Chapter 1.4

Joss Moorkens and Ana Guerberof Arenas

Artificial intelligence, automation and the language industry

Abstract
Widespread disruption to the language industry from artificial intelligence (AI) such as machine translation (MT) has been predicted for many years, but now that these technologies are being deployed, the effects are varied and, at times, unexpected. Neural MT, in particular, can produce output of greater quality compared to previous MT paradigms, but not without errors, and the best way to interact with MT to produce quality translation is not entirely clear. The use of MT and other forms of AI in the language industry necessitates consideration of risk, of value and of environmental and social sustainability. In this chapter, we introduce definitions of AI and automation, follow developments in AI within the language industry, and then consider the direction in which these developments need to go and how we might get there.

Keywords
artificial intelligence; machine translation; translation automation; lights-out project management; machine learning.

Author affiliations:

Joss Moorkens, Dublin City University/ADAPT Centre, Dublin, Ireland. joss.moorkens@dcu.ie

Ana Guerberof Arenas, University of Groningen, Groningen, The Netherlands.
a.guerberof.arenas@rug.nl

1. Introduction

There is a widely accepted view that technology will change the world of work, with less agreement about *how* it will change and whether that change will be positive or negative. The techno-positive notion of the Fourth Industrial Revolution (or *Industrie 4.0*; Kagermann and Wahlster 2022) foresees a future of efficient, decentralised global production enabled by artificial intelligence (AI), digital organization and production and a flexible workforce. This contrasts with predictions by Frey and Osborne (2017) that 47% of jobs in the United States are at risk due to “computerization”, and with ruminations in the media (and from some researchers) that translation is an occupation at risk because natural language understanding by computers is just around the corner (Metz 2015). More recently, researchers have looked at the impact of large language models (LLMs) on the US labour market and suggested that interpreters and translators will be highly exposed to this technology, although not necessarily claiming that these professions will be fully automated (Eloundou et al. 2023). The reality of automation and the use of AI in translation is a little more complicated and nuanced than either of these opposing views (Vieira 2020), with some tasks automated and some augmented.

The common understanding of automation is that it involves a machine taking full control of a process. This was also the view of translation automation that was feared by Pierce and colleagues in the ALPAC report (1966: 28) on machine translation (MT), who wrote that "perhaps, the machines will make it, but I as a translator do not yet believe that I must throw my monkey wrench into the machinery in order to prevent my technological unemployment". So, do we require a monkey wrench now that the language industry increasingly relies on automation? That depends on whether automation is used to add value or just to displace jobs. In practice, this automation does not necessarily rule out a human's involvement -- automation, it turns out, is not 'all or nothing' -- but it will change the nature of their task. This is why subsequent research on the effects of generative AI on occupations tends to discuss exposure to AI rather than replacement by AI (Eloundou et al. 2023; Felten, Raj & Seamans 2023). At the most basic level of automation, this change involves the use of word processing and spell-checking, expanding to the use of translation memory (TM) and terminology, to post-editing of MT. Under pressure from marketplace competition and amid massive digital data growth, companies and institutions are increasingly trying to maximize translation efficiency and velocity (and reduce cost) by using technology predictively, using machine learning to find patterns in data gathered from legacy tasks. As we shall see in Section 3, this occurs in a number of different ways, at times risking organizational or environmental sustainability.

In the language industry, AI (at the time of writing) usually means either MT, the current paradigm being neural MT (NMT), where artificial neural networks use vast swathes of human (and often machine) translations to predict the most likely target text translation for a source text, or generative AI tools for producing text or images, again using artificial neural networks trained on huge amounts of previous text or images. The Stanford University-led *One Hundred Year Study on AI* cites "language translation" as one of the most important advances, thanks to the creation of new neural language models (Littman et al. 2021: 13). However, as the language industry looks to both make current offerings more efficient and to expand services, there are other examples of AI being applied throughout translation workflows, with AI-related services offered to customers ranging from recruitment to pricing, data services to chatbots. For example, companies offer data collection, data annotation and labelling, and transcription that is either carried out automatically or automatically distributed to an online crowd of bilingual workers (who are not necessarily qualified translators) on demand to "perform string-based translation, which means multiple resources can process strings of data from one single file without the context needing to be understood".[1] Moreover, the use of computers (not excluding mobile phones) connected through a network means that translation and interpreting services can be offered remotely, but it also means that employees are increasingly working in isolation from home without necessarily keeping to a fixed schedule (it is common for companies to "follow the sun" by offering a 24/7 service with workers operating in several time zones) and perhaps not even a fixed building or country. Thus, most of this internal and external communication is not merely inter-personal, but mediated via on-line work platforms,[2] which are more popular than ever following the COVID-19 pandemic. All of these factors are changing the social status and organization of professionals in the language industry and creating new forms of professional communication within automatic systems.

Translators and interpreters are generally considered to be change resistant (see Cadwell, O'Brien and Teixeira 2018), but research has shown them to be pragmatic and "quite willing to adopt new technology as long as it makes their work more efficient" (Koskinen and Ruokonen 2017: 21). In fact, this practical view is probably widely held in the language industry, so that, while the most dogmatic (and at times widely disseminated – e.g., Van der Meer 2021) voices might predict the imminent demise of human translation amid an age of intelligent machines and others condemn the

---

[1] https://www.pacteraedge.com/how-ai-localization-differs-traditional-localization

[2] Such as those of Lionbridge Geoworks, Gengo, Blend, Motaword, Smartcat and Stepes.

use of any form of MT, many translation employers judiciously apply MT to low-risk, fast turnaround texts based on tiered quality requirements. Of course, technology not only affects translators. It is also relevant to wonder how the application of AI to the entire translation workflow (e.g., project management, translation and interpreting, vendor selection) will shape the organization of the language industry and the impact this might have on the professionals currently involved in it.

In this chapter, we first define AI and automation, track previous and ongoing developments in AI within the language industry, then consider where we believe we need to go and the challenges in getting there. There is a major challenge in making all stakeholders in translation and language-production processes aware of the benefits and problems of automation at each decision point, along with related opportunities, while also being cognizant of the sustainable and ethical use of resources, both human and material. The use of AI can bring great and increasingly essential benefits. There is more content to translate than there are humans to translate it, and the need to communicate in many languages almost instantly has never been higher. However, technology also has many limitations, not always clearly outlined in the media or even scientific articles. Thus, automation needs to be implemented with care so that more stakeholders can share in the benefits while being aware of the pitfalls and the implications that AI is likely to have on their professional and personal lives.

2. Definition(s) and description(s)

A shared understanding of several concepts may be helpful when reading the rest of this chapter. In alphabetical order, these include AI, automation, conversational agents, data management, remote interpreting and speech technology, which are defined and described below.

Artificial intelligence and machine learning

Definitions of AI are often gathered into categories related to thinking or behaving like humans or thinking or behaving rationally. Russell (1999: 12) tends towards the latter as an ideal for AI research, aiming towards "the capacity to generate maximally successful behaviour given the available information and computational resources". This information may be in the form of rules or training data. Before the 2010s, the rules or data were generally human-readable and processed using rules or probabilities in what is sometimes called Symbolic AI. We can see Russell's proposal in practice in contemporary data-driven AI (including MT), in which the software outputs the largest predicted probability based on its training data – bilingual human translations for MT or huge amounts of mostly monolingual data for large language models (LLMs).

In recent times, AI research has centred around machine learning, which can be supervised, i.e. the system tries to learn how to carry out a task based on training data, as happens mostly with NMT, or unsupervised, where patterns are inferred from training data without an explicit instruction, as is the case with many generative AI tools. In both cases, we can say that machine learning is "an automated process that extracts patterns from data" (Kelleher, MacNamee, and D'Arcy 2015: 3). Machine learning using neural networks has become hugely popular since the mid-2010s and is constantly being applied to new use cases, not least in translation and, more recently, text and image generation. Neural networks are an example of sub-symbolic AI, in that information is "encoded" into numbers and processed iteratively via thousands of operations to look for connections between inputs. For NMT, this means a greater capacity to produce words in their appropriate context, boosting the fluency of output beyond previous MT approaches.

Automation

Parasuraman, Sheridan and Wickens (2000: 286) define automation as a "device or system that accomplishes (partially or fully) a function that was previously, or conceivably could be, carried out (partially or fully) by a human operator". This definition is broad enough to include complete automation of a previously manual process, partial automation of that process, or the automation of process steps that humans have not even done before. It also suggests gradations of automation, something the authors expand on in their levels of automation and action suggestion (2000: 287). At Level 1, the computer offers no assistance, leaving everything to the human. This scenario must now be exceptionally rare in translation – there are few translators using pen-and-paper or even purely mechanical typewriters. At Level 2, the machine offers a large set of alternative actions, with these options narrowed down as we continue to Levels 3 and 4. At Level 5, the human must approve an automated action, with the opportunity for human intervention shrinking to a time-restricted veto opportunity (Level 6), automated execution, informing the human (Level 7), informing the human if asked (Level 8), informing the human if the machine decides to (Level 9), until we reach Level 10, where the computer "decides everything" and "acts autonomously" (Parasuraman, Sheridan & Wickens 2000: 287). We can easily apply this to computer-aided translation (CAT), with spelling and grammar checking at Level 2, a limited number of TM and terminology suggestions at Levels 3 and 4, and automatically propagated MT that can be approved or edited at Level 5. The use of raw MT without review would equate to Level 10 (if we discount the contribution of human training data and consider MT to be fully automated). We will look into further automation steps in translation in the following sections.[3]

Conversational agents and chatbots

A conversational agent is a computer system designed to converse with a person. Some of them, such as chatbots or dialogue systems, follow a known pattern of questions and answers. Since communication may be multilingual, these conversational agents are capable of communication in different languages with varying degrees of success, depending on their training data and the language combination (Robertson & Díaz 2022). Recent use of LLMs for training along with new methods of incorporating human feedback in order to avoid offensive output (as described for ChatGPT by Ouyang et al. 2022) have produced effective publicly available chatbots that produce impressive if sometimes unreliable output (Bang et al. 2023).

Data management

As a consequence of new technologies, the value of data has grown, and more data are urgently needed for continued development. In order to maintain and gather data, the language industry can now offer services creating data (manually or using automation), "crawling" the internet for data, annotating and managing data, and evaluating data.

Remote Interpreting Services (RSI)

Remote interpreting "refers to the use of communication technologies to gain access to an interpreter in another room, building, town, city or country" (Braun 2015: 352). This service can be carried out via a telephone, which is now called telephone interpreting, by video conferencing, or through video conferencing platforms created specifically for this purpose, such as virtual booths.

Speech technology

Speech technology allows the transferral of voice to text (and vice versa) so that it can be manipulated and translated using other technologies. Automatic Speech Recognition (ASR), like MT, is not a new technology, but it has also greatly benefited from neural networks (Nassif et al. 2019). AI

[3] Note that there are many other similar typologies of automation. For a literature review and comparison between some of them, see Vagia, Transeth and Fjerdingen (2016).

can enable multilingual speech and apply it to text-based language technologies. For example, Apptek[4] offers technology that allows ASR in one language (speech to text), translates the text into another language using NMT and then applies neural speech synthesis to reproduce that speech in the other language (text to speech) with a synthetic voice.

3. Developments: Where do we stand?

3.1 Past

Development of what we retrospectively call Rule-Based MT (RBMT) had progressed for over ten years prior to the negative assessment of the rate of progress and output quality contained in the ALPAC report (1966), as mentioned in our introduction, in which the notion of MT post-editing was dismissed as absurd. However, elsewhere the augmentation of human intelligence with machines was proposed as an effective use of computing (e.g. Engelbart 1962). RBMT based on hand-coded and language-specific rules produced output that was of limited quality for general use, but for a specific field or domain it could be highly effective. For example, the 1976 deployment of RBMT and post-editing by Environment Canada, the *Metéo* English-French MT system for translating weather forecasts, was very successful (Hutchins 1995: 14) and RBMT was still used for limited domains or low-resource languages well into the 2000s. English-French translations from the Canadian Parliament were the source of early experiments with Statistical MT (SMT), a process that attempted to learn probable translation models from previous translations and then to apply these to new and unseen source texts. From the 1990s until 2015 or so, SMT became the dominant MT paradigm until output quality appeared to plateau. However, the idea that “existing translations contain more solutions to more translation problems than any other available resource” (Isabelle et al. 1993: 205) had farther-reaching implications with the introduction of the first CAT tools with TM in 1992 (Rothwell et al. 2023).

In the early 1990s, the work of translators and interpreters was mostly manual and translation workflows linear and asynchronous. Translators sat in front of a computer and translated material for print, software files or a film using dictionaries, disk-stored glossaries and style guides (probably at Level 2 or 3 of automation). Interpreters went to conferences and carried out their duty in-situ after preparing at home with a similar setup to translators. Technical aspects of projects, such as dealing with documentation, software and help, were performed by translators themselves with the help of a small number of project managers, engineers or desktop publishing specialists. Project management, vendor management and other managerial roles were in their infancy and were often shared with those translating, reviewing or interpreting. Communication between clients and providers still took place mostly face to face or by telephone, as did communication within the companies. Freelance translators and interpreters were also part of the ecosystem, but not in the proportions seen more recently. So-called “mom and pop companies”, small and specialized businesses owned by highly knowledgeable translators or people from related professions, offered translation, interpreting and language-related services. Even the timid advances of the first commercial CAT tools, incorporating TM and terminology databases, were quite rudimentary with translators almost manually creating the translation matches. We can safely say that the intelligence was almost one hundred percent human.

The global spread of the internet in the 1990s made it possible for this industry to open up to new locales, technologies, workflows, roles and complexities. Languages were no longer an afterthought but a key consideration so that all products could be shipped to many countries at the same time to increase revenue. Increases in scale, translation speed and complexity required technical

[4] See https://www.apptek.com/company/company-overview

ingenuity and in time produced translation data, parallel corpora initially for local leverage by individual translators within CAT tools. These data became the fuel for data-driven MT. CAT tools added new functionality to become relatively stable, feature-rich tools, offering automatic options for creating TMs, managing terminology and even reviewing colleagues' translations.

Although RBMT was used in some large organizations, the cost of coding rule-based engines was too high for smaller companies or freelancers. However, the move to data-driven SMT, the increasing availability of online data via the internet and open access repositories, and the gathering of TMs at scale made MT more accessible rather than something reserved for large organizations. In the mid-2000s, free online services made MT ubiquitous and by 2010, post-editing was reported to be the fastest-growing service by many Language Service Providers (LSPs) (DePalma & Hegde 2010). The advent of NMT, with improved fluency and overall quality of output in many languages, heralded large-scale application of AI in the language industry.

3.2 Present

Undoubtedly, MT in its different paradigms has taken centre stage in the language industry and it represents the main and – until the advent of generative tools – the most talked-about AI technology in use. However, in this section we also address the integration of MT into translation and localization workflows, the tools and platforms within which translation and post-editing is carried out, data gathered within these tools and platforms, and AI applications beyond MT that make use of these data. The bywords are efficiency and velocity for a number of reasons, such as the huge and growing amount of data produced daily, much of it time-sensitive, and user expectations for text and media of a fast turnaround and global availability. Software development has moved from irregular to agile to continuous delivery and deployment, with updated strings of code pushed to localization and updates published immediately after the integration of the target text.

Pressure on cost and speed together with the increased quality and fluency produced by NMT means that, for the first time, raw MT is now published (usually digitally) for low-risk, short-lifespan texts such as online reviews or tweets (Level 10 on the automation scale). Because of some disputed claims that NMT output can equal or surpass the quality of human translators (Läubli et al. 2020), one might have the impression that it is a stable and predictable technology. However, NMT output can produce mistranslations that are difficult to spot due to the increased fluency, output that exhibits gender and racial bias, and inexplicable "hallucinations" that bear no relationship to the source text. An increasing number of researchers, both in commercial and academic settings, are tackling these many remaining problems of language translation, but advances are so unpredictable that even those working in MT development find it difficult to know what exactly the future might bring.

Text generation tools such as ChatGPT and GPT-4 have been found to produce NMT-quality translation for well-resourced languages (Hendy et al. 2023), addressing context issues for well-resourced language pairs better than NMT systems (Castilho et al. 2023), and promising capability for automatic translation evaluation (Kocmi & Federmann 2023). Problems of hallucinations and bias are relevant to these systems, too, although the employment of reinforcement learning using (large and expensive amounts of) human feedback has mitigated the production of toxic output (Ouyang et al. 2022). Further, recent research has highlighted the importance of using new factors such as temperature (level of randomness of the system), task information and domain information when prompting to obtain better results (Peng et al. 2023). At the time of writing, these tools have not yet been effectively integrated into translation workflows, and it is not yet clear that this integration will differ from the ongoing integration of NMT. The possibilities that these systems might offer to translation customers are diverse, and time is needed to see the real impact of this technology in the industry. For example, customers might generate the output directly in the target language and then avail themselves of the services of LSPs to correct, fact check or embellish an existing text. This could lead to the emergence of new roles, i.e., from prompter to editor of text generated by GPT models, or the adaptation of existing roles, as with post-editing of MT.

At the moment, however, the integration of NMT into translation workflows takes place at two main levels. The first is through the routing of content based on content type, product line or other criteria in order to maximize the value of human input in translation organizations (these can be small or large). The second level is integration within individual translators' workflows.

At the organizational level, there are often variable workflows based on *fitness for purpose* (Bowker 2020). This translates into defined tiers of quality, whereby in the lowest tier, content that will have few end users is of low value and carries little risk may be machine translated with cursory or no review. Within this quality tier there may be further subdivisions, perhaps using automated quality estimation (again, the state of the art uses neural networks) to control the output, remove poor-quality segments for retranslation or post-editing and allow acceptable quality segments to pass through. Such quality estimation systems are not perfect, but they are beginning to see some industrial use. An example text here might be online travel reviews that will soon be superseded by new contributions or low-risk online software documentation that appears to be little used. The second tier of quality is for content of greater value, perhaps due to a longer expected shelf-life, a larger number of end users/readers or increased risk if mistranslated. This might be post-edited. Again, we can have subdivisions of post-editing quality, but research has shown that high quality is achievable, even if the text is likely to stick closely to the MT output, with post-editors instructed to use as much of the raw MT as possible. An example from Melby (2022) is maintenance manuals for heavy machinery that must be comprehensible (without being perfect) with no substantial errors.[5] Then, high-value and/or high-risk content will be translated by professionals, aided by technologies and usually within CAT tools. Examples here might be EU legislation, medical device translation or translations of literary fiction. Even organizations that employ highly paid translators at the premium end of the market might choose to automate translation of portions of their content. Of course, not all organizations have the appropriate content, ability, finances or interest in this tiered notion of quality, and there are examples of many who focus primarily on a low-cost strategy, routing all of their translation work to post-editing.

Although the ubiquity of NMT might give us the impression that such organizations are in the majority, and that MT and post-editing is the main activity in the language industry, many localization companies continue to offer 'human' translation as their main service. For example, the annual report from the European Language Industry states, "[a]mong language companies, standard human translation was still by far the dominant type of service", with post-editing ranking second (ELIS Research 2022). CSA Research (DePalma et al. 2019) reports year-on-year that post-editing represents roughly 4% of language industry turnover. However, the picture is not quite so clear. This brings us to the second level of MT integration: within the workflow of individual translators. The fuzziness of TM and MT integration, whereby both are served within the same CAT interfaces (or translation platforms), and the variety of ways that proposals from MT can be incorporated into a translation – post-editing of a full segment, sub-segment integration or "fuzzy match repair", MT as autosuggest or just MT as initial inspiration – means that the true level of MT use in "human" translation is impossible to measure. This is especially so if the MT is accessed through application programming interfaces

[5] Melby (2022), however, proposes "grades" of translation based on the product and use case, rather than on the production process.

(APIs), concealing the workflow process followed to produce their final translation. There are also open source NMT tools such as Opus-MT (Tiedemann & Thottingal 2020) that can be easily accessed through OPUS-CAT by a tech-savvy translator without the input of their client.[6] This is in many ways a welcome change, making the technology more accessible and "democratic", initiated by and under the control of a translator. However, this opacity when it comes to MT integration means that measuring its usage, particularly when there is sometimes stigma attached to this use, is difficult, making the percentages cited above appear naive or, at least, partial. This opacity is likely to increase further with the use of GPT models.

To a greater or lesser extent, CAT tools and particularly cloud-based translation portals or platforms are sites of interaction with AI that also allow the gathering not just of translation data, but also of translator activity data. Several large organizations and LSPs have built platforms, "digital infrastructure that can facilitate functionality and user interaction" (Rothwell et al. 2023), within which translation and editing can take place in a tightly controlled environment, with many types of data accessible to the platform manager. Of 6,925 respondents to a CSA Research survey (Pielmeier & O'Mara 2020), 89% said that they use such platforms for translation work. Some of these platforms work similarly to commercial CAT tools, but according to CSA survey respondents, many have limited functionality. However, such platforms allow the collection of many types of translation data. Bilingual parallel corpora from TMs provide the raw material for NMT training and for quality estimation (which also now uses neural networks; Shterionov et al. 2019), but other data can also be useful. This could include information on the job and domain type, translation quality measures, timeliness of the delivery, work volume, adherence to instructions, reliability, accuracy of layout and formatting, ratings of friendliness and communication skills, and feedback from the PM. Many of these attributes are used to compute Translated's T-Rank score based on over a million previous jobs (Cattelan 2017) and Wordbee quality ratings (Vela-Valido 2021) – metrics that are used to decide who gets offered which translation jobs, sometimes without any human intervention. While these work and reputation data are available to translation platform managers, they are rarely accessible to translators (Fırat 2021).

The use of AI in audiovisual translation has dramatically increased following the COVID-19 pandemic. Streaming services have been available since the mid-2000s but the figures grew exponentially during the pandemic and related lockdowns. Netflix has revealed that they "subtitled seven million and dubbed five million run-time minutes in 2021" (Marking 2022). The need for new entertainment material means increased translation velocity in many language pairs, often within vendor-specific subtitling platforms that incorporate NMT and content-specific wikis or cloud-based interfaces offering intuitive workflows (Díaz-Cintas & Remael 2021: 242–48). Content is translated, sometimes decomposed and distributed to multiple translators for a quick turnaround, into a pivot language (generally English) template file, then translated at speed with subtitle timings locked so that viewers can access new episodes or seasons simultaneously world-wide. Although there have been quite a few controversies about the use of this process for series such as Squid Games or Money Heist (Groskop 2021; Lange 2021) and official statements by the European Federation of Audiovisual Translators (AVTE),[7] and the French and the Spanish Associations of Audiovisual Translators, research has often shown that the use of NMT in subtitles improves the translation cycles without necessarily hindering the quality (Bywood, Georgakopoulou & Etchegoyhen 2017; Matusov, Wilken & Georgakopoulou 2019; Koponen et al. 2020). Of course, this research assumes that certain quality conditions have to be in place for the improvement to happen, and concern has been expressed about job satisfaction within an AI-enabled subtitling workflow (Moorkens 2020). Moreover, the advances

[6] See https://github.com/Helsinki-NLP/OPUS-CAT

[7] See https://avteurope.eu/avte-machine-translation-manifesto/

in speech technology that we discussed in the introduction mean that the possibility of automatically transcribing speech, translating it into any language as a subtitle or using synthetic voices to dub the media content is already being explored and, with varying degrees of quality, realized.[8] It also means that companies such as Netflix are looking into ways to perhaps simplify the original language to facilitate the use of NMT in subtitling (Mehta et al. 2020). Automation in the AVT industry and the subsequent lowering of reimbursement rates have caused many professionals to abandon the field altogether. A historical collective agreement has been reached that includes minimum rates for the self-employed in Finland (Laurent 2023) to find solutions to the problem of the talent crunch in media localization (Stasimioti 2022).

Powered by AI and the need for growth, the industry is no longer offering only translation and interpreting services to their clients. AI related services such as data generation, data annotation, data validation, chatbot text generation, translation and testing, engineering (text, video, speech, audio) and synthetic data creation are required internally and offered externally by LSPs. This in turn means that the type of professional that companies look for is also quite diverse, with a vast range of skills. AI is also used to recruit this "new talent": rather than manually selecting CVs or testing transcribers, candidates may be asked to submit application data through platforms that take the candidate through a series of automatic steps to test them for the role for which they have applied. In a similar way to crowdsourcing marketplaces like Amazon Mechanical Turk,[9] some localization companies offer crowdsourcing platforms that select the personnel automatically, and they also diversify their portfolio so that they do not necessarily look for qualified translators and interpreters for linguistic tasks.. As mentioned, if customers implement GTP systems as part of their multilingual content, new roles are likely to emerge in localization companies..

As discussed, in the past, the communication in the language industry between the different collaborators happened at the workplace either face to face in the office or asynchronously via email or telephone. However, this communication is increasingly computer-mediated, especially since the COVID-19 pandemic. Nowadays, not only freelancers, but also employees in the language industry spend a high percentage of their week working using a computer from home, meaning that automated and AI services can be used to control and surveil remote work and work-related communication. One company, for example, uses workplace messaging data for sentiment analysis, to see whether a worker's communication tends to be positive or negative (Prodoscore 2022). Such monitoring has been associated with "decreased job satisfaction, increased stress, decreased organisational commitment and increased turnover propensity" (Ball 2021). Since November 2022, ChatGPT can also be used to write and/or automate tasks, although this has been met with quick reactions from large corporations that have banned its use internally, an example being Amazon (Schwartz 2023) among many others.

If the past was marked by small-to-medium specialized language companies, the present is characterized by the merger of these companies into large multinational corporations with finance-focused shareholders and a translation-agnostic board of directors. Graham (2019) tracked 82 acquisitions of translation and interpreting companies between 2016 and August of 2019. The changes were so rapid that a person working for such a company that had been bought, sold or merged many times could be quite unsure of who they are working for at present, apart from their immediate supervisors or colleagues. This has completely changed the culture within the industry and has resulted in a great depersonalization of the workplace. This disconnect between owners and workers

[8] https://www.apptek.com/technology/text-to-speech-technologies
[9] https://www.mturk.com/

has resulted in reduced trust. Pym (2004: 5) writes that "to build trust is also to reduce transaction costs". Conversely, in remote work scenarios with little trust, AI facilitates increased and more invasive monitoring in the absence of solidarity in the "bulk" area of the language industry.

3.3 Trends

The previously mentioned Stanford *One Hundred Year Study on AI* report (Littman et al. 2021) suggests that the main opportunity for AI in the coming years is to augment human capabilities, and thus research and development should focus on effective human-AI collaboration, harking back to the work of Engelbart (1962, as mentioned in Section 3.1) and the original motivation behind TM (Kay 1997). Lommel proposed an augmented translation scenario in 2017, with translators (or linguists) in an empowered central role, rather than simply post-editing a static MT output or sending and receiving files. In their follow-up report, Lommel and De Palma (2021) name up to seven core technologies: enhanced translation memories, adaptive NMT, quality estimation (automatic), automated content enrichment (the inclusion of metadata to provide more detail about a given word or group of words), intelligent terminology management (more context-based terminology), lights-out project management (automatic workflow without human intervention, as occurs in current automatic job assignment, curiously leaving freelance translators isolated from their main point of contact) and translation management systems (the system that connects the different core areas and the professionals).

Lommel's (2021) prediction for what he calls Responsive MT similarly seeks to focus on improved human-AI collaboration, adding automatic adaptation for subject domain based on user feedback, automatic matching with requirements for applicability and usability, and awareness of context beyond the sentence level. The latter is an active topic of research, with Castilho (2022) looking at the context span required to solve different categories of ambiguity, from neighbouring sentences to whole text level, and Jiang et al. (2022) creating BlonDe, an evaluation metric for the document level. The use of generative tools may be useful in addressing Lommel's suggestions.

Gartner Research (Elliot et al. 2020) predict that a tiered approach based on risk and content value will continue in what they call the "localisation hyper-automation" era, with increased use of AI and automation throughout translation workflows. Human translators' attention will be on complex, strategic and critical texts, leaving simpler work to NMT. While for the most part this sounds likely, the deployment of risk prediction is likely to be tricky. Soloviev (2022: 7) proposed "risk-driven sampling" as part of a quality estimation process, but at the higher level of translation risk (where there is legal risk or risk of injury or death; Canfora & Ottmann 2018), automation would appear unwise for the time being.

In line with augmented translation, Interpreting services, although later than the translation sector, are also beginning to experiment with computer-assisted interpreting tools, which might include speech transcription or NMT to assist interpreters (Saeed et al. 2022). Interpreting skills have also been found useful for interlingual respeaking (Szarkowska et al. 2018), the process of repeating what a speaker says in a different language using ASR. This is then used to create subtitles for accessibility purposes or for understanding another language. Training of interlingual respeakers is still in its infancy, and therefore its deployment in the industry is still minimal at the time of writing. The SMART project led by Elena Davitti[10] combines speech to text, interlingual respeaking and MT to deliver real-time subtitles. This is one of several ongoing projects that seek to combine different processes such as speech-to-text, respeaking, MT and audio description. However, the high quality of

[10] https://smartproject.surrey.ac.uk/about/

best-performing ASR systems aligned with GPT models means that many of these tasks might be automated as machines produce higher accuracy than humans (Radford et al. 2022).

As discussed, NMT has resulted in improved output for many language pairs, especially in terms of fluency, but there are also ongoing challenges. Some of these include MT for poorly-supported language combinations (sometimes called low-resource or long-tail languages, meaning that there is little digital parallel or monolingual data available), domain and terminology mismatches (production of consistent terminology is still a problem) and interpretability (with previous paradigms, it was easier to identify in which step of the process an error might occur, but NMT is highly opaque). NMT developers are exploring different avenues for low-resource languages, with big tech companies focusing mostly on massively multilingual NMT (e.g. Bapna et al. 2022, NLLB team et al. 2022) and academic teams tending to focus on bilingual systems, creating extra training data and exploring different MT models (Haddow et al. 2022). These are just a few examples of the many areas in academic and commercial research that look to improve the state of the art of NMT. These efforts are testimony to the progress of and challenges towards the provision of fully automatic, usable MT despite the mixed media messages, but also to the resources dedicated to the improvement of NMT quality. As of 2022, we have had the added possibility of translating using generative tools and chatbots using new parameters and prompts. It is, thus, difficult to predict what the quality limit of MT might be in five years' time.

3.4 Sustainability

There are legitimate concerns about the environmental, social and cultural sustainability of the language industry. Translation jobs are likely to be computer-based, with associated energy and manufacturing requirements, but the increasing use of AI for MT, LLMs and beyond represent the largest environmental concern. Amodei and Hernandez (2018) show how the largest AI models from 2012 to 2018 increased in size by an average factor of ten, year on year: when combined, a 300,000-time increase in total. Strubell, Ganesh and McCallum (2019) estimate that the largest transformer neural models emit the same amount of carbon dioxide as five cars over their entire 20-year lifetimes. Schwartz et al. (2020) call this focus purely on performance without concern for cost or efficiency Red AI. Red AI is encouraged by reviewers for the largest AI and MT conferences and competitive shared tasks that pit research teams against each other to produce the optimum output based on a single metric (which rarely, but occasionally, is efficiency). The alternative, proposed by Schwartz et al. (2020: 59), is Green AI, which will still produce “novel results while taking into account the computational cost, encouraging a reduction in resources spent.” Wu et al. (2021), who include the manufacturing and operational cost of equipment in their calculations (neural networks require powerful computers and state-of-the-art graphical processing units), report that transformer-based translation models of a fixed size became 25% more efficient over a two-year period. However, this increased efficiency is counterbalanced by the continually increasing size of language models.

Taking up this call, several researchers have proposed methods to make NMT and AI tools more efficient (e.g. Jooste, Haque, and Way 2022; Shterionov and Vanmassenhove 2022). Shterionov and Vanmassenhove (2022) also note that carbon emissions can differ greatly, depending on the location, as some regions use more renewable energy than others. Dodge et al. (2022) add in the effect of the time of day when training takes place and the location of the data centre used for cloud services. While the interest in Green AI and reporting of energy use for machine learning research has increased, there is no question that Red AI is still in the ascendant, particularly for research by well-funded industrial teams. Large technology companies employ ethics researchers who often focus on sustainability, but recent events have demonstrated that when these teams become too critical of internal Red AI research, there are likely to be repercussions (Wakabayashi & Metz 2022). There is also a contrast between those of the view that AI has a part to play in moving to a future of net zero emissions of carbon dioxide (to which many local and national governments have now committed) and those who believe that a more drastic change of trajectory is needed. Cronin (2017), for example, has proposed that translation must scale down to a craft, as the current globalized language industry, with its dependency on technology and production of superfluous translation, is unsustainable. Heilinger, Kempt and Nagel (2023) have similarly questioned whether sustainable AI is possible in an economy targeted at perpetual growth.

Concerns have been expressed by Firat (2021), Moorkens (2020) and others about the sustainability of the language industry due to AI-enabled employment practices in which work is rigidly controlled and monitored, work processes (such as post-editing within limited translation platforms) are unilaterally imposed on translators and remuneration is dropping or not increasing in line with inflation (Vieira 2020; do Carmo 2020). There are anecdotal reports of translators leaving the industry and claims of a talent crunch in subtitling, one of the areas of the market most affected by these practices. On the other hand, Durban (2022) has reiterated the need for high quality translators at the premium end of the market, and Rothwell et al. (2023) and others highlight the growing range of roles relating to translation for which linguists qualify.

From its presentation in the media, we might have the perception that AI really *is* artificial. However, the truth is that AI, NMT or any data-driven machine learning technology is created and sustained with a lot of human data, ideas, annotations and tasks (including prompting). It appears that the work of translators and interpreters is not decreasing with automation, but there is a wage or earnings dispersion considering the size of companies, and, further, that the definition of translation, or what "linguists" do, is widening (Pym & Torres-Simón 2021), presumably because professionals diversify to take on jobs that pay better. However, it is a concern that the tasks done by translators, interpreters and other language professionals might be included in a lower-paid segment of the market than more technical tasks, even if these only involve annotating or evaluating data. The prediction of Felten, Raj and Seamans (2023) is that the effect of generative tools will be to broaden this devaluing of cognitive work to many other fields. If data is the new gold, and interpersonal and management skills are in demand, then this gold and its creation need to be valued and fairly remunerated, beyond only the technical or management part of leverage and reuse.

There are also concerns about the use of synthetic data, uncleaned data or MT post-edited data in the training of new models as it might be cheaper or faster than engaging with professionals, but also in the use of NMT post-editing in general. Apart from the impact on the sustainability of the profession, this can have an impact on the quality of the languages involved, with potential risks to cultural sustainability for languages that are struggling to survive. We could be facing an unnecessary degradation of lexical, syntactical and general richness of languages when translated. Studies have observed that when it comes to creativity in translation, NMT has a high tendency to *reproduction* and errors as opposed to the creative choices by professionals, but also that NMT post-editing can constrain creativity (Guerberof-Arenas & Toral 2022). There have also been some attempts to automatically identify the effects of MT on textual elements (Castilho & Resende 2022; Daems, De Clercq & Macken 2017; Toral 2019; Vanmassenhove, Shterionov & Way 2019). These studies show that it is possible to see the traces of MT in final post-edited texts: human translations appear to be richer (lexically and syntactically). More studies on the impact of these data hybrids can be valuable to decide when NMT might be less advisable or should be used with discretion. In summary, any use of AI in the language industry that will shape the profession needs to be a collective effort considering all the people involved, rather than allowing AI to "shape society independently of human choices and values, in a manner that humans are helpless to control, alter or steer" (Littman et al. 2021: 35).

4. Needs: Where do we have to go?

As we have described in the previous sections, the pervasive nature of AI and the "magical" allure given to any advance made under this umbrella suggest that humans do not have sufficient opportunity to intervene in its development and blurs the direction to which a group of very diverse individuals in this industry might need or might want to go. While the sense in business ethics is that an organization should consider people and the planet alongside profit (Henriques & Richardson 2013) and that a common good approach to running an organization can be compatible with business success (Melé Carné 2020), there has been an absence of clear corporate-social responsibility in the language industry, concentrating rather on speed, growth and profit (Moorkens & Rocchi 2020). Employees at companies that focus on business sustainability have been found to apply more effort to their work and to be far less likely to leave (Wheelen et al. 2018). It is also quite clear, in view of the situation our planet is immersed in from an ecological point of view, that careful thought has to be put into implementing practices in the language industry when it comes to AI that will sustain the ecosystem of the profession and the industry, but also of AI itself.

4.1 Resources

The three resources on which the language industry is dependant are humans, energy and data. Section 3.4 outlines the changes necessary for energy in the context of a transition to net zero emissions and for the continued availability of high-quality contemporary translation data. The ability to produce text automatically using generative tools means that an increasing proportion of language data will not be human created, making the identification and curation of high-quality data more challenging. One effect of the appearance of publicly available text and image generation tools has been a larger conversation about data harvesting practices, such as the use of web-crawled data for system training. At the time of writing, the European Data Protection Board have created a task force to investigate data gathering practices for generative tools, the Italian privacy regulator briefly banned ChatGPT in Italy due to possible privacy violations and other countries' regulators have committed to investigating methods of data collection and storage (Mc Gowran 2023).

The main resource is, of course, human, and there needs to be work that fulfils basic needs regarding salary, conditions and holidays, but beyond this, work should also be enriching, enabling growth and advancement, and offer recognition for a job well done (Moorkens 2020). Unfortunately, many of us are under the impression that the decisions that affect the language industry are made from the top down (Esselink 2022), on the basis that technological advances are inevitably dictated by economic interests. A future in which translation work continues through the levels of automation as discussed in Section 2, limiting the input from humans, is unlikely to be motivating for translators. Calhoun (2022) makes the distinction between adaptable and adaptive automation, with the former allowing the human to use their expertise to decide how automation is applied and the latter involving a level of automation that is automatically assigned. She found that adaptable automation results in better task performance with more control, but less perceived workload, suggesting that a cost strategy, automating as much as possible to reduce short-term costs, may not maximize long term benefits.

Satisfying translation work also necessitates usable interaction with tools and technologies, whether this is within versions of existing CAT interfaces, translation platforms or newly developed interfaces. While major CAT tool developers have put great efforts into their user experience, there are large differences in the user experience in the recent proliferation of cloud-based platforms. Calhoun's (2022) findings encourage customizable interaction with MT and technology rather than restriction to a single interactive mode. The future must be increasingly participatory, whereby those that provide data have a say in their use, translators are empowered to discuss changes to work practices and the use of new technologies, and users of translation are considered when translation that could put them at risk. There have been proposals such as the control of data within a translator-controlled commons with different levels of access and reuse available for respective categories of cost (Moorkens & Lewis 2019). Without a motivated workforce, the supply of timely high-quality translation data will dry up.

4.2 Research

In recent years, there has been more research and, not coincidentally, visibility of the value added of professionals when it comes to the interaction with technology. The interest in research projects and literature in Translation Studies related to the impact of MT and AI has grown hugely. Translators and Interpreters' Associations are also involved in making the work of professionals more visible. While more research and visibility are needed, these should also be in the right context and delivered in the right tone for a wider audience. It is not within the translation and interpreting sector where we have to show the value added by language professionals, but further beyond. Real interdisciplinary work is needed, as AI is interdisciplinary by nature, and more visibility of the impact of the language professionals should reach the users of translation services.

There are some initiatives that look to use technology to empower translators following a bottom-up rather than a top-down approach. As mentioned in Section 3.2, the OPUS-CAT project seeks "to empower translators to utilize NMT independently in their work, so that translators can have a larger role in driving NMT adoption"(Nieminen 2022). Translators can run an NMT engine from their own computer and connect to their preferred CAT tool. The data is confidentially stored since it is not sent to external servers or third-party solutions. Research has a part to play in creating and improving interaction with translation technology in general and MT in particular. Innovations from academic research (e.g., Green et al. 2014; Flanagan 2015) have been successfully integrated into tools in the past. Research and surveys also play a valuable role in highlighting problems with work practices and conditions, providing a voice and advocating for translators regarding their interaction with technology and ergonomic wellbeing (Ehrensberger-Dow and O'Brien 2015). Research has been valuable in highlighting gender bias in MT and will be needed to evaluate recent and coming applications of AI in translation such as job allocation (Herbert et al. 2023) and platform work (Gough et al. 2023). In summary, we believe that interdisciplinary research that has the wellbeing of all stakeholders at its core is a priority not only for Translation Studies but for any AI research endeavour.

4.3 Training

Training of translation students and practitioners will play an important role in the future of the language industry. Many translation or interpreting graduates do not become linguists, but rather move into a role within the language industry, rising to roles of power and influence. The opportunity during training to highlight ethics and sustainability issues should be taken so that graduates have a basis to evaluate the repercussions of decisions during their career. Training not just in the use of technology, but also in understanding and communicating the underlying processes of technology should empower graduates and practitioners to advise on their appropriate (and inappropriate) uses, allowing them to fill expert roles relating to translation workflows. The advent of projects such as MultiTraiNMT (Kenny 2022) and DataLitMT (Krüger & Hackenbuchner 2022) is particularly useful in this regard.

We believe that training on technology and possibly programming has to go hand in hand with ongoing training in translation and interpreting strategies, and further with an increased focus on the development of critical skills such as written and oral communication, ethical principles, authorship and copyright knowledge, workflow and management skills, and creativity, either through the creation of specific modules or with specific tasks that foster distinct human capabilities (Schjoldager et al. 2022). The increasingly popular concept of transversal skills, meaning broadly applicable skills relevant beyond a single job, task or discipline, could present a useful framework for translator education to maximize employability of graduates by focusing on less automatable tasks (UNESCO 2014).

5. Conclusion

The progress of MT, particularly since the advent of NMT, and the subsequent availability of chatbots and generative tools based on LLMs have brought opportunities, hopes and fears about AI to the fore in the language industry. As breakthrough instantiations of AI, NMT and generative tools can produce remarkably useful output of a quality well above previous MT paradigms. The term AI tends to be used for applications of machine learning, as that has become the focus of AI research, and there are an increasing number of applications of machine learning and neural networks in translation and interpreting workflows alongside the use of NMT. In this chapter, we have covered diverse activities and research initiatives from data annotation and labelling to computer-assisted interpreting. Throughout the description of how AI is implemented in the language industry and the various levels of automation, we have described real implementations of AI in the language industry which makes clear how far we are from full (Level 10) automation. AI has at its centre professionals, researchers and users in a very complex web of roles and applications, perhaps not easy to disentangle and with evident power imbalances. Because of this, the need for knowledge of how this technology is created, used and planned has only increased as the technology improves, together with the need for knowledge dissemination.

The method of AI deployment in the language industry can ultimately depend on the organization's strategy. A low-cost strategy might be good for market share (see Wheelen et al. 2018), but maximizing automation may also risk hindering sustainability and alienating employees. Such a strategy would involve moving as much translation work as possible to post-editing within a constrained interface that can collect large amounts of activity data, and then reusing this data for pricing, job assignment, automating workflow steps and even automating decisions about levels of automation. A value strategy whereby an organization looks to differentiate its offering and maximize value internally might involve building and valuing long-term relationships with internal and freelance workers, deploying MT for low-risk and low-value content, perhaps including some post-editing as bilaterally agreed with translators, within a customizable interface, and supporting high-quality content production.

In *The Myth of the Machine*, Mumford (1967: 9) argues for technology that is "broadly life-oriented", working with the unique abilities of humans: primarily language, alongside originality and creativity. This call is echoed by the more recent Stanford report that recommends an augmentation approach to use of AI (Littman et al. 2021). It is difficult to clearly measure the use of MT and AI in the language industry, and perhaps there is a tendency to focus on negative anecdotes without awareness of successful deployment of AI that benefits all stakeholders. There are huge opportunities for the appropriate use of MT and AI with cognizance of environmental and social sustainability, respectful gathering and reuse of data, minimization of risk and foregrounding engagement and comprehensibility for end readers or users. This will require an understanding of responsibilities, ethics and MT literacy on the part of those deploying AI, and real efforts to bring a participatory approach to the collection and reuse of language data.

References

ALPAC. 1966. *Languages and machines: computers in translation and linguistics.* A report by the Automatic Language Processing Advisory Committee, National Academy of Sciences. Washington D.C.: National Research Council.

Amodei, Dario & Danny Hernandez. 2018. AI and compute. *OpenAI*. https://openai.com/blog/ai-and-compute/.

Ball, Kirstie. 2021. *Electronic monitoring and surveillance in the workplace: literature review and policy recommendations.* Luxembourg: European Commission. https://data.europa.eu/doi/10.2760/5137 (accessed 27 September 2022).

Bang, Yejin, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet Do, Yan Xu & Pascale Fung. 2023. A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity. arXiv. doi:10.48550/ARXIV.2302.04023. https://arxiv.org/abs/2302.04023.

Bapna, Ankur, Isaac Caswell, Julia Kreutzer, Orhan Firat, Daan van Esch, Aditya Siddhant, Mengmeng Niu, Pallavi Baljekar, Xavier Garcia, Wolfgang Macherey, Theresa Breiner, Vera Axelrod, Jason Riesa, Yuan Cao, Mia Xu Chen, Klaus Macherey, Maxim Krikun, Pidong Wang, Alexander Gutkin, Apurva Shah, Yanping Huang, Zhifeng Chen, Yonghui Wu, Macduff Hughes. 2022. Building machine translation systems for the next thousand languages. *arXiv*. http://arxiv.org/abs/2205.03983 (accessed 15 July, 2022).

Bowker, Lynne. 2020. Fit-for-purpose translation. In Minako O'Hagan (ed.), *The Routledge Handbook of Translation and Technology*, 453–468. Abingdon: Routledge.

Braun, Sabine. 2015. Remote interpreting. In Holly Mikkelson & Renée Jourdenais (eds.) *The Routledge Handbook of Interpreting*, 352–367. Abingdon: Routledge.

Bywood, Lindsay, Panayota Georgakopoulou & Thierry Etchegoyhen. 2017. Embracing the threat: machine translation as a solution for subtitling. *Perspectives* 25(3), 492–508. doi:10.1080/0907676X.2017.1291695.

Cadwell, Patrick, Sharon O'Brien & Carlos S. C. Teixeira. 2018. Resistance and accommodation: factors for the (non-) adoption of machine translation among professional translators. *Perspectives* 26(3), 301–321. doi:10.1080/0907676X.2017.1337210.

Calhoun, Gloria. 2022. Adaptable (Not Adaptive) Automation: Forefront of Human–Automation Teaming. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 64(2), 269–277. doi:10.1177/00187208211037457.

Canfora, Carmen & Angelika Ottmann. 2018. Of ostriches, pyramids, and Swiss cheese: Risks in safety-critical translations. *Translation Spaces*. John Benjamins. doi:https://doi.org/10.1075/ts.18002.can.

Carmo, Félix do. 2020. 'Time is money' and the value of translation. *Translation Spaces* 9(1), 35–57. doi:10.1075/ts.00020.car.

Castilho, Sheila. 2022. How Much Context Span is Enough? Examining Context-Related Issues for Document-level MT. *Proceedings of the 13th Conference on Language Resources and Evaluation (LREC 2022)*, 3017–3025. Marseille, France: European Language Resources Association (ELRA).

Castilho, Sheila, Clodagh Mallon, Rahel Meister & Shengya Yue. 2023. Do online machine translation systems care for context? What about a GPT model? In Mary Nurminen, Judith Brenner, Maarit Koponen, Sirkku Latomaa, Mikhail Mikhailov, Frederike Schierl, Tharindu Ranasinghe, Eva Vanmassenhove, Sergi Alvarez Vidal, Nora Aranberri, Mara Nunziatini, Carla Parra Escartín, Mikel Forcada, Maja Popovic, Carolina Scarton & Helena Moniz (eds.) *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, 393Tampere, Finland. https://doras.dcu.ie/28297/.

Castilho, Sheila & Natália Resende. 2022. Post-Editese in Literary Translations. *Information* 13(2), 66. doi:10.3390/info13020066.

Cattelan, Alessandro. 2017. T-Rank. A decision support system for project assignment. Paper presented at the META-FORUM 2017, Brussels, Belgium. http://www.meta-net.eu/events/meta-forum-2017/pdf/02-Session-06-06-Alessandro-Cattelan.pdf.

Cronin, Michael. 2017. *Eco-translation: translation and ecology in the Age of the Anthropocene.* Abingdon: Routledge.

Daems, Joke, Orphée De Clercq & Lieve Macken. 2017. Translationese and post-editese: how comparable is comparable quality? *Linguistica Antverpiensia, New series: Themes in Translation Studies* 16, 89–103.

DePalma, Donald A. & Vijayalaxmi Hegde. 2010. *The Market for MT Post-Editing*. Boston: Common Sense Advisory.

DePalma, Donald A., Helene Pielmeier, Paul O'Mara & Robert G. Stewart. 2019. *The Language Services Market: 2019*. Boston: CSA Research.

Díaz-Cintas, Jorge & Aline Remael. 2021. *Subtitling: concepts and practices*. Abingdon: Routledge.

Dodge, Jesse, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole DeCario & Will Buchanan. 2022. Measuring the carbon intensity of AI in cloud instances. *2022 ACM conference on fairness, accountability, and transparency*, 1877–1894. Seoul Republic of Korea: ACM. doi:10.1145/3531146.3533234 (accessed 15 July 2022).

Durban, Chris. 2022. Translation, Time, Technology—Who's Counting? *Teaching Translation vs. Training Translator*s, 8. Olomouc, Czech Republic: Palacký University. https://tifo.upol.cz/TIFO2022_book.pdf.

Ehrensberger-Dow, Maureen & Sharon O'Brien. 2015. Ergonomics of the Translation Workplace: Potential for Cognitive Friction. *Translation Spaces* 4(1), 98–118. doi:10.1075/ts.4.1.05ehr.

ELIS Research. 2023. European Language Industry Survey 2023. Brussels: European Union of Associations of Translation Companies. http://elis-survey.org/wp-content/uploads/2023/03/ELIS-2023-report.pdf

Elliot, Bern, Anthony Mullen, Jessica Ekholm & Adrian Lee. 2020. *Market Guide for AI-Enabled Translation Services*. https://www.gartner.com/en/documents/3994429/market-guide-for-ai-enabled-translation-services (accessed 28 September 2021).

Eloundou, Tyna, Sam Manning, Pamela Mishkin & Daniel Rock. 2023. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv*. http://arxiv.org/abs/2303.10130 (28 March, 2023).

Engelbart, Doug. 1962. *Augmenting Human Intellect: A Conceptual Framework*. Washington DC: Air Force Office of Scientific Research. https://www.dougengelbart.org/content/view/138/#2 (accessed 21 September 2022).

Esselink, Bert. 2022. Thirty years and counting: a global industry growing up. *Journal of Internationalization and Localization* 9(1), 87–95.

Felten, Edward W., Manav Raj & Robert Seamans. 2023. Occupational Heterogeneity in Exposure to Generative AI. *SSRN preprint*. https://www.ssrn.com/abstract=4414065 (2 May, 2023).

Fırat, Gökhan. 2021. Uberization of translation: impacts on working conditions. *The Journal of Internationalization and Localization* 8(1), 48–75. doi:10.1075/jial.20006.fir.

Flanagan, Kevin. 2015. Subsegment recall in Translation Memory — perceptions, expectations and reality. *Journal of Specialised Translation* 23, 64–88.

Frey, Carl Benedikt & Michael A. Osborne. 2017. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change* 114, 254–280. doi:10.1016/j.techfore.2016.08.019.

Gough, Joanna, Özlem Temizöz, Graham Hieke & Leonardo Zilio. 2023. Concurrent translation on collaborative platforms. *Translation Spaces* 12(1), 45–73. doi:10.1075/ts.22027.gou.

Graham, Dominic. 2019. Future of the deal: Macro level look at M&A. Paper presented at the ATC Language Industry Summit, London, UK.

Green, Spence, Sida I. Wang, Jason Chuang, Jeffrey Heer, Sebastian Schuster & Christopher D. Manning. 2014. Human Effort and Machine Learnability in Computer Aided Translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1225–1236. Doha, Qatar: Association for Computational Linguistics. doi:10.3115/v1/D14-1130.

Groskop, Viv. 2021 (Oct). Lost in translation? The one-inch truth about Netflix's subtitle problem. *The Guardian*. https://www.theguardian.com/tv-and-radio/2021/oct/14/squid-game-netflix-translations-subtitle-problem (accessed 15 October 2021).

Guerberof-Arenas, Ana & Antonio Toral. 2022. Creativity in translation: machine translation as a constraint for literary texts. *Translation Spaces* 11(2), 184–212. http://doi.org/10.1075/ts.21025.gue.

Haddow, Barry, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl & Alexandra Birch. 2022. Survey of low-resource Machine Translation. *Computational Linguistics* 48(3), 673–732. doi:10.1162/coli_a_00446.

Heilinger, Jan-Christoph, Hendrik Kempt & Saskia Nagel. 2023. Beware of sustainable AI! Uses and abuses of a worthy goal. *AI and Ethics (2023)*. doi:10.1007/s43681-023-00259-8.

Hendy, Amr, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify & Hany Hassan Awadalla. 2023. How Good Are GPT Models at Machine Translation? A Comprehensive Evaluation. *arXiv*. https://arxiv.org/abs/2302.09210.

Henriques, Adrian & Julie Richardson (eds.) 2013. *The Triple Bottom Line*. Abingdon: Routledge.

Herbert, Sarah, Félix do Carmo, Joanna Gough & Anu Carnegie-Brown. 2023. From responsibilities to responsibility: a study of the effects of translation workflow automation. *Journal of Specialised Translation* 40, 9-35.

Hutchins, W. John. 1995. Machine translation: A brief history. In Ernst Koerner & Ronald Asher (eds.), *Concise history of the language sciences: from the Sumerians to the cognitivists*, 431–445. Oxford: Pergamon.

Isabelle, Pierre, Marc Dymetman, George Foster, Jean-Marc Jutras & Elliott. 1993. Translation Analysis and Translation Automation. In Proceedings of the Fifth Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages. Kyoto, Japan. https://aclanthology.org/1993.tmi-1.17.

Jiang, Yuchen, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Sennrich, Ryan Cotterell, Mrinmaya Sachan & Ming Zhou. 2022. BlonDe: An Automatic Evaluation Metric for Document-level Machine Translation. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1550–1565. Seattle: Association for Computational Linguistics. https://aclanthology.org/2022.naacl-main.111.

Jooste, Wandri, Rejwanul Haque & Andy Way. 2022. Knowledge Distillation: A Method for Making Neural Machine Translation More Efficient. *Information* 13(2), 88. doi:10.3390/info13020088.

Kagermann, Henning & Wolfgang Wahlster. 2022. Ten Years of Industrie 4.0. *Sci* 4(3), 26. doi:10.3390/sci4030026.

Kay, Martin. 1997. The proper place of men and machines in language translation. *Machine Translation* 12(1/2), 3–23.

Kelleher, John D., Brian MacNamee & Aoife D'Arcy. 2015. *Fundamentals of machine learning for predictive data analytics: algorithms, worked examples, and case studies*. Cambridge: The MIT Press.

Kenny, Dorothy (ed.). 2022. *Machine translation for everyone: Empowering users in the age of artificial intelligence*. Berlin: Language Science Press.

Kocmi, Tom & Christian Federmann. 2023. Large Language Models Are State-of-the-Art Evaluators of Translation Quality. *arXiv*. https://arxiv.org/abs/2302.14520.

Koponen, Maarit, Umut Sulubacak, Kaisa Vitikainen & Jörg Tiedemann. 2020. MT for subtitling: User evaluation of post-editing productivity. *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, 115–124. Lisboa, Portugal: European Association for Machine Translation. https://aclanthology.org/2020.eamt-1.13 (accessed 13 July 2021).

Koskinen, Kaisa & Minna Ruokonen. 2017. 1 Love letters or hate mail? Translators' technology acceptance in the light of their emotional narratives. In Dorothy Kenny (ed.), *Human issues in translation technology*, 8–24. Abingdon: Routledge.

Krüger, Ralph & Janiça Hackenbuchner. 2022. Outline of a didactic framework for combined data literacy and machine translation literacy teaching. *Current Trends in Translation Teaching and Learning E*, 375–432. doi:10.51287/cttl202211.

Lange, Jeva. 2021 (Oct). Squid Game and Netflix's subtitle problem. *The Week*. https://theweek.com/culture/1005620/squid-game-and-netflixs-subtitle-problem (accessed 15 December 2021).

Läubli, Samuel, Sheila Castilho, Graham Neubig, Rico Sennrich, Qinlan Shen & Antonio Toral. 2020. A Set of Recommendations for Assessing Human–Machine Parity in Language Translation. *Journal of Artificial Intelligence Research* 67. https://jair.org/index.php/jair/article/view/11371 (accessed 30 September 2022).

Laurent, Lina. 2023. Historic AV translation collective agreement: includes minimum rates for self-employed. Translated by Ian Mac Eochagáin. Helsinki: Union of Journalists in Finland. https://journalistiliitto.fi/en/historic-av-translation-collective-agreement-includes-minimum-rates-for-self-employed/ (accessed 8 May 2023).

Littman, Michael L., Ifeoma Ajunwa, Guy Berger, Craig Boutilier, Morgan Currie, Finale Doshi-Velez, Gillian Hadfield, Michael C. Horowitz, Charles Isbell, Hiroaki Kitano, Karen Levy, Terah Lyons, Melanie Mitchell, Julie Shah, Steven Sloman, Shannon Vallor, Toby Walsh. 2021. *Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report*. Stanford: Stanford University. https://ai100.stanford.edu/2021-report/gathering-strength-gathering-storms-one-hundred-year-study-artificial-intelligence (accessed 28 September 2021).

Lommel, Arle. 2021. Responsive Machine Translation: The Next Frontier for MT. Boston: CSA Research. https://csa-research.com/Blogs-Events/Blog/responsive-machine-translation (accessed 28 September 2021).

Lommel, Arle & Donald A. DePalma. 2021. *Augmented Translation: How Artificial Intelligence Drives Productivity and Efficiency for the Language Industry*. Boston: CSA Research. https://insights.csa-research.com/reportaction/305013226/Toc.

Marking, Marion. 2022 (Jan). Netflix COO Reveals Scale of Dubbing and Subtitling Operations. *Slator*. https://slator.com/netflix-coo-reveals-scale-of-dubbing-subtitling-operations/ (accessed 25 January 2022).

Matusov, Evgeny, Patrick Wilken & Yota Georgakopoulou. 2019. Customizing Neural Machine Translation for Subtitling. *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, 82–93. Florence, Italy: Association for Computational Linguistics. https://aclanthology.org/W19-5209 (accessed 13 July 2021).

Mc Gowran, Leigh. 2023 (Apr). EU creates privacy task force focused on ChatGPT. *Silicon Republic*. https://www.siliconrepublic.com/machines/chatgpt-eu-privacy-task-force-ai-openai.

Mehta, Sneha, Bahareh Azarnoush, Boris Chen, Avneesh Saluja, Vinith Misra, Ballav Bihani & Ritwik Kumar. 2020. Simplify-then-Translate: Automatic Preprocessing for Black-Box Machine Translation. *arXiv* http://arxiv.org/abs/2005.11197 (accessed 13 July 2021).

Melby, Alan. 2022. Grades of Translation as a Marketing Strategy in the Era of AI. *Tool Box Journal* 343. https://www.tranquality.info/wp-content/uploads/2022/12/343rd-issue-of-The-Tool-Box-Journal_AMelby.pdf (accessed 22 December 2022)

Melé Carné, Domènec. 2020. *Business ethics in action: managing human excellence in organizations*. 2nd edition. London: Red Globe Press.
Metz, Cade. 2015 (Jun). AI's Next Frontier: Machines That Understand Language. *Wired*. https://www.wired.com/2015/06/ais-next-frontier-machines-understand-language/ (accessed 10 September 2022).
Moorkens, Joss. 2020. "A tiny cog in a large machine": Digital Taylorism in the translation industry. *Translation Spaces* 9(1), 12–34. doi:10.1075/ts.00019.moo.
Moorkens, Joss & Dave Lewis. 2019. Research Questions and a Proposal for the Future Governance of Translation Data. *Journal of Specialised Translation* 32, 2–25.
Moorkens, Joss & Marta Rocchi. 2020. Ethics in the translation industry. In Kaisa Koskinen & Nike K. Pokorn (eds.), *The Routledge Handbook of Translation and Ethics*, 320–337. Abingdon: Routledge.
Mumford, Lewis. 1967. *The myth of the machine. 1: Technics and human development*. New York: Harcourt, Brace, Jovanovich.
Nassif, Ali Bou, Ismail Shahin, Imtinan Attili, Mohammad Azzeh & Khaled Shaalan. 2019. Speech Recognition Using Deep Neural Networks: A Systematic Review. *IEEE Access* 7, 19143–19165. doi:10.1109/ACCESS.2019.2896880.
Nieminen, Tommi. 2022. OPUS-CAT: Free, secure, and confidential MT for professional translators. Paper presented at the CRITT Translation Colloquium, University of Kent. https://sites.google.com/site/centretranslationinnovation/conferences-workshops/colloquium?authuser=0#h.ct5sq6zcwdt.
NLLB team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk & Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation. arXiv. http://arxiv.org/abs/2207.04672 (accessed 25 August 2022).
Ouyang, Long, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike & Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *arXiv* https://arxiv.org/abs/2203.02155
Parasuraman, Raja, Thomas B Sheridan & Christopher D Wickens. 2000. A Model for Types and Levels of Human Interaction with Automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 30(3), 13.
Peng, Keqin, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang & Dacheng Tao. 2023. Towards Making the Most of ChatGPT for Machine Translation. *SSRN preprint*. https://www.ssrn.com/abstract=4390455 (accessed 27 March 2023).
Pielmeier, Hélène & Paul O'Mara. 2020. *The state of the linguist supply chain*. Boston: CSA Research.
Prodoscore. 2022. The Happiness Index. Prodoscore. https://www.prodoscore.com/wp-content/uploads/2022/05/Happiness-Index.pdf.
Pym, Anthony. 2004. Propositions on cross-cultural communication and translation. *Target. International Journal of Translation Studies* 16(1), 1–28. doi:10.1075/target.16.1.02pym.
Pym, Anthony & Ester Torres-Simón. 2021. Is automation changing the translation profession? *International Journal of the Sociology of Language* 2021(270), 39–57. doi:10.1515/ijsl-2020-0015.

Radford, Alec, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey & Ilya Sutskever. 2022. Robust Speech Recognition via Large-Scale Weak Supervision. *arXiv* https://arxiv.org/abs/2212.04356

Robertson, Samantha & Mark Díaz. 2022. Understanding and Being Understood: User Strategies for Identifying and Recovering From Mistranslations in Machine Translation-Mediated Chat. *2022 ACM Conference on Fairness, Accountability, and Transparency*, 2223–2238. Seoul: ACM. doi:10.1145/3531146.3534638. https://dl.acm.org/doi/10.1145/3531146.3534638 (12 September, 2022).

Rothwell, Andrew, Joss Moorkens, Maria Fernandez Parra, Joanna Drugan & Frank Austermühl. 2023. *Translation Tools and Technologies: A Practical Guide for Students and Translators.* Abingdon; Routledge.

Russell, Stuart. 1999. Rationality and Intelligence. In Michael Wooldridge & Anand Rao (eds.), *Foundations of Rational Agency*, 14, 11–33. Dordrecht: Springer Netherlands. doi:10.1007/978-94-015-9204-8_2.

Saeed, Muhammad Ahmed, Eloy Rodríguez González, Tomasz Korybski, Elena Davitti & Sabine Braun. 2022. Connected yet Distant: An Experimental Study into the Visual Needs of the Interpreter in Remote Simultaneous Interpreting. In Masaaki Kurosu (ed.), *Human-Computer Interaction. User Experience and Behavior* (Lecture Notes in Computer Science), 214–232. Champaign: Springer. doi:10.1007/978-3-031-05412-9_16.

Schjoldager, Anne, Hanne Dam-Jensen, Vilelmini Sosoni, Michael Tieber, Maarit Koponen & Ana Guerberof Arenas. 2022. The HAL panel: What if we had the perfect machine translation system? A discussion of the future scenarios in translation and technology. Roundtable discussion at the NETTT Conference, Rhodes. July 2022.

Schwartz, Eric Hal. 2023 (Jan). Amazon Employees Using ChatGPT for Coding and Customer Service Warned Not to Share Company Information With AI Chatbot. *Voicebot.ai*. https://voicebot.ai/2023/01/27/amazon-employees-using-chatgpt-for-coding-and-customer-service-warned-not-to-share-company-information-with-ai-chatbot/ (8 May, 2023).

Schwartz, Roy, Jesse Dodge, Noah A. Smith & Oren Etzioni. 2020. Green AI. *Communications of the ACM* 63(12), 54–63. doi:10.1145/3381831.

Sherionov, Dimitar & Eva Vanmassenhove. 2022. The Ecological Footprint of Neural Machine Translation Systems. arXiv. https://arxiv.org/abs/2202.02170 (accessed 30 September 2022).

Shterionov, Dimitar, Félix Do Carmo, Joss Moorkens, Eric Paquin, Dag Schmidtke, Declan Groves & Andy Way. 2019. When less is more in Neural Quality Estimation of Machine Translation. An industry case study. *Proceedings of Machine Translation Summit XVII: Translator, Project and User Tracks*, 228–235. Dublin, Ireland: European Association for Machine Translation. https://aclanthology.org/W19-6738.

Soloviev, Kirill. 2022. PEMT human evaluation at 100x scale with risk-driven sampling. Paper presented at the 15th Conference of the Association for Machine Translation in the Americas, Orlando, USA.

Stasimioti, Maria. 2022 (May). What to Do About the Talent Crunch in Media Localization. *Slator*. https://slator.com/what-to-do-about-the-talent-crunch-in-media-localization/ (accessed 22 December 2022).

Strubell, Emma, Ananya Ganesh & Andrew McCallum. 2019. Energy and Policy Considerations for Deep Learning in NLP. *arXiv*. http://arxiv.org/abs/1906.02243 (accessed 9 September 2022).

Szarkowska, Agnieszka, Krzysztof Krejtz, Łukasz Dutka & Olga Pilipczuk. 2018. Are interpreters better respeakers? *The Interpreter and Translator Trainer* 12(2), 207–226. doi:10.1080/1750399X.2018.1465679.

Tiedemann, Jörg & Santhosh Thottingal. 2020. OPUS-MT -- Building open translation services for the World. In André Martins, Helena Moniz, Sara Fumega, Bruno Martins, Fernando Batista, Luisa Coheur, Carla Parra Escartín, Isabel Trancoso, Marco Turchi, Arianna Bisazza, Joss

Moorkens, Ana Guerberof Arenas, Mary Nurminen, Lena Marg, Mikel L. Forcada (eds.) *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation,* 479–480. https://helda.helsinki.fi/handle/10138/327852 (accessed 11 September 2022).

Toral, Antonio. 2019. Post-editese: an Exacerbated Translationese. *arXiv*. http://arxiv.org/abs/1907.00900 (accessed 29 June 2020).

UNESCO. 2014. Skills for holistic human development. UNESCO Education Policy Brief Vol.2. https://unesdoc.unesco.org/ark:/48223/pf0000245064.

Vagia, Marialena, Aksel A. Transeth & Sigurd A. Fjerdingen. 2016. A literature review on the levels of automation during the years. What are the different taxonomies that have been proposed? *Applied Ergonomics* 53 Pt A, 190–202. doi:10.1016/j.apergo.2015.09.013.

Van der Meer, Jaap. 2021 (July). Translation Economics of the 2020s. *MultiLingual*. https://multilingual.com/issues/translation-economics-of-the-2020s/ (accessed 28 September 2021).

Vanmassenhove, Eva, Dimitar Shterionov & Andy Way. 2019. Lost in Translation: Loss and Decay of Linguistic Richness in Machine Translation. *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, 222–232. Dublin: European Association for Machine Translation. https://www.aclweb.org/anthology/W19-6619.

Vela-Valido, Jennifer. 2021. Translation quality management in the AI Age. New technologies to perform translation quality assurance operations. *Revista Tradumàtica* (19), 93–111. doi:10.5565/rev/tradumatica.285.

Vieira, Lucas Nunes. 2020. Automation anxiety and translators. *Translation Studies* 13(1), 1–21. doi:10.1080/14781700.2018.1543613.

Wakabayashi, Daisuke & Cade Metz. 2022 (May). Dismissal Fuels Split In A.I. Unit At Google. *The New York Times*. https://www.nytimes.com/2022/05/02/technology/google-fires-ai-researchers.html

Wheelen, Thomas L., J. David Hunger, Alan N. Hoffman & Charles E. Bamford. 2018. *Concepts in strategic management and business policy: globalization, innovation, and sustainability*. Fifteenth edition. Harlow: Pearson.

Wu, Carole-Jean, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga Behram, James Huang, Charles Bai, Michael Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin S. Lee, Bugra Akyildiz, Maximilian Balandat, Joe Spisak, Ravi Jain, Mike Rabbat & Kim Hazelwood. 2021. Sustainable AI: Environmental Implications, Challenges and Opportunities. *arXiv*. https://arxiv.org/abs/2111.00364 (accessed 30 September 2022).