With Friends 2, Game of Thrones Slot Casino, and Harry Potter Puzzles and Spells.

One of the keys to making its “forever games” (as the company calls it enduring properties) so successful is a strong emphasis on adapting its

products for markets around the world. After all, those who speak English as either their first or second language make up less than one-twelfth of the world’s population. A game experience that’s available only in English misses out on most of the potential user base. Presenting games in multiple languages makes solid commercial sense. But, as Zynga executives told AI Business, successful localization is about more than just translating text within a game.

Call it “culturalization”

“We try not to use the term localization; we try to use the term ‘culturalization,’” said Bernard Kim, Zynga’s president of publishing. “So not only do we shift the way that the product reads out with regards to the language, but we try to adapt the game dynamically for the local market with regards to gameplay and engaging features. For example, in Korea gamers might engage more times throughout a day and get more obsessed with one of our games. We try to create offers and challenges for those players that are culturalized for the local market.”

Vineet Rajosi Sharma, Zynga’s head of strategic partnerships, told AI Business that it was important to make players feel that games have been designed especially for them. “If I’m a user, and I interact with an app, it should feel like it’s made for me,” he said. ➤



► Making a game that simultaneously appeals to millions of users in multiple countries while feeling personalized isn't easy. Kim, a veteran of the games industry, said that he's had plenty of experience of running into challenges with IP that doesn't translate well around the world. One lesson he has learned is that localizing a title for a particular market can't feel "tacked-on."

"[Gamers know if] a game is one they could play for a long time," Kim said. "The feel, the language, the way that things are presented – if any of this is done at the last minute for certain markets, you will see that [reflected negatively] in the engagement numbers."

To counter this, Zynga considers its localization (or culturalization) strategy right at the start of any new project. It's all about "broadening the funnel" in terms of the number of potential players.

"If you look at our forever franchises, they have global appeal," Sharma said. When choosing which concepts to greenlight, "we make it a huge priority to culturalize," he added. "All the aspects of representing the app in the market need to be understood."

Language: Important, but not the only factor

Language is, understandably, a key consideration for Zynga, which has studios based around the world, with offices in the United States, Canada, India, Ireland, and Turkey. Sharma shared a tip that can be used to make the localization process simpler: cut down on the amount of text. "If you put too much text in the game, then you're basically making your job harder," he said. "And nobody likes to have to read that much text either." This is especially true for the smaller screens of mobile devices.

To make matters more challenging, any text has to be translated in a

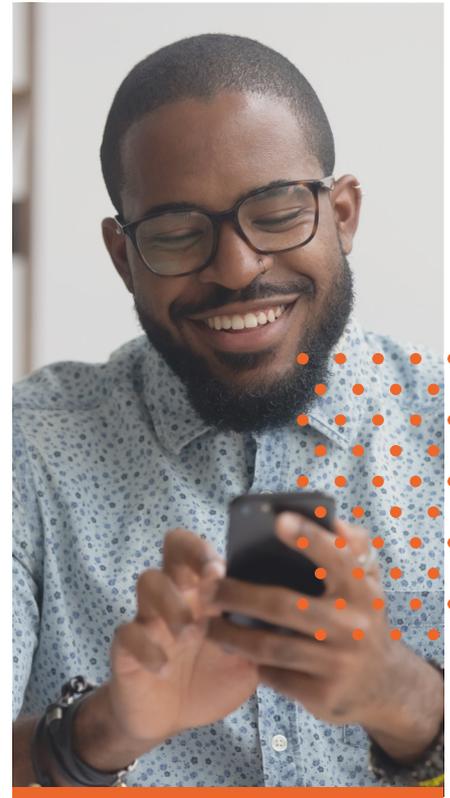
manner that conveys the original title's 'voice.' Along with in-game text and notation, this also includes elements of fictional universes like currencies, measurements, date formats, and more. Then there are app store descriptions, marketing materials, and other accompanying text which must be translated. "There's a translator, whether that's in-house or outsourced," Sharma explained. "What we really hone in is [the quality of the translation.] We have localization QA, or 'international QA' as we call it, that ensures that the translations that are coming in are coherent and make sense."

Written and spoken language is just the tip of the iceberg for the broader culturalization efforts. For instance, there are certain symbols, ways of addressing gender, and titles that have to be considered. The same is true for integrations with locally popular social media platforms, whether for marketing purposes, authentication, or social sharing. Kim said that one of the notable details that varied across markets was the "shape of an in-app purchase package" inside a game, often changed to make it more appealing for players in a particular location. In the past, the company has held back games – such as its Harry Potter title in certain Asian markets – to ensure optimal experience.

"This is exactly why we have human translators, and they specialize in our games that they're helping with and work with those games for a really long time," Kim said.

The importance of localization in a free-to-play world

Localization, Zynga's executives noted, is even more important in the free-to-play gaming sector. Unlike a premium title that charges a fixed amount to purchase, 'freemium' games make their money by bringing back recurring players. That means it is not sufficient to focus on surface-level localization; for example, running



a good local marketing campaign but making minimal changes in-game. Freemium, Kim said, ensures quality because in a crowded marketplace, gamers will only return to the titles they have a strong affinity for. That typically means games developed using a localized approach. With some 3.8 billion smartphone owners globally, according to Statista – around 48.33% of the world's population – localization is a smart investment for a mobile games studio to make.

"Overall, we see culturalization as a creative endeavor," Kim said. "It goes all the way from the first time a user hears about our game – whether it's an ad creative or a local marketing campaign – to after they become a player, with CRM (customer relationship management) outreach. Is that by emails or notifications? What time is it, because you don't want to send a push notification to a person that's sleeping, or send them in a language that's not appropriate? All these factors comprise culturalization."

Using language services to improve Wikipedia

“Where machine learning comes in ... is using models to make the experience of readers and editors better,” Albon said. “For example, we have models that look at, and try to predict, the quality of an edit, and whether or not we think an edit is a damaging edit or not.”

Chris Albon, director of machine learning at the Wikimedia Foundation, discusses the project’s mission and the importance of local data.

The world’s largest user-generated encyclopedia should be accessible as widely as possible. Can localization services help make a valuable resource even better for users everywhere? That’s the question that faced the creators of Wikipedia, which boasts more than 40 million articles in 300 different languages, catering to a monthly audience of 500 million users. Their answer: absolutely.

The most obvious way to ensure that every Wikipedia version is as comprehensive as possible would be to aggregate them all, then carry out automated translation to, for example, translate English articles (English language being the biggest Wiki around) into other languages with smaller numbers of users.

The Wikimedia Foundation has incorporated some limited aspects of this, using a Foundation-supported content translation tool that provides an initial machine translation of a Wikipedia article that a human editor then reviews, making improvements as needed, before publishing. This is crucial, since unsupervised machine translations can cause problems. In one notable error, the words “village pump” in English were translated to “bomb the village” in Portuguese. By keeping a human in the loop, such errors can be minimized. To date, the automated tool has been used to help translate over half a million articles.

However, as Chris Albon, director of machine learning at Wikipedia, told AI Business, rather than going full-tilt with

automated translation, the organization has adopted an approach “that actually takes into account local communities’ opinions a lot better.”

Localized machine learning

Wikipedia uses machine learning models on a far more widespread basis than many users might think. “Where machine learning comes in ... is using models to make the experience of readers and editors better,” Albon said. “For example, we have models that look at, and try to predict, the quality of an edit, and whether or not we think an edit is a damaging edit or not.”

Wikipedia employs hundreds of “bots” of various complexity that are developed to do everything from suggesting possible article topics to cracking down on vandalism, so that pages remain readable. But while Albon and his colleagues (who do not personally write content for Wikipedia) only want the best for the free encyclopedia, they also don’t want to risk embedding their own biases in the AI models they build.

What would bias mean, when it comes to Wikipedia? Simply put, a person in San Francisco’s idea of what Wikipedia needs, or what constitutes a good Wikipedia article, can’t automatically be applied to a Wikipedia user in Kenya with a different set of cultural expectations.

“The Wikipedia that I’ve obviously read by far the most is English language Wikipedia,” Albon explained. “So my ideas of what would be a good article, >



➤ or what a particular topic of article should be, or how that article should be structured, or what kind of tone should be used in that article are all based on my experience with a single community in the Wikipedia project. And that does not wholly apply to all the different communities. We all are under this broad consensus of being part of the Wikipedia project, but when it comes to what is a good quality edit, or what's a damaging edit, or how you handle [a particular topic], that's very local. That's something that should be governed by the local community in that local language."

Localization is part of the Wikipedia mission

Wikipedia currently operates around 100 machine learning models that are trained entirely using local community data. These are deployed on around 30 different Wikis. "The reason that localization is so critical to us at the Foundation is that it's really wrapped into the whole project," Albon said, referring to the goal of providing everyone in the world with a free encyclopedia in their own language.

The reason for training each model on local data is that users in different countries may have entirely different expectations about what constitutes a good Wikipedia article.

"Where this becomes important is that what the Swahili community might think of as being important in articles, or unimportant in articles, or important things to link to, or unexplored topics, is different [from] something like English Wikipedia because of the local context," Albon said. "For example, in English Wikipedia, something like a matatu (a privately owned minibus) is not really [considered important]. But in Swahili Wikipedia, there's a lot of interest. That is something that is fundamental to the daily lives of folks in Eastern, Southern and Western Africa ... We don't want to impose [our view] if we were to create a model using English



Wikipedia, and then apply it to Swahili Wikipedia."

One of the machine learning models used at Wikipedia is designed to suggest when a word in an article should link to another page. Because, feasibly, the majority of nouns on Wikipedia may have their own pages, and it is possible to create an article filled entirely with hyperlinks – cluttering a page and drawing attention away from the links that might be relevant to a potential reader. As a result, Wikipedia has taken the stance that it should not link to what it refers to as "common knowledge," meaning information that is considered so rudimentary that visitors will almost certainly know it. But there's a problem: what is common knowledge may vary based on the region. It's therefore crucial that an AI model deployed for this task takes a localized approach to the assumption of common knowledge.

The gold standard

"Our gold standard for model production is that when we train a machine learning model, we use data from the community to which that model is being served," Albon said. "For example, if Swahili Wikipedia comes and says, 'I want a model to

detect edit quality,' our gold standard of what we want is to then gather training data of what a good article would look like, from that community. We'll then train the model and serve it back to them."

Improving the quality of these different Wikipedias will make them more useful to users, which, in turn, attracts more users. These new users can then provide their feedback, which makes the locally trained machine learning models smarter and even more localized. It's a unique 'virtuous circle' approach to localization that is making a valuable online resource even better.

"That's why it's so important to get the community involved in stuff like this," Albon said. "Because if we don't do that, it is going to be one set of folks from a certain region with a certain notion of what Wikipedia should look like imposing that will [on everyone else]. People talk about laundering bias through machine learning models. That is exactly what we don't want. We want the community to have a say in their models, to be part of the process of creating those models, and part of the process of governing those models. At the end of the day, for us, it is about using those [AI] models to make that community experience better."

12



Inflection points

Discussing the relationship between machine learning and language with Rafa Moral and Kajetan Malionowski of Lionbridge

Translation has long been a part of artificial intelligence. In 1933, more than two decades before AI was established as its own discipline, the Soviet scientist Peter Troyanskii presented his “machine for the selection and printing of words when translating from one language to another” to the Academy of Sciences of the USSR. At the 1964 New York World’s Fair, IBM showed off a machine translation tool able to translate between English and Russian in a rudimentary way.

While it’s easy to view machine translation as just another of many sub-disciplines which fall under the umbrella of AI, Rafa Moral, VP of innovation at Lionbridge, argues that it’s one of this technology’s core themes. “AI is all about imitating human intelligence, and there’s nothing more human than language,” Moral told AI Business. “So if a researcher in AI is trying to make something that behaves like a human, that conceptually implies dealing with language. I think this is why from the very beginning, language was a topic – possibly the main topic – for AI.”

Supercharging machine translation

Like other disciplines of AI, machine translation has undergone many changes throughout its history. These paved the way for the extraordinary advances seen in recent years, as the technology has finally graduated from promising R&D to fully fledged tools.

One breakthrough development is the concept of word embedding. Part of the machine learning and deep learning toolset, word embedding refers to an approach of word representation that means that words with similar meanings can be represented in similar ways. The process of word embedding naturally captures the meaning of words in a way that can greatly improve the translation process, setting it apart from previous approaches in which different words had entirely different representations from one another, even if they were similar in their meaning.

“Before embedding, words were a text string,” Moral said. “With the concept of embedding came the concept of vectorization, and how we could transform words into numbers, and those numbers could include semantic

meaning. Embedding is about identifying which number or vector you’re going to associate with a word, based on context.”

Neural MT: A clear inflection point

One of the major breakthroughs in recent years has been the advent of neural machine translation, which Moral calls a “clear inflection point” for the industry. Neural machine translation (NMT) differs from previously used rule-based MT (RbMT) and statistical MT (SMT). It’s a completely automated approach to machine translation that relies on neural networks, the brain-inspired tools that power much of modern machine learning. NMT offers better, more accurate translations by considering the context of the words used, instead of analyzing every word in isolation. It looks not only at words, but the relationship between them.

Technology like NMT has helped supercharge localization services offered by companies like Lionbridge, which use it as part of their arsenal for providing high quality translation of documents and other literature. >

► “There is more content than we could ever translate, hundreds times more than all translators of the world are capable of handling,” Kajetan Malinowski, product leader and strategist at Lionbridge, told AI Business. “In that sense, machine translation is helping us revolutionize the industry, helping us to manage more content, and helping humans focus on where they can add more value.”

Malinowski, like Moral, says that the advent of approaches like NMT have been critical for automated translation. But part of high-quality translation is about knowing which approach to use – some are more suited for particular types of content than others.

Knowing which approach to use for optimal results

“There are different types of content for different purposes,” Malinowski said. “There’s content that is very short-lived and low visibility, which doesn’t require super high-quality translation in order to be usable. Then there’s other content, like marketing copy, or commercials, or very specific kinds of legal or healthcare or regulatory documents that require a very high quality of translation. What machine translation allows us to do is to handle more content, and to direct humans to those more ‘value added’ tasks.”

At Lionbridge, for example, the company uses AI-based tools to help determine which approach to use for optimal results – before even a single word is translated.

“We have a hybrid process that takes the best of each approach, and tries to identify all these problems: lexical density and diversity, complexity and variability, quality,” Malinowski said. “This is our starting point. Once we have clearly defined the content of our customer, we can identify which is the



Once we have clearly defined the content of our customer, we can identify which is the best process, which is the workflow, for that type of content. That includes the best translation memory to use, the best machine translation engine to use for language or domain, and the right set of tools for that particular content type.

best process, which is the workflow, for that type of content. That includes the best translation memory to use, the best machine translation engine to use for language or domain, and the right set of tools for that particular content type.”

Humans in the loop

In many cases, as described in the article on augmented translation on [page 5](#), the best option involves a combination of machine translation and human expertise.

“[Machine translation engines] have gotten so good that they are able to cover all of these kinds of low visibility content types, where the quality just needs to be good enough,” Malinowski said. “Then, if it needs to be slightly better, humans can step in and help improve that content. That allows

humans to focus on the high value content that requires more attention. It needs to read better; it needs to have a certain flow, it needs to inspire certain emotions in the target audience. Machine translation cannot do that – yet.”

While the subjectivity of language, and the ambiguity around what constitutes good translation, continues to be a challenge for even the best machine translation tools, modern translation benefits enormously from advances in this field. “Humanity’s generating exponentially more content than we generated just one or two years ago, and it’s growing all the time,” Malinowski said. “Machine translation, and all the AI and machine learning tools we have, are helping us to create more content, to speak to more audiences, to better handle that content.”

14

How Lionbridge helped Volvo Cars create a unified online marketplace

“We wanted to move away from this scattered look and feel, and we wanted to build one brand and one platform.”

The company used to be responsible for approximately 100 different websites

When Volvo Cars decided to shift to 100% digital online sales and electric vehicles by 2030, they teamed up with Lionbridge and business services and cloud solutions provider Avanade to make the necessary changes. These included overhauling their presence online, and re-imagining their website as a digital marketplace, rather than simply a marketing tool. This necessitated the creation of a user experience that would inform and please potential customers all over the globe.

The problem to solve

Prior to the overhaul, Volvo Cars had a strong web presence – but this was extremely fragmented. The company was responsible for approximately 100 websites in a total of 45 languages. These web pages were created by one central team, but they were then passed to in-market teams to translate. This impacted overall brand coherence, since different markets wound up with different website setups and sales strategies.

By centralizing its website, Volvo Cars hoped to achieve more top-down control over the localization efforts, with central oversight of translation, and direct contact with those carrying out this all-important work. “As we [moved] to online sales, we [needed] to own the message again,” said Cecilia Ernby, production lead for global online digital at Volvo Cars. “We wanted to move

away from this scattered look and feel, and we wanted to build one brand and one platform.”

The solution employed

To optimize its translation process, Volvo Cars utilized the Smartling Translation Management System, reducing the time it dedicated to managing translation tasks by approximately 1,000 hours. Part of this involved centralizing the use of translation memories that were previously unique to each market. The company also reworked the translation workflow, with an added step that ensured content was properly optimized for local Search Engine Optimization (SEO) requirements.

Centralization and localization: A force to be reckoned with

As a result of the changes, the Volvo Cars website now appears uniform across different markets. That means that customers who are looking at product pages in the UK, in Germany, or Belgium will all experience the same user journey – albeit with high quality translation of content. This makes it easier for Volvo Cars to launch campaigns and products across multiple territories.

“So many campaigns launched at the same time, which would not have been possible with our previous setup,” Ernby said.

This case study presents a valuable demonstration of how localization and centralization can work together to provide a cohesive user experience, without sacrificing any of the benefits that local customers demand.

1,000
hours of work eliminated
through new systems

The future of language technology

Looking at the potential of wearable devices, augmented reality, and incredibly large language models

There's a classic, possibly apocryphal, story about the limits of machine translation. As it goes, an early MT tool was asked to translate the English sentence "The spirit is willing, but the flesh is weak" into Russian. It came back as: "The whiskey is strong, but the meat is rotten."

Machine translation and language technology have come a long way since the early days of the industry. But there's still plenty more distance for language technology to travel, as it becomes the essential part of our lives that its advocates have always believed it will be.

Some of the most impressive demonstrations of next-gen language technology are rooted in hardware. These innovations, often involving the fast-growing category of wearable tech, promise that translation will become seamless and ubiquitous, without end-users even having to think about it. This is edging close to the 'universal translator,' long hypothesized in science fiction, offering instant understanding of any language.

The road to the universal translator

One example of this technology in action is language translator earbuds, which combine speech processing with machine translation to provide real-time translation of spoken sentences. Notable examples include the Pilot Smart Ear Buds (which also provide a written transcript of conversations), Google Pixel Buds, and WT2 Plus AI Real-time Translator Ear Buds.

On the visual end of the spectrum is Google Translate's augmented reality (AR) feature. Called Google Lens, it

allows users to point their smartphone camera at a sign, menu, or other piece of text in another language, and see a translation overlay on their phone display – often in the same typeface as the original text. All of this is carried out using on-board machine learning tools, meaning that the Lens doesn't require an Internet connection to work. In January 2021, the company revealed that the standalone Google Lens app has been downloaded more than 500 million times on the Google Play app store. Google Lens supports the majority of Google Translate languages, which currently number more than 100.

As advanced as Google Lens might seem, it represents the first generation of AR-based translation tools. Lens' initial release was in late 2017, and more widespread support didn't arrive until later. But the

situation could change dramatically with the advent of popular head-up displays, such as smart glasses. These wearable devices will free AR from being confined to phones and tablets, and make augmenting wearers' surroundings far more seamless and ubiquitous. From a localization perspective, it means that translations will appear much more easily.

Under-the-hood advances

Public-facing demos are certainly impressive reminders of the advances in language technologies. But for those who require high quality translation, the real innovation remains under-the-hood, baked deep into software, rather than smart earbuds or smart glasses.

As explained in our "What is augmented translation?" feature on [page 5](#), the near future of translation >



16

➤ technology will almost certainly involve a hybrid, augmented approach that involves both AI and human translators. Domenico Lombardini, CEO of Absolute Scientific Translation (ASTV), a language service provider that specializes in the intellectual property, legal, and life science sectors, sees the popular conceptualization of humans vs. machines as a mistake.

“While new technologies can to some extent suppress or reduce human intervention, it is equally true that they create new professional figures and new working styles,” Lombardini told AI Business. “More and more translators are currently becoming post-editors of machine translation output, but their contribution is nevertheless of fundamental importance, because the machine is still prone to errors. Thus, the training of linguists is a sine qua non condition of the final quality of translations.”

Even when humans remain the final arbiters of translation quality, AI tools can still be extremely useful in augmented translation tasks – such as creating recommender systems that suggest the best translator for a particular project, based on their specific skill set. Rather than handing

over the entire translation task to an algorithm, professionals can use AI to analyze the contents of the document, singling out keywords and stylistic requirements. They can also tag and categorize the text, before cross-referencing this information.

Like the next generation of grammar or spell-checkers, machines will be able to increasingly assist human linguists with translation jobs. That could mean something as simple as correcting a typo – but it could also pull linguists up on stylistic inconsistencies or anomalies, SEO points, or potentially alert them of cultural sensitivity issues that they (and the company that has hired them) may wish to be aware of.

Large-scale language models

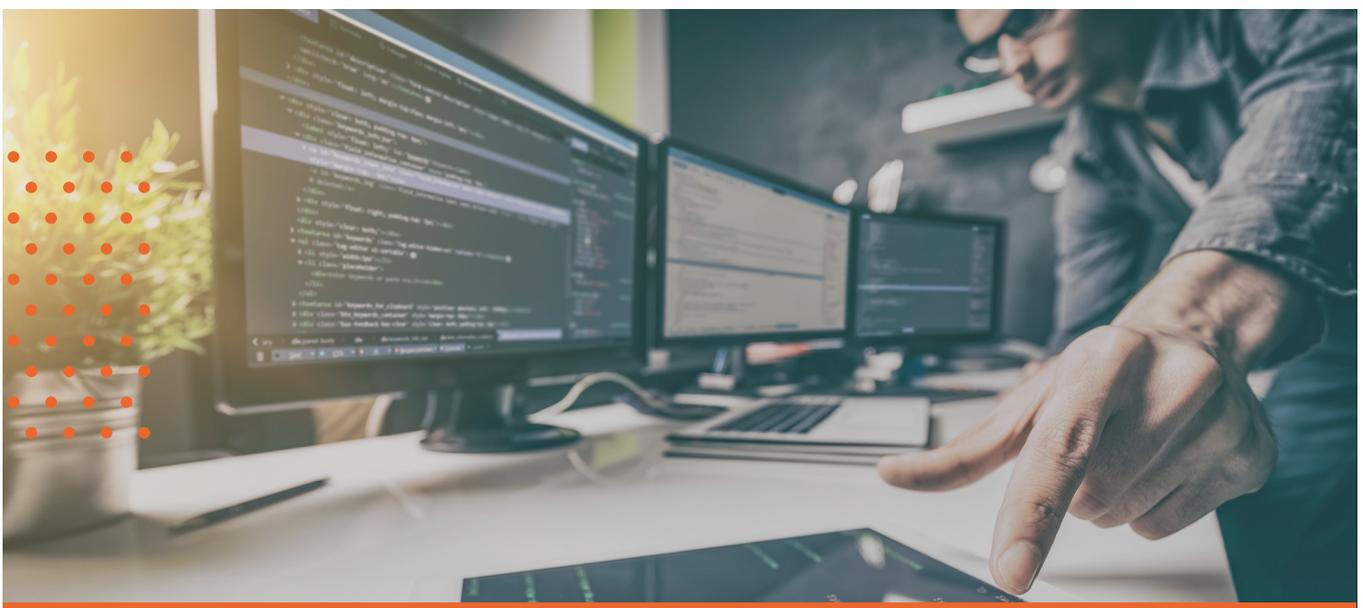
Then there is the use of AI language models for other tasks that could aid with localization requirements. One of the most significant AI advances of 2020 was the creation of GPT-3, a massive language model developed by OpenAI. GPT-3's neural network boasts 175 billion parameters, while the model itself was trained on 45 terabytes of data, representing a significant amount of the content available on the Internet. By ingesting so much data, GPT-3 can carry out

an impressively accurate job when it comes to predicting which words are statistically likely to follow on from others – for instance, knowing that “orange” is often followed by “juice.”

By leveraging its enormous corpus of training data, GPT-3 can be used for a wide range of applications – including translation tasks that can shift the language between different ‘voices,’ like turning legal text into plain English, or vice versa.

Tools such as GPT-3 open up new possibilities for language technologies: not just translating existing text, but generating entirely new copy. “AI can help us do much more with language than just translation,” said Kajetan Malinowski, product leader and strategist at Lionbridge. “It can help us write content, too. We’ve been testing the recent GPT-3 model, which allows for multiple use-cases – of which content generation is one of the most interesting.”

In a world where users expect more sophisticated and personalized messaging, innovations like large language models will be able to localize text not just for countries or regions – but potentially, for individual users.



The future of language technology

How close are machines to replicating human language abilities, really?

In a notable 2020 research paper, “A Set of Recommendations for Assessing Human-Machine Parity in Language Translation, ([web link](#))” the authors argued that “the quality of machine translation has increased remarkably over the past years, to the degree that it was found to be indistinguishable from professional human translation in a number of empirical investigations.”

Two of the co-authors, Samuel Läubli of the University of Zurich, and Sheila Castilho of Dublin City University, spoke with AI Business about the power and potential of machine translation.

AIB: You point out in your paper that machine translation has advanced significantly in the past several years. What are these advances primarily due to?

Samuel Läubli (SL): I see three main factors: A democratization of resources, more compute power, and increased commercial interest. Initiatives like Paracrawl have made ample amounts of training data available to basically everyone who’s interested in creating machine translation systems. Since neural approaches to machine translation are conceptually simpler than earlier approaches, such as statistical ones, and multiple implementations are freely available, it is easier than ever before to build a machine translation system for academic, or even commercial use.

Training neural models with large amounts of training data requires a lot of compute power, but since the same amount of money buys you a lot more of it compared to just a few years ago, this isn’t a showstopper for organizations or individuals other than big players like Google, Microsoft,

or Amazon anymore. Then there’s commercial interest. While it’s been shown that it was already possible to save lots of time – and, as such, money – in translation processes with the predecessor technology, neural modeling allowed for the generation of truly fluent sentences for the first time, which created a ‘wow effect’ among the general public – including decision makers.

AIB: What is the best measure of assessing translation accuracy?

Sheila Castilho (SC): That is an almost impossible question to answer. There are thousands of works on translation evaluation and many authors disagree on a method and even what translation quality is. To me, the best way to measure translation quality is define what quality is for that translation piece first. Then you design your methodology from there.

SL: I agree with Sheila. If the intent of a translation is for users to solve a problem with their laptop, having professional translators find and categorize linguistic errors in this translation may not be the way to go. I’d like to add one point though: Whatever you do, involve humans at some point. If increases in translation quality have shown one thing, it’s that automatic means of assessing translation quality, such as **BLEU** ([web link](#)), become less and less reliable.

AIB: What are the most exciting and promising translation technologies appearing on the horizon?

SC: I think everything that is mixing different approaches is interesting. For example, interactive post-editing and speech.

SL: In addition to mixed-initiative translation, I see two imminent changes on the horizon. First, machine translation systems will be able to consider more context, i.e., translate full documents rather than chopping documents into sentences, translating each sentence in isolation, and then [gluing] them back together. This problem is essentially solved ([web link](#)) in research settings, and will, once commercialized, lead to better coherence.

Second, there will be more control over what a machine translation system produces. Machines already produce translations with stunning quality – but decisions about style, for example, can’t be influenced by their users. In the future, you won’t only be able to control whether you want a polite rather than an informal translation ([web link](#)) of your text, but many other features such as verbosity or gendering.

This interview has been edited for clarity and length.





AI Business
eBook Series



LIONBRIDGE

Lionbridge partners with brands to break barriers and build bridges all over the world. For 25 years, we have helped companies connect with their global customers and employees by delivering translation and localization solutions in 350+ languages. Through our world-class platform, we orchestrate a network of passionate experts across the globe who partner with brands to create culturally rich experiences. Relentless in our love of linguistics, we use the best of human and machine intelligence to forge understanding that resonates with our customers' customers. Based in Waltham, Massachusetts, Lionbridge maintains solution centers in 23 countries.

Learn more at www.lionbridge.com



For more info

