Introduction

In July 2022, I fell off an electric scooter on my way to work. I had used Bristol's scooter hire scheme a few times without any incidents. On that day, however, as I was already on the university campus looking for a place to park, a dip in the road surface sent me flying into the air. I was wearing a helmet, but the helmet did not stop me from suffering a severe laceration to my chin. A kind member of the university's security services witnessed the fall and took me to hospital in a campus security vehicle. Once in hospital I was quickly given first aid and then had to stay in the waiting room in between assessments, X-rays and eventually the procedure to close the wound. As I waited there with a bandage around my face, I could see many patients arrive. Some were checked in and told to wait. Others were told to use a non-emergency service instead. One of the arrivals caught my attention. A woman approached the reception desk and started to communicate with the staff using what looked to be a machine translation app like Google Translate. The woman held her smartphone close to her face and would sometimes turn it around for the hospital staff to see the screen. The staff seemed confused at first. One of them grabbed a paper from the desk and I wondered if a professional interpreter would be called, but the woman continued to use the app on her phone. There was some pointing and gesticulating on both sides, and the hospital staff would sometimes speak more slowly and louder than normal. A few minutes later the interaction concluded and the woman left the emergency department.

While I tried not to stare, I was sitting nearby and could not help but notice how natural it was for the woman to use her phone to communicate over a language barrier. It looked as if she had done that before. The app allowed her query to be addressed quickly and easily. She did not have anyone with her who could assist her. So had it not been for the app, she would probably have needed to wait until the hospital staff called an interpreter or located a colleague who spoke her language.

It is not surprising, therefore, that individuals turn to machine or artificial intelligence (AI) translation tools for language assistance. This technology plays an increasingly important role for multilingual communities. But like many other digital tools, machine translation has advantages and disadvantages. It cannot be

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

blindly trusted. In 2019, researchers at the Olive View-UCLA Medical Center assessed the demand for language support at an emergency department in Los Angeles. They found that seventy-five per cent of spontaneously written discharge instructions provided in Spanish contained obvious translation errors,[1] most of which because of machine translation.[2] One message read, in Spanish, 'your United States is normal', an error that resulted from the translation system's inability to recognise that 'US' was an abbreviation for 'ultrasound'.[3] Although this technology is constantly advancing, its use can be dangerous. In the scene I witnessed in hospital, mistranslations could have downplayed the severity of the woman's symptoms or outright misinformed the hospital staff.

The woman I saw in hospital is not alone. Many of those who are fortunate enough to have access to a smartphone or computer, and an internet connection, will have at least considered using a tool like Google Translate. This type of tool is often used for tasks of little consequence such as checking what others are saying on social media or reading product reviews on shopping websites. Many uses of these tools also pose a higher level of risk, whether it be in hospitals, in the legal system, in police questioning or in other interactions with potentially life-changing outcomes. For many of those involved in these interactions, whether as service users or providers, machine translation may often seem like the only practical way of communicating.

Consider a nurse who is on a night shift doing the ward round in a hospital in the UK. Consider the possibility that a Polish-speaking patient under the nurse's care needs assistance. No one on the ward other than the patient speaks Polish. Like in most medical settings in the UK, the hospital should have a process in place for the nurse to request assistance from a professional interpreter who could mediate the conversation, in person or over the phone. It is late, however, and the wait for the interpreter could be long. The nurse also knows that the patient is not in a critical condition. All that is required is to ask for permission to check the patient's blood pressure, to explain a medication that will be administered, and to ask if the patient needs help with mobilising to the toilet and back into bed. The nurse decides to use Google Translate to ask those questions and to understand the answers. The translations may have some obvious errors, but care is delivered successfully, and both nurse and patient are happy.

A scene like this one was reported by a nurse in a questionnaire I distributed to UK residents in 2019.[4] The nurse's assessment of machine translation in this case was positive. The decision to use Google Translate in this type of interaction is nevertheless fraught with risk. Misunderstandings could affect the patient's ability to provide informed consent for medications and medical procedures. The nurse in turn could be liable for any adverse effects that the use of Google Translate might have for the patient, which could have legal

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

consequences not only for the nurse personally, but also for their superiors and for the authority that runs the hospital. Luckily, none of these eventualities were reported, but these are some of the potential outcomes to consider when relying on machine translation to provide a service, and there are many cases where the use of this technology has gone awry.

Consider a traffic police officer who stops a motorist on a highway in Kansas in the United States. The motorist is a Spanish speaker with low English proficiency. The officer does not speak Spanish and decides to bring the motorist into the patrol vehicle to ask a few questions through Google Translate using the vehicle's laptop. As required by law in the US, the officer asks for consent to search the motorist's car. The officer finds illegal substances in the car and the motorist is charged with a crime. From the perspective of the officer, Google Translate served its purpose. But the motorist challenged the charges in court. The argument from the defence team was that consent for the search was invalid because of the language barrier. The judge ruled in favour of the motorist. The use of Google Translate was not enough to demonstrate that the motorist could understand what he was being asked, so the officer was found to have searched the vehicle without consent and the evidence had to be suppressed.[5] This widely discussed case shows that even when at face value machine translation seems to work, its use can have significant repercussions. By admission of the police officer himself,[6] a human interpreter would have been preferable to Google Translate as a method of communicating with the motorist. But in fairness to professionals who find themselves in a similar scenario, decisions around what action to take in the face of a language barrier can be highly complex, and the convenience of machine translation just adds to this complexity.

From Science Fiction to Reality

Using technology to translate between natural languages has for a long time been a human ambition. The idea dates back at least to the 17th century when numerical representations of language were proposed by Descartes and Leibniz.[7] Imaginary prototypes with impressive powers have been described in popular culture over the years, but in real life machine translation tools have significant limitations. They can nevertheless convert texts and speech between a growing number of languages at speed, even if with errors. Although this feat of language engineering has changed how we communicate, the popularity of translation tools is a relatively recent phenomenon, not least because initial attempts at developing these systems were not very impressive.

At first, in the mid 20th century, machine translation development was restricted to small-scale experiments.[8] Research advanced in the second half of the century, but in 1966 machine translation researchers were dealt

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

a significant blow when a United States committee issued what would come to be widely known as the ALPAC report.[9] This report questioned the technology's potential and prompted cuts to research funding in the United States.[10] Development slowed down after the report, but several systems were launched in the 1970s and 1980s, including METEO, a system designed to translate weather forecasts, and SYSTRAN,[11] a well-known provider of machine translation to this day.

As internet accessibility gathered pace towards the end of the 20th century, the social status of machine translation unsurprisingly changed. The first online, free-of-charge machine translation system was launched in 1997.[12] At the turn of the millennium this system had at least a million users.[13] At the time of writing over two decades later, the ability to translate texts automatically is a staple of digital living. The functionality is embedded in web browsers, websites and in social media applications. It is within reach on our phones, tablets and smartwatches through applications such as Google Translate, Apple Translate and Microsoft Translator. Machine translation is also part of the capabilities of chatbots such as ChatGPT and Gemini. From time to time, big internet companies release usage statistics for their translation products and the figures are always staggering. To pick one such statistic, just in the month of April 2021, Google Translate processed over twenty billion webpage translation requests.[14] Given this large usage volume, for many people around the world machine translation is likely to be the main method used to consume multilingual content. This technology is central to conceptions of language and of translation, and it is not inconceivable that for most people 'translation' is or may come to be understood as machine translation by default.

Uses of machine translation for communication only recently started to be examined in translation (written) and interpreting (spoken) studies. Some early work on the topic can be found under the heading of other subjects,[15] but there are several reasons why translation and interpreting have arguably been slow to catch up with the mainstream status of machine translation tools. For one thing, the work of professional translators has been directly affected by machine translation deployment. Translators are often expected to edit machine translations to boost their productivity, and their fees can be adjusted as a result, sometimes in ways that are financially unfavourable to them.[16] The use of machine translation as a productivity booster has been ongoing for quite some time. At the Pan-American Health Organization, for instance, language professionals were editing machine translation output over a decade before the technology became available free of charge on the internet.[17] Machine translation has therefore introduced significant changes to human translation processes, which for many years have been more pressing issues in translation and interpreting research compared to more recent uses of machine translation as a communication or reading tool.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

Translator and interpreter education is also a factor in why translation scholars have for some time focused on the language industry over other machine translation use settings. Although translation and interpreting education can play a wide range of roles, it has often had vocational purposes. Degree programmes in translation are usually expected to prepare students for professional practice, a goal that can shape teaching methods as well as research agendas. As the popularity of machine translation continues to grow, however, the remit of translation and interpreting research is widening. It is not just future language professionals who need information about this technology. There is a need for society at large to be attuned to the strengths and limitations of translation tools.

This book, therefore, is about uses of machine translation that take place outside the language services industry, in scenarios like the ones cited above involving providers of critical services and the individuals they serve. The book is focused on machine translation as a tool for information consumption and distribution – for understanding and being understood. I approach this topic primarily as a translation scholar, but throughout this project I have attempted to broaden out that perspective to consider the inherent interdisciplinary nature of the subject.

The project is grounded in direct accounts from critical professionals. It covers five sectors: healthcare, social assistance (i.e., social care/services), emergency services, legal services and the police. I refer to the professionals working in these sectors as the 'service providers'. I call the individuals they serve the 'service users'. While these services are not necessarily or directly publicly funded, I refer to them as 'public services' at different points in the book as a reflection of the reach and importance of these services for any member of a community. This is also why I call these services 'critical'. I note, however, that there are important public services, such as education, which are outside the scope of the selected sectors. I also note that the selected services were not equally prominent in the data gathered for the project. Healthcare and social assistance had larger participant groups (see Notes on Methodology) so many of the examples analysed in the book are from these sectors, even though all the above sectors are represented.

All professionals who took part in the project were based in the UK and most of the public-domain examples I discuss are from predominantly English-speaking countries, especially the UK and the United States. The discussion is therefore inevitably restricted to these contexts. At the same time, English-dominant countries are fertile ground for research on this subject. Given the status of English as an international language, the ability to speak English may be expected of newcomers to these countries, which often shapes institutional approaches to translation and multilingual communication. My previous work has shown that, in the UK, the use of machine translation is extremely

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

common.[18] Meanwhile, levels of political and financial support for language services in the UK are often low, as I will discuss in Chapter 1. This type of context is particularly vulnerable to machine translation misuse.

The communicative contexts examined here call for complex decision-making, but the objective of the book is not to spell out what to do in different scenarios. My intentions are broader, not least because part of the challenge posed by machine translation is that its use consequences are dynamic. The technology may be lifesaving in some cases. In others, translation errors can be highly detrimental and yet difficult to spot for individuals who do not speak both the source (i.e., input) and target (i.e., output) languages. Engaging with machine translation responsibly is therefore often a matter of being sensitive to its risks and discerning the best course of action based on circumstances that will change across professional, linguistic and cultural contexts. This means that there are no easy, overarching formulas to follow. There are, however, important principles to consider in relation to uses of these tools in high-stakes environments. This book will eventually arrive at these principles, but first a few explanations are required.

What Is Machine Translation After All?

First, machine translation needs to be defined. Although on first impression defining this technology may seem unproblematic, it is by no means a simple task. Part of the challenge lies in the fact that machine translation is somewhat amorphous. Often it is the main functionality of dedicated applications. That is the case of tools such as Google Translate, Microsoft Translator, Apple Translate, DeepL and other similar applications that are available as mobile apps, web applications, or both. Machine translation can also be an ancillary feature. For example, web browsers and websites can have integrated machine translation functionality, which allows the website visitors to access the content in different languages. As mentioned, machine translations can also be obtained from chatbots like ChatGPT. Here the underlying technology is generative AI or, more specifically, a large language model. Large language models are not the same as machine translation models. For one thing, large language models can be applied to a wide range of tasks and not just to language translation. The steps taken to develop the two types of model also differ, although all such models are based on machine learning – that is, models that 'learn' patterns and associations by processing existing data.

Machine translations produced by all these tools and models are within the scope of this book. My interest here is not the underlying mechanism used to generate the translations, but rather how the translations are used. The focus on use means that what I refer to as machine translation can in fact be the product

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

of slightly different technologies depending on the specific system – for example, a dedicated machine translation tool or a large language model. On the other hand, tools that do not provide automatically generated translations of either text or speech are not covered by the discussion. For example, uses of grammar checkers, electronic dictionaries or pronunciation tools that do not also provide automated translations are not the focus of my analysis, although when these technologies are used together with a machine translation tool this broader context is considered.

Terminological inconsistency is another complicating factor in discussions of this subject. In this book, I use the terms 'machine translation' or 'AI translation'. I use these terms interchangeably and sometimes I refer to them as simply 'the technology'. Other terms that can be used to refer to machine translation include automated translation, automatic translation or simply online translation. Mechanical translation was also a term used in early research on the topic.[19] For the most part machine translation is the term that became prevalent in computer science and then in translation studies, so it is the term I use here most often even though this term may well fall into disuse, a possibility I return to at the end of the Introduction.

The issue of how to define translation itself – whether human or otherwise – can also be a source of confusion. In academia and the language industry, translation often refers to written texts, whereas interpreting is used for spoken content. There are at least two reasons why this distinction is not as straightforward as it seems. First, non-linguists may not be familiar with this distinction at all. A sign of this is the many blog posts where language service providers try to explain this difference to prospective clients.[20] Official government documents sometimes do not make this distinction either. Terms such as a 'phone translation service' can be found in public policy documentation,[21] even though phone calls usually involve speech and in academic contexts go under the heading of interpreting rather than translation. Second, technologies themselves can blur this distinction. Many machine translation tools, especially those that are available as mobile apps, include speech recognition and synthesis. This means that users can speak and hear in a different language a version of what they said. When speech is used as both the input and output of a machine translation tool, this can be referred to as machine interpreting. Speech may be used intermittently in the same conversation, or it may be used just as the output or just the input – for example, if users speak but then just show written translations to their interlocutors on a screen, or if they type the input and play an audio version of the output out loud. Technology therefore facilitates hybrid forms of communication that are increasingly likely to involve mixtures of written and spoken language. For these reasons, I am not too strict in trying to distinguish whether uses of machine translation constitute interpreting or translation.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

Similarly, when I refer to machine translation tools or machine translation as a type of technology, I take machine interpreting and other uses of speech to be included in the discussion. I do, however, draw on separate frameworks and theoretical traditions that are specific to translation or interpreting whenever that is relevant to the analysis. More crossover between these two fields is something that colleagues and I suggested in previous research and which examples I discuss in this book also call for.[22] I note, however, that uses of machine translation for sign language interpreting are outside the scope of what I set out to cover.[23]

Lastly, in delineating what I mean by machine translation it is worth touching on the issue of hardware. The hardware used to access machine translation will often consist of devices used for everyday tasks, such as desktops, laptops, smartphones, tablets, smartwatches or smart home devices. Dedicated machine translation devices are also available. They are often referred to as instant, portable or pocket translators. Some of these dedicated devices may look like a smartphone. Others are simply a pair of earbuds. I do not restrict the discussion to a specific type of hardware, so the physical presence of machine translation in all these forms is considered. This physical presence is in fact important for understanding the technology's social consequences, especially in relation to its portability and how it might influence decision-making, a topic I revisit in Chapter 2 and again later in the book.

In summary, my understanding of machine translation in this project is quite broad. It includes different types of machine translation technologies, different ways of interacting with these technologies (e.g., using written texts and/or speech) and different methods of accessing them (e.g., via a website, in a mobile app or using a dedicated piece of hardware). This broad approach to the topic presents challenges – not least terminological – but it is necessary if the discussion is to do justice to the wide range of possibilities presented by machine translation when it is used as a communication tool.

How to Keep Up with Technological Change?

Technologies evolve fast, so it will not take long after publication of this book for the translation technology landscape to change. In any case, this book is not about specific incarnations of machine translation. The book is informed by machine translation tools that are available at the time of writing, but the discussion is hoped to transcend these specific tools.

Looking back at how machine translation has evolved, key methodological paradigms in the development of the technology have included rule-based systems, statistical systems, neural systems and, more recently, large language models. Statistical systems became more prominent than rule-based systems in the 1990s.[24] Neural systems then became more prominent than statistical ones

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

from the mid 2010s.[25] Since tools like ChatGPT gained popularity in the early 2020s,[26] large language models and neural machine translation systems are now both widely used.

Rule-based systems require developers to formulate specific rules for the program to follow. In other words, linguists and programmers 'teach' computers how to translate by designing a formula that the program can apply. Statistical systems, by contrast, learn by example. They use existing data – namely, translations, preferably high-quality ones – as a benchmark. When presented with new texts, statistical systems guess the translations that are statistically more likely to be correct based on the examples previously seen in the learning phase. Neural systems also learn by example, but their learning mechanism is much more sophisticated. Neural systems take more context into account. Like large language models, they can also learn linguistic associations multilingually rather than just for a specific language pair.[27] In the case of large language models specifically, a key characteristic of the learning process is its large scale. These models are expensive and require vast amounts of data,[28] hence their wider range of applications.

Translations generated by neural systems tend to have fewer errors than those of previous paradigms.[29] At the time of writing, research also shows that the output of translation-specific systems can sometimes be superior to that of large language models, especially for non-European languages.[30] The type of errors usually made by each type of system can also vary,[31] but the system is not the only factor that can affect the quality of the output. There will also be language-[32] and text-specific variation,[33] to mention just two additional factors. So irrespective of the underlying technology, machine translation outputs will almost always be dynamic. The same system and methodology can produce different results, for example when they encounter different text types or, in the case of large language models, depending on the prompt (i.e., the question or instruction provided to the tool).

For my analysis in this book, this means that change can be taken for granted. Technologies often converge and become more integrated over time,[34] and this phenomenon is under way with machine translation too. The arrival of large language models may well in time obviate the need for dedicated translation tools. The term 'machine translation' itself may eventually be phased out. Machine translations may come to be understood as just one of the possible outputs of AI systems that are increasingly embedded in online applications and thereby in our everyday lives. Our interaction with these systems is also likely to change. Users can already tailor the output of large language models to their needs by asking questions and giving instructions. This type of human–computer interaction is likely to become more sophisticated as user interfaces become more seamless and sensitive to the nuances of human communication.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

None of these possibilities is particularly problematic for my objectives here. Irrespective of the shape taken by future translation technologies, by accepting change upfront I can move the discussion to a level of analysis that focuses on the social dynamics that surround these tools. Machine translation's social contexts are not static either, but it is only by analysing the technology's presence in these contexts that its impact on society can be understood.

The Structure of the Book

I structure the book into the following chapters:

Chapter 1 reviews different factors that can affect the risk-benefit ratio of using a machine translation tool. The chapter examines key concepts, such as usability and AI literacy, and proposes a typology of machine translation use followed by a discussion of specific examples.

Chapter 2 provides an analysis of different theoretical frameworks that can inform the understanding of the social status of machine translation technology. This chapter draws on a range of socially oriented theories of technology as well as on theories of ethics. The focus of the chapter is on post-phenomenology and on virtue ethics. The chapter looks at how these two traditions can provide useful instruments for dealing with the inherent uncertainty that surround high-stakes uses of machine translation like those discussed at different points in the book.

Chapters 3–5 are based on a UK survey of the professional sectors selected for analysis. In Chapter 3, I discuss uses of machine translation that pose ostensibly low levels of risk. I examine how different communication methods can be combined in these contexts. I also highlight differences between what I call ancillary communication and communication that is central to the provision of a service. I examine the dynamic nature of these two communication types and the ways in which they often merge and overlap.

Chapter 4 is about high-risk use contexts. I look at specific reasons why machine translation is used in these contexts. These reasons included service user preferences, urgency, notions of person-centred care, and human language services that often fell short of expectations. This chapter also examines machine translation use guidelines and the issue of consent. I consider two types of consent, namely using machine translation to seek consent and whether the use of machine translation itself is consensual.

Chapter 5 is about trust. I examine the concept of trust and how some of its formulations can be applied to uses of machine translation. Unlike in other chapters in the book, here I consider project participants who had not used machine translation before. I analyse their perceptions of the technology, why they had not used it and whether they would trust it. This chapter also looks at

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001

the question of professional development and at the types of information that machine translation users would like their workplace training to include.

Chapter 6 is a case study of multilingual communication practices in social work. I interviewed UK social workers to delve deeper into their experience of dealing with language barriers and of using machine translation tools. This chapter discusses concepts such as linguistic capital and low-resource languages. It also looks at how different social and communicative needs can overlap and have a compounding effect on the service users.

Finally, the Conclusion considers the data presented in the previous chapters and outlines some of the broader implications of the book. I summarise different types of machine translation use based on the extent of reliance on the technology and on its use context. I also propose principles to be considered in the deployment of machine translation tools and in wider discussions of their social effects.

The book is intended to be of use to a diverse audience including researchers, critical service professionals or a more general readership interested in multilingual communication and AI. Those seeking information emerging directly from the data gathered in the project may wish to focus on Chapters 3–6. Those seeking a discussion of relevant theories may wish to focus on Chapter 2 and Chapter 5. Real-life examples are provided throughout the book. The study methodology is explained in the chapter Notes on Methodology, after the Conclusion.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 22:00:57, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009769372.001