Conclusion

In June 2025, as I was approaching the final stages of writing this book, I shared some of the data collected for the project at a meeting in the UK Parliament. The meeting was organised by the All-Party Parliamentary Group on Modern Languages in the House of Lords. My talk examined some of the examples I discussed in this book, including cases where pharmacists used machine translation to explain medication dosing and where social workers used it when nothing else seemed practical. That was not my first time at a parliamentary event. In January 2023, I had spoken at a similar meeting where I shared some previous data I had collected. The data I shared in 2023 included the example I discussed in Chapter 4 in which a mental health professional used Google Translate when detaining a patient and explaining their rights. Parliamentarians and staff from different government departments were present in both meetings, but the second meeting felt different. Although in early 2023 the need to make decisions about AI was coming fast on the horizon, it was probably not yet at the top of the pile in politicians' in-trays. In 2025, the topic felt more urgent. Artificial intelligence more generally was a subject of major political and societal concern.

Concerns raised by the information which I and others shared in the 2025 meeting prompted one of the attendees to ask: wouldn't it be safer for communities to learn English? Wouldn't this be in their interest? Those who see the value of multilingualism may find this question frustrating. The premise of the question is problematic for its potential Anglocentric angle, not to mention the apparent assumption that non-English speakers experience language barriers by choice. But for several reasons I would argue that it is important to consider questions of this nature. First because many individuals, especially those with high linguistic capital (see Chapter 6), will ask themselves the same. Visions of a monolingual society are not necessarily ideological. Many of those who speak widely spoken global languages will not be well placed to comprehend the challenges posed by language barriers. English speakers are particularly likely to be in this privileged category. Second, the question asked in Parliament is relevant because the straightforward answer to it is 'yes' – it is safer to speak the local language than to not speak it. It is in individuals' interest

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

to not be reliant on machine translation systems if they can help it. But what the question misses is how difficult and time-consuming it can be to learn a language. It also misses the fact that language access is not only for newcomer communities who will stay in the country. It is for the benefit of any individual, whether they be a tourist who walks into a pharmacy or someone fleeing danger who cannot wait until they master a new language to access critical services. Perhaps one way to turn the question around would be to ask what kind of society we wish to be: one that imposes knowledge of the local language as a precondition for different types of assistance, or one that is able to help, protect and care for others irrespective of their language?

While many would have no qualms about the first type of society, I would argue that AI can be part of the second. But as is usually the case with technology, there are no benefits without risks, and the risks in the case of machine translation are particularly difficult to navigate. As I mentioned in the Introduction, this is a complex subject for which there are no hard and fast solutions. My goal with my presentations in the UK Parliament, and with this book more broadly, was never to solve a problem, but rather to analyse and document it – to understand what machine translation represents in public service contexts that can be highly consequential for large numbers of people. Below I summarise how uses of machine translation in these contexts can be categorised. I also outline the types of decision-making considerations that these uses call for.

Uses of Machine Translation

The uses of machine translation discussed in this book can be placed along two axes. The first one represents the extent to which users rely on machine translation as a language mediator. The second axis represents the communicative context, which varies between low- and high-stakes communicative purposes. These two axes are illustrated in the graph shown in Figure C.1. The examples discussed in previous chapters cover a wide area of this graph. The level of reliance on AI will be low in some cases because users may have some knowledge of the language or may be able to rely to a larger extent on other communication methods. By contrast, machine translation can also be the main communication or information access method in use.

Quadrants A and B largely coincide with the types of ancillary communication discussed in Chapter 3: using machine translation to tell the time, to give directions around a hospital or for small talk. Quadrants A and B may also include uses of machine translation that take place before a human linguist gets involved, such as researching a subject before an appointment or just telling others that a human interpreter is on the way. Whether the reliance on machine translation is low (quadrant A) or high (quadrant B), in all such cases the types

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

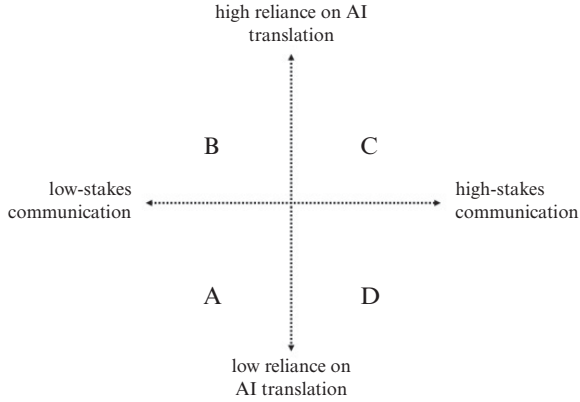


Figure C.1 Extent of AI translation use verses low- and high-stakes communication.

of communication and information access it facilitates would have been unlikely to happen as effectively before the technology became available. These uses of machine translation are also unlikely to compete with the work of human linguists given their low risk and ad hoc nature.

Quadrant C, on the other hand, is more problematic. Here machine translation is relied upon for high-stakes communication. This may happen for several reasons. Maybe it was impractical for linguists to provide timely assistance – for example, in emergencies. Maybe the risks of using machine translation were miscalculated. Perhaps a traffic police officer decided to ignore the professionally translated consent form they had in their vehicle and decided to use machine translation instead. Or maybe a healthcare manager thought that information leaflets could be easily translated with AI and that investing in professional translations was not necessary. Whatever the reason, and whether the users appreciate the risks or not, the use of machine translation in quadrant C is unideal.

Quadrant D also represents high-stakes communication but here communication is less reliant on AI. Maybe the machine translations are checked by bilingual staff. Or perhaps a human interpreter mediates an interaction but, like some of the social workers in Chapter 6, the interactants are not confident in the interpreter's ability so AI is used as a second source of translations. Although many of the project participants had had negative experiences with human linguists, AI is unlikely to be safer than qualified language professionals, so using it as a way of verifying human translations is a precarious strategy. The

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

use of machine translation in quadrant D may be secondary to other communication methods, but it is not necessarily for that reason less risky.

The dotted line dividing the quadrants in Figure C.1 represents their fluid relationship. As discussed in Chapter 3, what starts as ancillary communication can easily morph into core care. The project's participating professionals may start an interaction believing that it is fine to rely on machine translation and that there is no need for a professional linguist to be involved. They may then be surprised by the complexity of the circumstances and by the level of risk posed by potential misunderstandings. At that point they will need to decide whether to carry on relying on machine translation anyway or whether to interrupt what they are doing until a linguist becomes available.

Several factors will weigh on these decisions. Urgency is a significant one among them. When professionals are taken by surprise and must think on their feet, they may need to use any means available to communicate. But not all uses of machine translation are marked by urgency. When AI is used to translate letters or leaflets, for instance, a lack of linguistically proficient oversight is particularly problematic.

The difference between an urgent interaction and an information leaflet is captured by the categories of machine-translated communication outlined in Chapter 1. In many cases translation users will have had the time to check that the translations were accurate, well formed, respectful and, crucially, that they inspired confidence and trust. The contexts where this type of checking is possible, or indeed necessary, will often involve asynchronous, unidirectional communication – for example, leaving a pile of leaflets on a reception desk, attaching announcements to a notice board, issuing statements on a website or providing individuals with documents to which they are not expected to reply. This type of communication can also involve audio and video content – for example, recorded telephone messages played to callers or instructive videos published online.

Asynchronous, unidirectional communication is important to get right because information communicated in this way will often come across as the official position or advice of the institution. Although public service workers are always acting on behalf of their organisations to some degree, when information is officially distributed asynchronously and unidirectionally, it is often the institution itself that is issuing the material. The tolerance for errors in these cases is lower while the level of risk posed by inaccuracies is likely to be higher because it may be difficult for the information consumers to seek clarifications. Additionally, these messages often address more than one individual. They will have many ears and eyes on them, so their impact is likely to be significant. The fact that content of this nature is consumed by many individuals bolsters the case for having the content professionally checked by a linguist. Getting the content right once will allow it to be used and reused

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

many times. The costs associated with ensuring that translations of this nature can be trusted are likely to be lower than the legal costs that mistranslations may entail.

In most everyday contexts, the likelihood that whatever we say may be legally scrutinised is probably minimal. But in the contexts analysed in this book, this possibility needs to be carefully considered. Most interactions I examined in the project could be legally significant if the service provided fell short of expected standards and the service user were negatively affected as a result. For the social workers featured in Chapter 6, this was a key reason for avoiding machine translation. But as mentioned in Chapter 4, machine translation is used even to translate consent forms. The legal value of a consent form is intended from the outset. The signed form acts as a written record confirming that individuals have agreed to something that could have some type of adverse effect on them. Uses of AI to translate consent forms are therefore alarming. Given the relatively static nature of these forms, and the fact that the same consent text can be used multiple times, the risks posed by machine translation here are not only concerning but also probably avoidable. Getting the forms professionally translated is desirable and cost-effective.

Not all written records can be timely translated by a linguist, however. In Chapter 4, I mentioned the case of a healthcare professional who used machine translation to code information into a patient's medical record. This use of machine translation is potentially highly consequential because inaccuracies could be undetectable and influence decisions that can affect the individual's health. It may be that language barriers take healthcare professionals by surprise, as mentioned at different points in the book. Maybe the patient overestimated their level of proficiency in the local language and difficulties arose mid-consultation when an interpreter was not available at short notice. Or maybe the patient's condition was not life-threatening, and the recommendation made by the healthcare professional was unlikely to pose significant risks. In any such case, if medical advice, or a diagnosis, rests on a patient's medical history, adding details of this history to the patient's record may be legally or professionally necessary. It may be preferable to add machine translations to the patient's record in this case than to not document the rationale for the decision at all.

The overtness principle discussed in Chapter 1 would at this point need to be considered. Transparency is one of the principles described in discussions of AI and machine translation use in professional contexts.[1] Whenever machine translation is used, it is important for its use to be declared. Potential consumers of the information need to be made aware of how the translation was generated. When machine translation tools are used in synchronous interactions that take place in a shared physical space, the presence of the tool will in most cases be obvious. This is not necessarily the case for written records, so the healthcare

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

professional mentioned in the scenario above would probably mitigate some of the technology's risks – or at least allow the risks to be identified or verified later – by adding the untranslated version of the information to the patient's record together with the machine translation and a note declaring that the translation was automatically generated and should therefore be approached with caution.

There are, then, some general principles that can be followed when navigating the different quadrants shown in Figure C.1. In brief, these principles include:

1. The stability of the information – how variable or static it is – and thereby its potential to be reused;
2. The urgency with which the content is needed;
3. The number of eyes and ears that will consume it; and
4. Its intended or potential legal value.

Specifically, for content that is static, non-urgent, reusable, and which may have a legal status affecting multiple individuals, there will be very few justifiable reasons for using machine translation at all. Furthermore, where machine translation *is* part of the communication strategy, it will be important to ensure:

5. That the risks of not using it are likely to be greater than the risks of using it, including in terms of risks to privacy and confidentiality;
6. That its use is overt; and
7. That it is as far as possible consensual, and that affected individuals have played a part in the decision-making.

As mentioned in Chapter 3 and in Chapter 4, it is also important:

8. To consider whether the use of machine translation can be complemented by other communication methods in ways that reduce the overall level of risk; and
9. To attempt to ensure that any content to be translated is as simple and plainly expressed as possible.

Service providers and their teams may need to be alerted to these principles, and policies can be useful in this process. Policies may raise awareness of previously unforeseen dangers. They may also specify who to contact for help. They may define the types of staff training that are required or how often the policy itself needs to be revised. Policies may also define lines of accountability and help managers to ask the right questions if they decide to procure machine translation systems.[2] But I would argue that policies cannot be expected to eradicate uncertainty. Policies on their own are unlikely to be of much value.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

Individuals will always require prudential judgement, and for that they are likely to need appropriate frameworks of support.

Supporting Front-Line Workers

When I started working on this project, I was very aware of the risk that the research could be hijacked by attempts to criticise public service professionals. Asking professionals whether they use machine translation at work is not easy. It is, in a way, asking them to admit to resorting to potentially substandard practices, so the question involves delicate considerations. Especially in my interviews with the social workers – where I spoke to participants directly – I sensed some initial hesitation in their answers. Many of them had never been advised by their superiors on the question of machine translation use, so they were talking about practices which, albeit necessary, were potentially unorthodox. Their desire to exchange information and get to know more was nevertheless stronger than any trepidation they might have had. I sometimes had to remind them – and myself – that my role was to listen rather than to advise or provide my own analysis on the spot. When I sensed that they were making the wrong assumptions – for example, that they could use machine translation whenever needed since it was being researched and was therefore probably accepted – I would clarify any points post-interview after they had had the chance to speak their minds. As mentioned in Chapter 6, many of them would at that point reflect on our conversation. Sometimes the study itself prompted them to think about questions they had not considered. Throughout the project, it was clear that not only were the participating professionals operating under pressure, but also that they had very little guidance from their employers on machine translation use.

A policy can remind professionals of principles like those outlined above, but it would be unlikely to reliably prescribe what these professionals should do every time. What should a paediatrician do when a parent only feels comfortable to ask questions once the telephone interpreter, booked weeks in advance, is no longer on the line? What should a social worker do when, following an anonymous tip-off, a language barrier takes them by surprise? What should a nurse do when anxious family members come to the ward asking questions about a loved one in a different language? Rather than judging these professionals with the benefit of hindsight, it is important to support them in the process of living well with AI. As discussed in Chapter 2, a virtue-ethics approach to this subject involves deliberate moral habituation: 'repeated moral practice of right (or nearly right) conduct' (emphasis omitted).[3] Engaging critically with the dilemmas discussed in this book is a process. There is always room for revision and improvement, which are likely to be best addressed in the context of training and education.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

Although proposing a syllabus for workplace training is outside the scope of this project, the AI/machine translation literacy development initiatives mentioned in Chapter 5 would be good places to start. Users need to be familiar with the probabilistic nature of machine translation models. A good way of thinking about it is that these models are always making an informed guess based on the data available, and the data available can be incomplete, inaccurate or just incompatible with the use context. The possibility of AI mistranslations also needs to be considered alongside the needs of different services or sectors. Remote communication, as described in Chapter 1, may pose greater risks compared to exchanges that take place in a shared physical space, where body language, pictorial guidance and other communication methods can be more easily combined. Synchronous communication is also likely to be marked by urgency which, as mentioned, prompted many of the participants to act fast and decisively in difficult circumstances. Training needs to be about supporting individuals to operate within these circumstances and make the best decisions possible. It would be unfair, not to mention counterproductive, to penalise individuals for adopting their own coping strategies or for failing to implement ideal practices that may not reflect their daily realities.

Who Is Accountable?

I addressed the question of accountability at different points in the book. I first looked at it in Chapter 1 in relation to organisational factors in machine translation use. I also addressed it in Chapter 2 in relation to how humans are ultimately responsible for uses of technologies despite the technologies' own persuasive properties. Although lines of accountability are at present often unclear, defining accountability is important because it will be precisely the accountable body, individual or group who will be invested in ensuring that AI is used responsibly.

At a practical level, however, it is also true that it may be difficult for accountability to be categorically imposed. Decisions on who is accountable may be subject to different agreements and responsibilities. It can also be subject to legal disputes. But what the discussion provided here makes clear is that technology providers are highly skilled at ensuring that they, for one, are not legally responsible for the negative consequences of their products.

Developers may argue that holding them accountable for the inaccuracies of their tools is inconsistent with how the technology works because their models are always guessing and are highly likely to make mistakes. The terms and conditions of use of machine translation tools may recognise in their small print that the translations can be incorrect. The small print is also likely to disclaim any responsibilities that the developers may be expected to have. But these same companies often make bold and inflated claims about the performance of

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

their products. Although it may be impossible for them to guarantee machine translation accuracy, transparency is within their power. Transparency will often clash with financial interests, however, which poses a great challenge to safe applications of AI more broadly. Discussions of concepts such as surveillance capitalism,[4] AI capitalism[5] and techno-capitalism[6] expose how profit-seeking environments can be inhospitable to privacy, inclusivity, social justice and many other values that are necessary for living well with technology. So expecting large, for-profit corporations to hold themselves accountable is unlikely to bear fruit.

Governments and public organisations too have institutional responsibilities, especially in assessing the risks created by the technologies that they deploy or that they fail to securely make available. Although individual malpractice will arise and should be addressed, if staff have never had the opportunity to discuss uses of machine translation and are left to their own devices, expecting them to take individual responsibility for misuses of machine translation would be an act of institutional omission.

In the UK, the government department that oversees technologies and their use is the Department for Science, Innovation and Technology.[7] In response to a recent parliamentary question about policy discussions on machine translation, a government minister from this department stated that 'the vast majority of AI systems – including machine translation – are best regulated at the point of use and in the context in which they are deployed'.[8] This statement is in many ways sensible. The context-dependent nature of AI use decisions is after all one of my central arguments in this book. But would it not be important for the government to take a proactive role in supporting this decision-making? Do separate public services have specific teams tasked with considering the question of AI and language barriers? Who should the professionals I consulted, or their managers and their managers' managers, speak to about continuing professional development and the risks that machine translation poses to them and to their service users? Perhaps most importantly, who is in charge of ensuring that teams within relevant sectors have the resources they need to deploy AI appropriately and to seek more reliable language services where necessary?

The question of funding is often unavoidable in discussions of this subject. Another risk of which I was all too aware when embarking on the project was that any argument in favour of *some* uses of machine translation, however restrictive or qualified, could be used to justify cost-cutting measures. Many governments and public institutions around the world are going through challenging times financially. Managers will have many competing priorities on their hands, not least the provision of the service they need to deliver, whether it be healthcare, policing or social services. But some of the most beneficial uses of machine translation are additional to the work of linguists – uses that are

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

likely to fall into quadrants A and B in Figure C.1. If these uses do not directly compete with professional language mediation, it follows that they are unlikely to lead to significant savings beyond those accrued by easier access to information in ways that, before machine translation, would have been harder or impossible. While the presence of machine translation in quadrants A and B is potentially welcome, in quadrant C it is a coping mechanism that should not be pursued as a strategy. As a general principle, machine translation should ideally add rather than replace. Powerful institutions need to ensure that its benefits are not mis-sold or abused. Accountability starts at the top.

Recognising Risks and Benefits

Translation and interpreting are for the most part unregulated occupations – in the UK and in many other countries.[9] In most contexts, no one needs a licence to provide language services, which over the years has called for a certain level of self-assertion among translation and interpreting academics and practitioners. Statements on how language services should be left to the experts are often made at translation and interpreting conferences and events. I have seen translators provide cautionary analyses of what happens to businesses that decide to do translation on the cheap. I have also seen academics sound the alarm about the risks of using translations that are not of the highest possible standard and have myself made this point in relation to many communicative contexts. But the truth is that translation and interpreting experts alone do not have all the answers. Analyses of this subject call for knowledge and information that those making the decisions on the ground will be best placed to provide. Some of the information the project participants kindly shared with me points to two potential fallacies in debates on this topic. The first is that paid language services provided by humans are always trustworthy. The second is that not using machine translation is always safer than using it.

The first fallacy often stems from assumptions that all humans who provide language services are qualified and have gone through rigorous training. This is true in many cases, especially where formal qualifications and training are available. But as discussed in Chapter 6, low levels of linguistic capital and international prestige affect the entire ecosystem of a language, including the availability of training and thereby of qualified professionals. Cost-cutting measures and outsourced models of service provision also deteriorate linguists' working conditions. Worse working conditions make language jobs less attractive to qualified professionals and may make the work of those who do accept these jobs less reliable. Many of the experiences described by the professionals consulted in this project are a sobering reminder that human language services are not guaranteed to be effective. From interpreter no-shows to cases of mistranslation in interpreted interactions, the project participants were not

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008

short of negative stories to tell. The fact that machines can be untrustworthy does not mean that humans can always be trusted – not while we continue to operate within organisational frameworks that have cost efficiency as a core value. The market logic at the root of this problem cannot be ignored.

Similarly, the fact that qualified human linguists are more trustworthy than machines does not mean that machines are unhelpful. At several points in this project I came across instances where machine translation plugged gaps. It was used to calm people down and to reassure vulnerable individuals. It was used for introductions, to reschedule appointments and to provide basic urgent advice. In brief, it allowed communication to happen when nothing else was available.

Machine translation therefore has benefits. Its availability is highly appreciated, which makes the question of its use a difficult one to address, especially in contexts where human language professionals have been historically undervalued. For those who research this topic, there is a risk that whatever is said about it can be received dogmatically by different groups. Arguments in favour of some machine translation use may be met with resistance from translation and interpreting scholars and practitioners. Arguments against machine translation may be met with resistance from technology enthusiasts and developers, especially those who have a vested interested in promoting their products. But risks and benefits cannot be decoupled. The challenges I documented here cannot be addressed by extremes. The balance will tip in favour of using machine translation in some cases and of avoiding it in others. Using it responsibly means accepting this complexity and seeking to cultivate the wisdom to contextually determine where the balance lies.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 10 Nov 2026 at 07:59:04, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.008