5 Trust

About two thirds of the professionals who took part in the project had not used machine translation at work. Some of them avoided it on purpose. Many simply had not needed it. Some declared no prior use of machine translation but then stated that they had in fact used it. Although the study description indicated that the project was about any written or spoken uses of machine translation in one of the selected sectors, in some cases participants mistakenly assumed that their use of it was not of interest: 'I have used Google Translate before at work but not on written text' (Healthcare and Social Assistance). There were also participants who remembered after the fact that they had used machine translation before:

> Yes! In fact, I have used Google Translate to communicate with an Italian woman before who has dementia and can no longer understand English as well as she used to. This helps to overcome the communication barrier. Automatic translators can be extremely helpful for situations like this and even help another person who speaks the same language to understand more clearly. (Healthcare and Social Assistance)

Since these participants initially declared having no direct experience with machine translation at work, they were not included in the sample of users analysed in Chapter 3 and Chapter 4. The extent of machine translation use revealed in this project is therefore a conservative estimate. A larger number of professionals are likely to have used it, whether knowingly or not. But more importantly, lapses of this nature show that, for many of these professionals, uses of machine translation are not particularly memorable. Thinking of this technology is not a priority for them given the broader service that they need to provide, which gives a clue into how some of the high-risk examples discussed in Chapter 4 would have come about. Challenging circumstances are often met with little preparatory reflection on what machine translation is and what consequences it may have. The opinions of non-users are therefore a valuable window into the types of mental dispositions that precede the machine translation uses previously discussed – into what leads someone to trust machine translation as a viable communication solution.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

In what follows, I examine the concept of trust and how it can be applied to uses of machine translation as a communication tool. I identify some of the trust-influencing factors considered by those who had not used machine translation at work. I also discuss workplace training by turning back to the opinions of machine translation users. I asked them to reflect on their experience of using the technology to suggest what kind of training they would have found helpful. I analysed their answers to identify any practices or types of information which, in their view, may be able to improve their trust judgement abilities.

What Is Trust?

Like a few other concepts discussed in the book so far, trust has been examined somewhat disparately across fields. Its most well-known, cross-disciplinary definition is probably the one proposed by Denise M Rousseau and colleagues, who refer to trust as 'the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another'.[1] Trust here is about accepting that we are vulnerable in some way and concluding that others are acting or intend to act in our interest. Trust in machine translation in this sense would be the product of calculated decisions. If users do not consider the risks, then their reliance on machine translation can be described as 'blind trust' or,[2] as I put it in this chapter, trust that is uncritical. In other words, if users are unaware that they are running a risk, then trust is not needed because users in that case see no vulnerability. They proceed without fear or hesitation.

The 'intentions or behaviour of another' in the definition by Rousseau and colleagues can have several proxies. In the context of crisis communication, translations that are logical, free of obvious errors and written in a clear and direct style have been found to be more likely to inspire trust.[3] In a similar context involving a natural disaster, translations have been identified as an influencing factor of both trust and distrust.[4] Those affected by the disaster may rely on translations to make important decisions about their movements or actions.[5] Conversely, some individuals may be less trusting and avoid placing themselves in situations where they would need to rely on translations for their safety.[6]

In a broader sense beyond the world of translation, trust can be described in relation to at least four levels: individual, interpersonal, relational and societal.[7] Individual trust is about a personality trait – an individual's propensity to be trusting. Interpersonal trust is about the process whereby a trustor decides to trust a trustee. Relational trust concerns the way in which trust can emerge from the trustor–trustee relationship itself – how the relationship dynamic may be conducive to trust or distrust. Societal trust is about collective attitudes – a society's trust in politics or science, for instance. All these levels of trust are relevant to the discussion in this book. Individuals who tend to be more

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

trusting may be more likely to fall victim to machine translation's pitfalls. Those who are more distrusting may be more likely to miss out on its benefits. Trust is also central to relationships between the professionals investigated here and the public they serve, and machine translation is a possible feature of this relationship. Lastly, societal trust influences the general value attributed to AI and to other technologies, which can in turn filter down to individual decisions.

The interpersonal level of trust is the one that directly concerns the discussion in this chapter. Interpersonal suggests that trust is placed by one person in another. Indeed, philosophers may argue that authentic trust only exists between wilful agents,[8] so trust in technology on this basis would be metaphorical.[9] The value of treating technologies as potential trustees is nevertheless widely recognised and there is a growing body of literature on trust in AI.[10]

Some of the trust-inducing characteristics of a trustee are particularly relevant to AI systems, such as the system's explainability and transparency. The obscure nature of how AI systems arrive at their results, and the fact that this process is difficult to explain, can lead to distrust of these tools.[11] Excessive information about a system's inner workings can also be a problem. It may be confusing and thereby erode trust.[12] Since my focus in this book is on communication practices rather than on specific incarnations of AI, in this chapter I draw on a broad model of trust that is well suited to technology while also being directly relevant to information consumption and dissemination, namely the model of trust in digital information proposed by Kari Kelton and colleagues.[13] This model, which I call 'the Kelton model', suggests that trust is needed in some specific conditions, including when the information is potentially consequential and when there are no standards to ensure its quality.[14] Standards, or a lack of them, are particularly relevant to discussions of trust in machine translation.

Uncertainty and a Lack of Standards

A standard can be understood as measures that allow products and processes to be assessed based on consistent criteria, such as efficacy or safety. The keyboard I am using to type these words, the light bulb in my desk lamp and the glass in the window in front of me all have met specific standards before being available for purchase. There are methods that can be used to consistently evaluate machine translations, but unlike keyboards, light bulbs or glass panels, machine translation systems have no accepted standard of accuracy, and thereby of safety, that needs to be met before they are released to the market.

A widely known machine translation evaluation method is the BiLingual Evaluation Understudy (BLEU) score.[15] This score automatically compares the machine translation output with pre-existing translations. It shows how

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

close the machine output is to high-quality translations that can be regarded as a model – that is, a reference. The use of BLEU, or automatic evaluation more generally, has several shortcomings.[16] Nevertheless, automatic evaluation is a useful diagnostic method, especially in the process of developing machine translation systems. Low BLEU scores can indicate issues with the data used to develop the system or with its development steps. Higher BLEU scores following attempts to improve the system can indicate that the adjustments have probably worked.

Human evaluation tasks – when humans manually assess translations – can also draw on established methodological frameworks. Multidimensional Quality Metrics is a widely used such framework.[17] It outlines a set of evaluation criteria – for example, how accurate or stylistically acceptable a text is – that humans can use to assess machine as well as human translations. The criteria can be customised since not all assessment components are relevant to all evaluations.[18] But the idea is that by drawing on a common framework, it is possible to rely on agreed benchmarks to decide whether a translation is adequate or fit for purpose.

Machine translation providers may use any of the above methods, or a combination of them, to assess the quality of their machine translations. But none of these methods can be considered a standard for declaring a machine translation system safe. As mentioned at different points in the book, even if machine translation tools are tested, it is always possible that users will need to translate something new, unusual or which might not conform to the examples used to develop or test the system. Although it is possible to extrapolate from a test to gauge how well a system is likely to perform 'in the wild', showing that a system has had positive test results in no way removes the need for carefully considered judgements about whether the system can be trusted.

In addition to the use of consistent benchmarks or criteria, standards are also upheld when a relevant body or organisation enforces reliable practices. For example, TV broadcasters in the UK are subject to standards set by Ofcom, a national regulator of communication services.[19] The Ofcom broadcasting code seeks to ensure that the information disseminated on news programmes is accurate and that any errors are rectified publicly and in a timely manner.[20] While this type of regulation has limitations, Ofcom has enforcement powers at its disposal.[21] Its regulatory work is likely to increase public trust in the accuracy of broadcast news. Similar, albeit less enforceable, guidelines are emerging for uses of machine translation and AI. But for several reasons these guidelines are limited in their capacity to reassure machine translation users that translations can be trusted. Organisational policy may specify, for instance, that any machine translations published on the organisation's website need to be checked and approved by a qualified linguist. Such a measure, if always

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

followed, would go some way towards mitigating the risk of mistranslations. But not even this measure would always guarantee high levels of accuracy. There would still be a need for trust in the translations and in the individual checking them.[22]

Processes of Trust Development

The Kelton model of trust identifies processes through which trust is developed as well as characteristics of the information that lead to it being perceived as trustworthy. Trustworthiness, according to this model, refers to the objectivity, validity and stability of the information as well as to its accuracy, coverage, believability and currency (i.e., how up to date it is).[23] As for how trust is developed, individuals may rely on experience and recommendations as well as comparing information sources or considering their reputation.[24] Having a positive experience with the information obtained on a website may lead the user to trust this website in the future. The reputation of the website publisher, and of the author or producer of the information, can also influence a potential trustor, as can comparing different sources and checking whether they corroborate each other.[25]

For individuals who possess the meta-literacy competencies mentioned in Chapter 1, judging information trustworthiness will often be intuitive. Cautious information consumers may want to check multiple sources and consider the sources' reputation before deciding whether to trust the information. They might also consider characteristics of the information itself, such as how plausible it is, whether it seems to be up to date and well documented. For several reasons, these principles and techniques are limited when the potential trustee is machine translation.

Take the principle of experience. In the summer of 2023, the Department of City Planning in New York City suggested on its website that the Communist Party of China could influence planning decisions in New York.[26] The Chinese version of the website included passages such as 'the Chinese Communist Party and the public can make wise decisions on every project that passes public review'.[27] The suggestion that a political party from a different country was involved in city planning processes was embarrassing for the department. The reference to the Communist Party was a Google mistranslation of the abbreviation used for the City Planning Commission, CPC. The error took the Department of City Planning by surprise: 'This is the first time anybody who I work with has encountered an issue like this', one of the department's representatives said.[28] Staff at the department had therefore not identified or been made aware of mistranslations of this nature until that point. Their personal experience can be described as an unreliable precursor of trust – or possibly of uncritical trust, since it is not clear whether the department had

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

considered the risks. Whichever the case, the perception of a positive experience can be deceiving. Those who cannot themselves identify machine translation errors may conclude that 'no news is good news'.

Regarding sources and their reputation, some of the most popular machine translation developers in the Western world are tech giants such as Microsoft, Google, Meta, Apple, Amazon and OpenAI. Most of these companies arguably do not have a good reputation in terms of their ethical and social responsibility records.[29] But their reputation in relation to machine translation is a more nuanced matter. On the one hand, 'Google Translate' or 'Google translations' have been used synonymously with 'translations of poor quality'.[30] On the other, these companies are widely trusted, even if unwarrantedly, as custodians of multilingual information. Machine translations provided by these companies are visibly present on the websites of governments, hospitals and other official institutions. As mentioned in Chapter 4, even courts of law publicly declare using them. If respected institutions' use and acceptance of machine translation were to be taken as a measure of the technology's – or its provider's – reputation, users would be right to think that Google Translate is a reputable and thereby probably trustworthy system.

The question of reputable sources becomes even more complex when machine translation deployment is covert or at least not immediately apparent. The machine translation tool CBP Translate, used by Customs and Border Protection officers at the US border, is powered by Google.[31] Yet the website of the Department of Homeland Security includes statements such as 'CBP developed CBP Translate, a mobile and web application that will assist CBPOs [CBP officers] in communicating with travelers.'[32] The developer of the tool is described here as CBP itself. This statement is technically correct: CBP developed, or procured the development of, the application and its user interface. Members of the public would however be forgiven to presume that the translations are generated by the US government when in fact they are generated by Google. Many machine translation applications are powered by external cloud-based providers.[33] The ambiguous status of the provider in these cases can confuse trust judgements.

Ambiguous providers can also undermine attempts to compare sources. Comparing translations provided by different systems is one technique that can be employed in the process of developing machine translation literacy. The idea, especially for those who are unfamiliar with machine translation or with how it works, would be to note how systems trained on different data produce different results.[34] It would also stand to reason that if different systems generate the same translations, the translations may in principle be more likely to be accurate. This principle is limited. Different systems can make similar mistakes, whether by chance or because the systems have similar training

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

datasets. Additionally, checking the outputs of different systems may be unadvisable for reasons of privacy and confidentiality. If a civil servant has access to an officially procured machine translation system, they should ideally avoid consulting additional publicly available systems which may not offer the required level of data protection, an issue I return to later in this chapter. It is nevertheless logical to presume that the larger the number of systems that generate the same translation, the more trustworthy the translation is likely to be. Service users may decide to cautiously rely on this principle to check information they come across on the websites of public institutions. But this strategy is only sensible if the translations are indeed independently generated. Wide, potentially covert, adoption of a small number of hugely popular providers compromises the already limited value of contrasting machine translation sources. Translations may appear to corroborate each other when in fact they are all produced by the same system or by similar versions of it.

Certainty, therefore, is not usually on offer. A lot about trust in machine translation comes down to the circumstances of each use context. If machine translation is needed in the first place, and if users cannot overcome a language barrier without it, then using it will to a greater or lesser extent be a leap of faith. The greater the level of users' awareness of this uncertainty, the less susceptible they will be to machine translation's risks. Perceptions of the technology's trustworthiness, including those held by non-users, can open pathways to risk mitigations and to machine translation literacy. The accounts discussed below are hoped to point to some of these pathways. They reveal the nature of expectations, of misconceptions and of calls for support and reassurance in relation to machine translation use.

Uncritical Trust

The participants who declared no prior use of machine translation were asked the following question: *Would you consider using automatic translators in your work? Please provide brief details of what could influence your decision*.[35] While I do not aim to provide a quantitative analysis of this question, it is worth noting that definitive negative answers – for example, 'No, I wouldn't trust it' – were rare. It was more common for participants to say that they would use machine translation if needed. Many seemed prepared to use it in any scenario.

Non-users of machine translation were not invited to describe their specific professional role, so when reproducing their responses in this chapter I only mention their sector. Some of them showed unqualified willingness to use machine translation even for purposes where information accuracy is important, such as taking medical histories or coding information into someone's medical record: 'Definitely yes [I would consider using machine translation] – . . . when asking patients about their medical history if English isn't their first

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

language' (Medical/Healthcare); 'Yes, for non-English people who have paperwork that isn't in English that needs adding to their UK medical records' (Medical/Healthcare); 'Yes to translate accident report forms from foreign drivers' (Legal Services).

Distrust was not, therefore, the reason why these participants had not used machine translation. When asked the question, most of them simply thought of contexts that would prompt them to use it. Although these participants may come to consider the risks and benefits of machine translation if faced with the need to use it, the fact that it did not occur to them to mention caveats is itself informative.

In any professional sector, some individuals will be more open to experimentation than others. An important precursor of uncritical trust is the individual's attitudes to technology. Attitudes have been identified as an important factor capable of influencing human–technology relations. Based on data from the first decade of the 21st century, a study about the internet in Great Britain found that positive attitudes to technology were stronger than age as a predictor of trust in the internet.[36] It is well known that age correlates with digital exclusion,[37] so the stronger predicting power of attitudes is noteworthy.

Individuals who have more positive attitudes to new communication methods may be more likely to trust machine translation or to accept its use even when the experience may at first seem unnatural. In response to the question about whether they would use machine translation and what would influence their decision, one of the participants responded, 'Yes definitely as it may make things easier. I'm always up [for] trying to learn more' (Healthcare and Social Assistance). Another participant said, 'Yes, [I would consider using machine translation] as it would be helpful, but I would have to get used to it, so you feel more confident, like most things, getting used to something new needs time to become more familiar and happy with it' (Healthcare and Social Assistance). This type of openness to learning new things is essential in the world of work as practices evolve and new technologies are introduced. But positive attitudes can also stray into uncritical acceptance. Individuals may feel under pressure to adapt to new norms and forgo important instincts in the process.

As mentioned, social and individual levels of trust are interrelated. Artificial intelligence technologies are increasingly present in the social and political discourse. As a hugely popular application of AI, machine translation is a feature of this discourse too. In a previous study, I looked at the coverage of machine translation in written English-language news.[38] I found that news outlets tended to oversell the capabilities of machine translation tools.[39] Attention to this type of coverage, and to the media landscape more broadly, is likely to influence individuals' attitudes to machine translation. A recent study has empirically demonstrated this opinion-shaping power of the media for generative AI. The study showed that attention to the media influences

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

perceptions of the social norm – that is, the extent to which one believes that friends, family members and others in society are likely to approve of AI uses.[40] This type of perception was a strong predictor of individuals' intentions to use AI.[41] In other words, trust is contagious. If we think that those around us will approve of our actions, this shapes and validates our own behaviour.

Uncritical trust is not, therefore, coincidental. Any misguided decisions to use machine translation will be the product of several complex factors, some of which have been discussed in previous chapters, such as convenience and institutional budgets. Those who trust machine translation without considering the risks are not necessarily or intentionally irresponsible. Most of my questionnaire respondents who used machine translation for the potentially dangerous purposes discussed in Chapter 4 were doing their best to care for the communities they served. Those who had not used it, too, often thought of service users when explaining what they would consider in their decision-making. I discuss more of their answers below as I look at some of the specific mechanisms leading to trust and distrust of machine translation tools.

Trust and Distrust in Machine Translation

Unlike those who displayed a propensity to trust machine translation uncritically, the participants whose views I discuss below acknowledged, even if implicitly, that using machine translation involved risks. Some of them focused on untrustworthiness: why they would not use machine translation. Others explained that they would trust machine translation in some cases and not others. I consider decision-making factors raised by both these types of participants. I start with perceptions of accuracy, which can be a factor in both trust and distrust. I then look at perceptions of the social norm, at the questions of privacy and confidentiality, and lastly at human empathy, which tended to be associated with distrust of machine translation systems.

Seeking Accuracy Assurances

Some professionals thought that machine translation was not accurate enough to be used in their work and were categorical in this respect: 'No, wouldn't trust it to give patients an accurate translation' (Medical/Healthcare); 'No, I would not trust them to be correct and would like to translate via a person' (Healthcare and Social Assistance); 'I would not [consider using machine translation]; my work involves assessing evidence to establish any criminal liability. To use a system which is not foolproof could result in improper interpretation of evidence and authorisation of actions which can lead to the loss of a suspect's liberty or of the victim's chance to get justice' (Legal Services).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

More commonly, participants would use machine translation if they could be certain of its accuracy: 'If I could trust the grammar then I can see a use for it in translating letters and reports into a patient's language, and allowing information leaflets to be translated into any necessary language easily and cheaply' (Healthcare and Assistance); 'Only if it had been fully tested for accuracy and reproducibility' (Police); 'Yes I would [consider using machine translation], but would need to see proof that it is correct first' (Medical/Healthcare); 'I would [consider using machine translation] but only if accuracy could be guaranteed – the risk in healthcare is high if there's a mistranslation' (Medical/Healthcare).

Requests for certainty in relation to machine translation are not new. Previous research has shown how users express desires for machine translation to be 'perfect' and '100% accurate'.[42] For many of the reasons discussed earlier concerning a lack of standards, and because of the dynamic nature of language, the prospect of guaranteeing accuracy is largely unrealistic. Views that underestimate translations' inherent malleability and subjectivity are nevertheless common, especially among those who do not have a direct experience of translating or interpreting for others.

When I teach translation to undergraduate students, a significant proportion of class time is devoted to discussions about how translations are not necessarily intrinsically correct or incorrect. Students can find this unsettling. Some of them will have been asked to translate texts while learning to speak or gain proficiency in a non-native language. In doing so, they will have often been marked down for phrases that do not convey the meaning of the original text or which might not conform to grammar conventions. So they come to class expecting clear boundaries between right and wrong. This expectation is, of course, valid. One of the premises of this book, after all, is that using machine translation involves risks, and chief among these risks is the possibility of misrepresenting a message and causing some type of harm in the process. But even in the public service settings discussed here, accuracy is context-dependent.

In her analysis of interpreting practices at a hospital in California, Claudia Angelelli discusses several examples where interpreters complement what is said by healthcare professionals. In an interpreter-mediated interaction between a nurse and a patient, the nurse asks the interpreter, 'Can you ask her [the patient] about chronic illnesses, diabetes . . . and all that?' Instead of using an equivalent term for 'chronic illnesses', the interpreter decided to break this phrase down for the patient:

INTERPRETER: Mrs Mesa, has a doctor ever told you even twenty years ago here or there that you had diabetes?
PATIENT: No.
INTERPRETER: That you had high blood pressure?
PATIENT: No.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

INTERPRETER: That you had heart disease?
PATIENT: Noooo.
INTERPRETER: That you suffered from liver problems? Kidney problems? Stomach problems?
PATIENT: No.
INTERPRETER: Have you ever been operated on or hospitalized? Here or there?
PATIENT: Noooo.
INTERPRETER: You never been sick?
PATIENT: Well . . . I was sick but . . . it was de . . . nervous depression . . . I did not suffer from anything else.[43]

The interpreter could have used an equivalent phrase for 'chronic illnesses' in Spanish (e.g., *enfermedades crónicas*). But they chose to elaborate on the question. The nurse may have prompted this intervention, which the use of 'all that' would imply. In any case, those with expectations of complete accuracy may favour an interpreting strategy that closely corresponded to what the nurse had said. But the interpreter's objective here was to obtain the required information, and for that a more tailored strategy may be necessary. Patients may be reluctant to share certain personal details, especially concerning their mental health. Due to cultural factors or personality, both the patient and the healthcare provider may have different communication styles – for example, being more guarded or direct. Without the type of mediation illustrated above, important information can be missed or withheld. Therefore, the measure of accuracy in an interaction depends on the communicative objective, and this is true for both human- and machine-mediated communication.

It would in principle be possible to ask a large language model to offer the type of mediation illustrated above. It is not far-fetched to imagine that, as machine learning improves, it may start taking cultural factors into account and attempting to respond accordingly by rephrasing questions and perhaps even by using face recognition to read emotional states.[44] Even if we set aside the ethical implications of this possibility,[45] accuracy would still be a moving target. Adjusting the original message is necessary in some contexts. Human mediators have the advantage of being able to empathise in these cases, as I will discuss later in the chapter. Humans understand that discussing one's mental health can be uncomfortable because they are equipped to know what this feels like. But neither AI models nor humans can be *guaranteed* to get this type of mediation right every time. The participants who asked for proof of accuracy beforehand in order to trust machine translation cannot, therefore, have their request fully and rigorously fulfilled – not if accuracy is understood as successfully achieving a communicative goal.

The marketing strategies adopted by some technology developers can make the notion of guaranteed accuracy even more problematic. Those who have something to gain from the use of machine translation can rely on questionable

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

methods to convince others that their systems can be trusted. In September 2016, researchers at Google published a blog post announcing a new version of Google Translate. At the time, Google was releasing translation models that had been developed using artificial neural networks. As mentioned in the Introduction, this methodology was superior to the one that preceded it. Google researchers sought to demonstrate their success by asking human evaluators to compare translations produced by both the old and the new types of model. Human translations were also assessed. The assessors used a scale that ranged from zero to six, where zero, the researchers explained, corresponded to 'completely nonsense translation' and six corresponded to 'perfect translation'.[46] The results were displayed in a chart with a dotted line at the top indicating the 'perfect translation' level. All translations, including the human ones, failed to reach the 'perfect' level, but the chart showed that the new type of model was superior to the previous one. It also showed that, for some language pairs, human and machine translations had very similar levels of quality.

These results were criticised at the time.[47] Google's graph was arguably misleading. It suggested that language translation was a finite problem – that there is a finish line which, once crossed, will allow translations to be declared 'perfect' with little regard for different purposes or contexts of use.

Google is not the only machine translation developer to frame translation in this way. In December 2024, the co-founder of Unbabel, an AI-powered language company, declared that 'humans are done in translation'.[48] In a contribution to an article in the *Economist*, Unbabel's co-founder predicted that within three years human translation would be almost entirely unnecessary in the language industry.[49] Isaac Caswell, a researcher at Google, reports in the same article that, for some languages, 'the problem of translating one sentence to another is "pretty close to solved"'.[50] These views are reflected in the article's title, which says, 'Machine translation is almost a solved problem: But interpreting meanings, rather than just words and sentences, will be a daunting task.'[51] The real problem, of course, is that, without meanings, words and sentences have no communicative value – for practical purposes, it could be argued that they are not words and sentences at all. The gist of what AI developers perhaps mean to say is in any case not unreasonable. For languages that are well represented in AI training datasets, machine translation systems are very good at literal-level language conversions. They are increasingly good at interpreting what words and idioms mean in context too. But this does not remove risk nor the need for considered trust judgements.

Google Cloud asks its customers to display a message to potential users of its translations. The message says: 'Google disclaims all warranties related to the translations, express or implied, including any warranties of accuracy'.[52] Once again the legal teams at technology companies seem more cautious than technical or marketing departments. Translation is not a problem that can be

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

declared objectively solved. Propagating such a view has consequences. It taps into, and potentially amplifies, expectations that machine translation systems can be guaranteed to be accurate. As Google's lawyers would seem to agree, they cannot.

Perceptions of the Norm

Many of the communicative tasks examined in this project take place in collective working environments. The behaviour of others in these environments can influence trust judgements, especially when work colleagues are perceived to set the pattern for how language barriers are handled: 'Yes [I would consider using machine translation], recently we had a Polish gentleman with dementia and other staff used Google Translate' (Healthcare and Social Assistance); 'Colleagues have used it before when needed, so there is no reason why I wouldn't use it' (Healthcare and Social Assistance).

The experience of others can also help to shape strategies for using machine translation more effectively: 'Yes, occasionally this might be helpful [if] working with someone who can't speak English well, although it's important to use them carefully because they aren't always accurate – my colleague used one and we noticed that if her pronunciation wasn't very clear and the sentences weren't short, the translation would be nonsensical' (Healthcare and Social Assistance).

Perceptions of what is normal or accepted are also clear in responses that refer to rules as a guiding principle. Some participants indicated that they would not use machine translation because they were not allowed to: 'No, they are not recommended under NHS policy. An approved translation service should be used' (Medical/Healthcare); 'We have translator services within our service which we have to use and are not able to use automatic translators. So sadly, no' (Medical/Healthcare).

Many indicated that they would use machine translation if this was endorsed by their superiors: 'Yes [I would consider using machine translation systems] but they would need to be approved by the police or I could be risking getting in trouble for using unauthorised systems' (Police); 'Yes if they were supported by the hospital, i.e. validated and approved' (Medical/Healthcare); 'It would be up to my manager' (Medical/Healthcare).

Participants who refer to what they are allowed to do take a deontological approach to decision-making (see Chapter 2). Behaviours that are officially or formally considered 'right' can outweigh individual preferences, as suggested by the participant quoted above who, in their words, would 'sadly' not use machine translation because professional language services had to be used instead. The lack of a specific position taken by employers can be interpreted in opposing ways. If machine translation is not explicitly mentioned, more

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

cautious members of staff, or those who fear 'getting in trouble', will probably conclude that by default machine translation use is prohibited. Conversely, if machine translation is not explicitly banned, some staff may interpret this as green light to use it. In either case, institutional silence gives more weight to the behaviours of others as indicators of the norm.

Of the project's 828 participants who *had* used machine translation at work, fifteen per cent declared that machine translation was one of the communication methods recommended by their employer.[53] Some descriptions of what was recommended indicated that the recommendations were not necessarily formal directives but rather suggestions made by colleagues, or actions that the team must resort to in times of need. A healthcare facilities manager described the recommended procedure as: 'Use the telephone service to access a human translator but this has a cost charged to the user [i.e., the healthcare provider] as it's an external company. That's why the majority of us use GT [Google Translate].' Decisions here are grounded in what can be considered a majority practice – what 'the majority of us' do.

The challenges of working on the front line can also empower professionals to take matters into their own hands even when they know that they are not following procedure. A healthcare support worker who had used machine translation explains: 'There is no official guidance on using automatic translators. This was done without the approval of senior management. The official procedure involves organising an in-person face-to-face translator to attend the ward.'

When advice is in place but is ambiguous or impractical, it can have the same assenting effect as no advice at all. One of the professionals quoted above mentions the fact that the UK's National Health Service does not recommend the use of machine translation. Yet many of the service's own websites are machine-translated, as discussed in Chapter 1. The disconnect between ideal and feasible practice helps to cement coping strategies as the de facto benchmark of what can be considered appropriate. Not all informal practices are problematic, but the professionals who mention examples set by colleagues, or those who defer to their managers, are pointing to the value of meaningful workplace guidance. While individual judgement is essential, failing to acknowledge machine translation at an institutional level creates the space for many of the risky uses of it previously examined.

Privacy and Confidentiality

An important caveat raised by non-users of machine translation was its potential to breach privacy and confidentiality. One participant asked, 'is it [machine translation] being used to monitor and spy on the police?' Those who mentioned concerns of this nature were most often prepared to use machine

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

translation if the technology could be deemed cyber-secure: 'Potentially [I would consider using machine translation] if there is a language barrier between me and a client or witness, but they [the machine translation systems] would have to be secure, for us to comply with data protection regulations, and approved by our compliance team' (Legal Services); 'Yes, provided that the patient data was stored safely and there were no privacy or confidentiality concerns. I come across many patients where there is a language barrier affecting their care and something like this could be greatly beneficial for both the patients and employees' (Medical/Healthcare); 'It would depend on the ethics and risk to confidential information, as client confidentiality is paramount. If these were met, and the client was ok with it, then maybe I'd use it' (Healthcare and Social Assistance).

Threats to privacy and confidentiality can materialise in different ways. Where machine translations are provided online, the provider will usually need access to the content in order to translate it. If a paid or bespoke service is used, the provider is likely to make assurances that the content is not used for purposes other than the provision of the service. When the tool is offered free of charge, assurances of this nature are not usually made.[54] In these cases, users may be giving the machine translation provider licence to use the content for the provider's own benefit.[55]

In addition to risks associated with the provider, any online application will be subject to the risk of cyber-attacks, and several recent incidents would suggest these attacks are on the rise.[56] Moreover, leaks can happen through inconspicuous means, including if the user unwittingly agrees to data transfers that put their information at risk.

In 2016, a German investigative journalist embarked on a project in collaboration with a data scientist from the German cybersecurity organisation Deutsche Cyber-Sicherheitsorganisation. As an experiment, the pair pretended to be a company looking to purchase web browsing data from Germany. They were told that data on German citizens was difficult to access, but they eventually succeeded. They managed to obtain for free a dataset with supposedly anonymised URLs. The data was available because the website visitors had installed a browser plugin that retained their browsing history with the intent of sharing it. The users of the plugin were probably unaware of how the plugin worked even though, to install it, they would need to have consented to its privacy policy. The journalist and her research partner attempted to identify those who had visited the URLs to demonstrate that the data was not in fact anonymous. They identified several high-profile individuals relatively easily using well-known data de-anonymisation methods. One of the individuals was a politician who had been filing her tax return and researching medication online. Another was a German police officer who had been using the web interface of Google Translate to translate details linked to a fraud

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

investigation. Any content inserted into the text box of Google Translate and translated by the system becomes part of the URL for the page. The Google translations were therefore included in the dataset. They contained details such as a specific IP address the officer was looking into. They also included the officer's email address, phone number, full name and his specific county and police division.[57]

The participants in my project who expressed concerns about the privacy implications of using machine translation were therefore right to be wary. While questions of privacy and information security are now in the public eye in relation to uses of large language models, relatively little research addresses the confidentiality and cybersecurity risks of using machine translation specifically.[58] Of the 828 participants who had used machine translation at work, 76.7 per cent selected an option declaring that they had used an openly available tool via a browser and 57.7 per cent reported using the technology on personal devices.[59] These are the perfect conditions for the type of data leak exposed in the German project.

These conditions are likely to constitute what is usually termed 'shadow IT',[60] when, as part of their work, members of an organisation access tools or online services that have not been vetted or procured by their IT department. Machine translation is a prime candidate to be used as shadow IT, especially in emergencies. The risks of shadow IT are a significant blind spot among the professional populations examined in this project. The privacy risks of using machine translation may in fact be a blind spot for society more generally. As mentioned in Chapter 1, some users feel that machine translation tools offer more privacy than professional human translators.[61] In a study about machine translation literacy instruction in Canada, students were surprised by the privacy implications of using tools like Google Translate and found knowledge of these implications to be among the most informative aspects of the course.[62]

In the settings examined in this book, machine translation users will often have to weigh up competing priorities. In an emergency, service providers may conclude that waiting could be more harmful than potential breaches of privacy. This possibility notwithstanding, the privacy and confidentiality risks of using any technology should arguably be central to trust judgements about the technology's use. Among non-users, privacy risks were a recurrent reason for distrusting machine translation. But these risks tended to go unnoticed by many of those who had used it.

The Human Touch

Be it when a doctor provides care to a patient, when a police officer questions a civilian or when a social worker speaks to a family, emotions are important in many service encounters. As discussed in Chapter 4, some professionals feel

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

that they can offer a more person-centred service in these encounters by using machine translation. Conversely, some of those who had not used machine translation mentioned its lack of empathy or emotions as a reason for avoiding it: 'No [I would not consider using machine translation], the type of communication delivered needs a human touch for emotions' (Medical/Healthcare); 'No because the human touch is essential for my work' (Healthcare and Social Assistance); 'No, because I am a psychotherapist. Communication is verbal and routinely much more likely to be viewed in terms of process than content, i.e. tone, form, delivery, timing, mood, etc.' (Healthcare and Social Assistance); 'Automatic translators could be handy but I think they lack the compassionate/empathetic factor' (Medical/Healthcare).

A lot about our experience of empathy concerns our preconceptions about who or what we are interacting with rather than the content of the interaction. In a scientific experiment, research participants were asked to read a letter. The participants were told that an immigrant worker in Israel had written the letter to complain to the authorities after being misidentified and assaulted by a group of security guards. In this imagined scenario, the worker was from the Philippines and had a low level of proficiency in Hebrew. The worker therefore had written the letter in English and had it translated to Hebrew. All participants in the experiment saw the same translation, but some of them were told it was a professional translation prepared by a human while others were told it was a machine translation. Those who were told it was a human translation thought the letter was more capable of conveying emotional otherness.[63] Machine translation was instinctively perceived to be less suitable to convey the experience of the worker.

In a different experiment, researchers asked licensed healthcare professionals to evaluate two types of written medical advice: messages that had been generated by AI and messages written by real doctors. The evaluators were not told which version of the advice they were reading. They found the AI-generated messages to be of higher quality. They also found the AI messages to be significantly more empathic than the human ones.[64] These results may seem surprising, but they corroborate a well-known phenomenon recently labelled the 'artificial-empathy paradox'.[65] Specifically, humans may find AI-generated texts convincing even in terms of the texts' capacity to convey empathy, but just the knowledge that the texts are artificially generated can cause a negative response.

Like in the experiment involving a single translation that had been attributed to different sources – human and machine – it is the idea of a machine-generated message that can have negative emotional connotations and not necessarily the message itself. The public portrayal of certain uses of AI is a good illustration of this emotional bias. In 2024, a UK tabloid published a story about doctors who had used large language models to reply to patient

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

complaints.[66] In a clearly emotional framing of the story, the letters were called a type of 'false apology'.[67] The Medical Defence Union, a UK organisation that indemnifies doctors against liabilities,[68] said that patients may find these uses of AI 'upsetting'.[69]

Irrespective of how accurate or secure AI models are, just the idea of using one can be perceived to erode trust between service users and providers. Openly declaring the use of AI may attenuate this negative perception or at least pre-empt accusations of deception.[70] But as shown in the Hebrew translation experiment, a message that is openly attributed to AI can also be subject to criticism.

Part of the reason for our bias against the capacity of machines to come across as empathic is the fact that not all aspects of empathy can be artificially reproduced. Researchers have identified three facets of empathy: cognitive, emotional and motivational.[71] Cognitive empathy concerns understanding or effectively recognising the emotions expressed by others. Emotional empathy is about sharing emotional states – experiencing what others feel by relating to what they are going through. Motivational empathy is about choosing to be empathic and helpful to others. Being empathic takes effort. It requires the use of our time and emotional resources, so those receiving empathy are likely to appreciate the attention invested in them.[72] While AI models may be able to accurately identify and indeed translate emotive language, they cannot feel what others are feeling or rely on this shared experience to make a conscious decision to help others.[73] The empathy they may be able to express is artificially motivated, hence it being described as 'false'. Humans can fake empathy too, but we cannot be entirely sure of whether someone is faking an empathic response. Artificial intelligence models, on the other hand, are known not to have any real feelings,[74] so in this sense they are always 'faking it'.

When machine translation is mentioned as both an enabler and an inhibitor of person-centred care, there are thus different mechanisms leading to these perceptions. The professionals quoted in Chapter 4 who believed machine translation allowed them to provide a more personal service were referring to their own capacity to be personal. Access to language professionals is limited and may involve delays. Some service providers therefore see machine translation as an opportunity to take ownership of multilingual communication and address service users' needs more promptly and directly. By contrast, the non-users who mentioned emotional awareness as a reason to distrust machine translation were referring to the inherent limitations of artificial empathy. To paraphrase one of the study's healthcare participants, machine translation cannot provide the emotional support that a human translator would. Artificial intelligence will possibly never be able to replicate humans' capacity to truly and consciously empathise, so the service offered by human linguists, albeit sometimes unavailable, is in this respect unique.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

Development of AI Literacy

I described the notions of literacy associated with AI and machine translation in Chapter 1. They are both related to a broader conception of information literacy, which has among its manifestations many of the processes of developing trust in digital information. Individuals can work towards developing competencies pertaining to all these literacies, and dedicated educational programmes exist for this purpose. Some of these programmes address the population in general and are openly available online, such as the Elements of AI programme, developed at the University of Helsinki, which includes courses on what AI tools do as well as on how they are developed.[75] Similarly, the Ethics of Technology course designed and delivered by the Massachusetts Institute of Technology covers topics such as privacy and consent.[76]

Reviews of AI literacy instruction in formal education have looked at how this subject is taught at colleges and universities,[77] as well as in primary and secondary schools.[78] Competencies that are specific to machine translation are covered in formal education too, especially in universities. The translation scholar Lynne Bowker, who first operationalised the concept of machine translation literacy, proposed an undergraduate machine translation course that covers topics including privacy, transparency and risk.[79] Students work through case studies that help them to discern the most suitable way of approaching different language tasks by considering the task's objectives and the consequences of potential mistranslations.[80]

A necessary component of machine translation literacy is arguably *translation* literacy. Individuals who are not familiar with the process of expressing concepts, ideas or feelings across languages are not necessarily attuned to the context-dependent nature of language and communication. These individuals may therefore have unrealistic expectations of language services. The ideas previously described of an objectively perfect translation, and hopes that machine translation systems can be guaranteed to be 100 per cent accurate, are potential signs of this lack of first-hand familiarity with translation and interpreting as communicative practices. The fact that prior critical engagement with these practices cannot be taken for granted justifies efforts to introduce translation and interpreting to non-linguists. Bristol City Council, in the UK, offers good-practice guidelines on how to speak to community members through an interpreter. The guidelines provide basic but sound advice such as 'Choose a quiet space if you can' and 'Explain what your role is [because] there are often no direct equivalents of some of our services in other countries.'[81]

A recent book goes into significant detail of this nature, mainly in relation to written translation. In *De-mystifying Translation: Introducing Translation to Non-Translators*, Lynne Bowker covers several aspects of language and

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

translation use that non-linguists may not necessarily realise. The book has a machine translation chapter which examines many features of the technology that merit attention, including its data-driven nature and how the textual data used to develop translation models have a direct effect on the models' results. Some models may generate sexist language by replicating what can be found in the data, for instance, whereas any model will be likely to underperform when translating languages that are not sufficiently represented in the data used to develop the model[82] (see Chapter 6).

Activities aimed at raising this type of awareness have a role to play in the professional sectors examined here. The qualifications required for practising professionally in the sectors covered by this project, especially health and social care, are likely to address questions of multilingual communication and cross-cultural practice.[83] As seen in Chapter 3, customer service qualifications also address some of these questions.[84] Machine translation may be a feature of this training, but many currently practising professionals will not have been exposed to formal education courses on AI or machine translation. A significant number of them will have left education years or decades ago when machine translation was not widely available. Any upskilling in these cases is therefore likely to be done on the job. Like other aspects of professional practice that may pertain specifically to an organisation – for example, the institution's policy on cybersecurity – approaches to the use of machine translation may in principle feature in workplace training and professional development.

The 828 project participants who had used machine translation at work were asked whether machine translation tools had been mentioned in any workplace training they had received. Ninety-three among them (11.2 per cent) said 'yes'.[85] This is a small number considering how consequential miscommunication can be in the service settings discussed in this book. Even among these ninety-three professionals, the training or advice received was not necessarily effective or appropriate. A social worker reported that, in their training, they were told to 'use them [machine translation systems] as needed'. A pharmacy dispenser was 'recommended to use Google Translate on provided iPad if unable to communicate'. This pharmacy dispenser had used machine translation when 'handing out or discussing medication/health conditions'. Responses like these suggest that not all training on machine translation that is currently provided equips professionals to critically consider the technologies' risks and benefits.

Those who had not received workplace training on machine translation were invited to describe what type of training they would have found helpful. Some of them took the opportunity to say that the use of machine translation should not be allowed in the first place even though they had used it themselves. Others did not have specific requests and would have been grateful for any training at all: 'Any training would be beneficial as we haven't had any' (Well-Being Advisor, Healthcare and Social Assistance). More commonly, participants

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

provided one of two types of responses: they described the types of training they would like to receive or instead said that training was not required.

Those who did not think training was required tended to focus on the procedural aspects of interacting with and making use of machine translation tools: 'They [machine translation systems] are generally easy to use. No need for training' (Physiotherapist); 'It's pretty self-explanatory so I don't think training is required' (Healthcare Assistant); 'None for me personally. People who are less able in regards to technology may need training in where to press, language selection box' (Support Worker, Healthcare and Social Assistance).

It is understandable that some of the participants would instinctively think of the technical procedures of operating a machine translation system. Many of them need to operate specialist technology in their work, which usually requires training. Machine translation tools, on the other hand, are intuitive. Although basic technical training may be useful for some members of staff, using machine translation is indeed unlikely to require specialist technical knowledge. But what is left unsaid by these comments is also noteworthy. Just the fact that it did not occur to these participants to think of potential risks and how to avoid them suggests that the ethical implications of machine translation were not at the front of their minds. This potential lack of attention to ethical questions speaks to some of the systemic issues around machine translation deployment which other participants decided to highlight.

Those who called for specific training had two prototypical requests. The first type of request asked for advice on which machine translation systems to use or which ones could be trusted: 'Most effective apps/websites that reduce risk of miscommunication' (Principal Social Worker); 'Knowing what services are available and what ones have been the most user friendly AND accurate and trustworthy' (Healthcare Support Worker). These requests echo questions that I often hear from friends and family in conversations about my research. I am often asked: which machine translation tool should I use? My answer is usually frustrating. As discussed above in the section Seeking Accuracy Assurances, no tool can be guaranteed to be effective. All of them should be approached with caution.

The second – and, I would argue, most important – prototypical request called for more transparent workplace procedures. In some cases, the study itself prompted participants to reflect on the importance of policies and of thinking through the potential outcomes of different communication methods: 'Nothing specific [i.e., no specific training] other than awareness that it's an option, although now [that] I think about it I think it would [be] useful to know if there are any legal considerations e.g. accuracy of [machine] translation leading to misunderstanding or invalidating consent' (Emergency Medicine Doctor).

Machine translation is such a common feature of everyday life that this emergency doctor had not previously considered the possibility of legal

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006

consequences. Some participants expressed stronger concerns in this respect, although they tended to emphasise slightly different aspects of ethics, professional conduct or information governance. Some of them wanted reassurance from their managers that their methods of dealing with language barriers were acceptable. Others just wanted the use of machine translation to be formally recognised. Irrespective of what these participants decided to emphasise, implicit in all their contributions was the fact that using machine translation was often an under-the-radar, informal practice. The information they asked for included: 'Permissions within the organisation and whether it needed to be declared' (Senior Manager, Medical/healthcare); 'A session that focuses on how GDPR [general data protection regulation] impacts on them [machine translation systems] to ensure practice remains information governance compliant' (Social Worker); 'The service we use crashes a lot as it needs a good signal. To help us know what [to] do if the [professional] translator service crashes!' (Theatre Scrub Sister); 'Maybe the situations in which it is acceptable to use and what information is allowed to be said across the translator, such as, [is] consent accepted via automatic translator?' (Cardiac Physiologist); 'Just even the mention of using them' (Call Handling Supervisor, Emergency Services).

As can be seen in these comments, some requests for information came from managers and supervisors. The fact that even those in managing roles were unaware of governance expectations speaks to a dearth of best-practice guidance but also to the potential difficulties of keeping up with changing technology-mediated behaviours in bureaucratic environments. The UK's NHS, for one example of such an environment, is known for the unwieldy nature of its various processes and divisions. In England alone, sites and facilities are locally managed by different bodies depending on the area of the country.[86] Perceptions of what is accepted may therefore differ between managing bodies. Internal policies also differ. Some policies may mention machine translation.[87] In an investigation conducted in 2023, most NHS policies did not.[88] Deriving a model of behaviour from such a fragmented landscape is difficult.

Many of the professionals quoted in this book trusted machine translation because they did not see an alternative. Others were deeply distrustful of it while many were inclined to trust it uncritically. These starkly different views are testament to the importance of initiatives aimed at bolstering the critical skills of professionals who might not have considered the ethical implications of using machine translation and who might not see a need for training or awareness-raising. Individuals cannot be presumed to possess the literacy and meta-literacy competencies described in Chapter 1. As discussed in Chapter 2, prudential judgement involves 'deliberating and choosing well'. It requires practice and training. In the Conclusion, I come back to principles to consider in this type of deliberation.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 09:14:21, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.006