investigate whether these mouse findings could have relevance for humans, the authors performed an analysis of clinical studies in which people with melanoma or lung cancer were treated with a type of immunotherapy known as immune-checkpoint blockade. Concomitant administration of painkillers that inhibit prostaglandin synthesis was associated with slightly improved response rates to immunotherapy, suggesting that pharmacological inhibition of prostaglandin E2 synthesis might reverse cross-resistance to targeted therapy and immunotherapy. This possibility will need to be investigated in clinical trials.

Elewaut and colleagues' work touches on the intriguing question of how the phenotype (characteristics) of cancer cells affects the cancer–immunity cycle, with a focus on cancer cells' effects on monocytes that interact with tumour-specific T cells. In this way, cancer cells shape responses to T-cell immunotherapy. Consistent with Elewaut and colleagues' findings in mice, a clinical study reported that the presence of interferon-activated monocytes in melanoma tissues is associated with positive responses to immune-checkpoint blockade[9]. Conversely, high levels of prostaglandin E2 in both human and mouse melanoma tissues have been shown to limit the proliferation of CD8$^+$ T cells[10,11], effectively disrupting the cancer–immunity cycle. These results fit into the wider context of work showing that a signalling pathway termed the GPCR–Gα_s axis, which is targeted by prostaglandin E2, is a key regulator of the killing functions of CD8$^+$ T cells in tumour tissues[12]. This provides a scientific rationale for considering the inhibition of this signalling axis in immunotherapy strategies that involve more than one treatment.

Anthony C. Buzzai and **Thomas Tüting** are in the Department of Dermatology, University Hospital and Health Campus Immunology Infectiology and Inflammation (GC-I3), Otto von Guericke University, 39120 Magdeburg, Germany.
e-mail: thomas.tueting@med.ovgu.de

1. Elewaut, A. *et al. Nature* **637**, 716–725 (2025).
2. Haas, L. *et al. Nature Cancer* **2**, 693–708 (2021).
3. Mellman, I., Chen, D. S., Powles, T. & Turley, S. J. *Immunity* **56**, 2188–2205 (2023).
4. Pittet, M. J., Di Pilato, M., Garris, C. & Mempel, T. R. *Immunity* **56**, 2218–2230 (2023).
5. Kruse, B. *et al. Nature* **618**, 1033–1040 (2023).
6. Wakim, L. M. & Bevan, M. J. *Nature* **471**, 629–632 (2011).
7. Joffre, O. P., Segura, E., Savina, A. & Amigorena, S. *Nature Rev. Immunol.* **12**, 557–569 (2012).
8. Coulombe, F. *et al. Immunity* **40**, 554–568 (2014).
9. Schlenker, R. *et al. Med* **5**, 759–779 (2024).
10. Morotti, M. *et al. Nature* **629**, 426–434 (2024).
11. Lacher, S. B. *et al. Nature* **629**, 417–425 (2024).
12. Wu, V. H. *et al. Nature Immunol.* **24**, 1318–1330 (2023).

The authors declare no competing interests.

Forum: Artificial intelligence

Striving for open-source and equitable translation

US technology company Meta has produced an AI model that can directly translate speech in one language to speech in another. Two scientists discuss the technical feats and ethical questions that underpin this advance in machine translation. **See p.587**

The paper in brief

- **Billions of people around the world regularly communicate online in languages other than their own.**
- **This has created huge demand for artificial intelligence (AI) models that can translate both text and speech.**
- **But most models work only for text, or use text as an intermediate step in speech-to-speech translation, and many focus on a small subset of the world's languages.**
- **On page 587, the SEAMLESS Communication Team[1] addresses these challenges to come up with key technologies that could make rapid universal translation a reality.**

Tanel Alumäe

Neat tricks and an open outlook

The SEAMLESS authors devised an AI model that uses a neural-network approach to translate directly between around 100 languages (Fig. 1a). The model can take text or speech inputs from any of these languages and translate them into text, but it can also translate directly to speech in 36 languages. This speech-to-speech translation is particularly impressive because it involves an 'end-to-end' approach: the model can directly translate, for example, spoken English into spoken German, without first transcribing it into English text and translating it into German text (Fig. 1b).

To train their AI model, the researchers relied on methods called self-supervised and semi-supervised learning. These approaches help a model to learn from huge amounts of raw data – such as text, speech and video – without requiring humans to annotate the data with specific labels or categories that provide context. Such labels might be accurate transcripts or translations, for example.

The part of the model that is responsible for translating speech was pre-trained on a massive data set containing 4.5 million hours' worth of multilingual spoken audio. This kind of training helps the model to learn the patterns in data, making it easier to fine-tune the model for specific tasks without the need for large amounts of bespoke training data.

One of the SEAMLESS team's savviest strategies involved 'mining' the Internet for training pairs that align across languages – such as audio snippets in one language that match subtitles in another. Starting with some data that they knew to be reliable, the authors trained the model to recognize when two pieces of content (such as a video clip and a corresponding subtitle) actually match in meaning. By applying this technique to vast amounts of Internet-derived data, they collected around 443,000 hours of audio with matching text, and aligned about 30,000 hours of speech pairs, which they then used to further train their model.

These advances aside, in my view, the biggest virtue of this work is not the proposed idea or method. Instead, it's the fact that all of the data and code to run and optimize this technology are publicly available – although the model itself can be used only for non-commercial endeavours. The authors describe their translation model as 'foundational' (see go.nature.com/3teaxvx), meaning it can be fine-tuned on carefully curated data sets for specific purposes – such as improving translation quality for certain language pairs or for technical jargon.

Meta has become one of the largest supporters of open-source language technology. Its research team was instrumental in developing PyTorch, a software library for training AI models, which is widely used by companies such as OpenAI and Tesla, as well as by many researchers around the world. The model introduced here adds to Meta's arsenal of foundational language technology models, such as the Llama family of large language models[2], which can be used to create applications akin to ChatGPT. This level of openness is a huge advantage for researchers who lack the massive computational resources needed to build these models from scratch.

Although this technology is exciting, several obstacles remain. The SEAMLESS model's ability to translate up to 100 languages is impressive, but the number of languages spoken around the world is around 7,000. The tool also struggles in many situations that humans handle with relative ease – for example, conversations in noisy places or between people with strong accents. However, the authors' methods for harnessing real-world data will forge a promising path towards speech technology that rivals the stuff of science fiction.

Tanel Alumäe is in the Laboratory of Language Technology at Tallinn University of Technology (TalTech), 19086 Tallinn, Estonia.
e-mail: tanel.alumae@taltech.ee

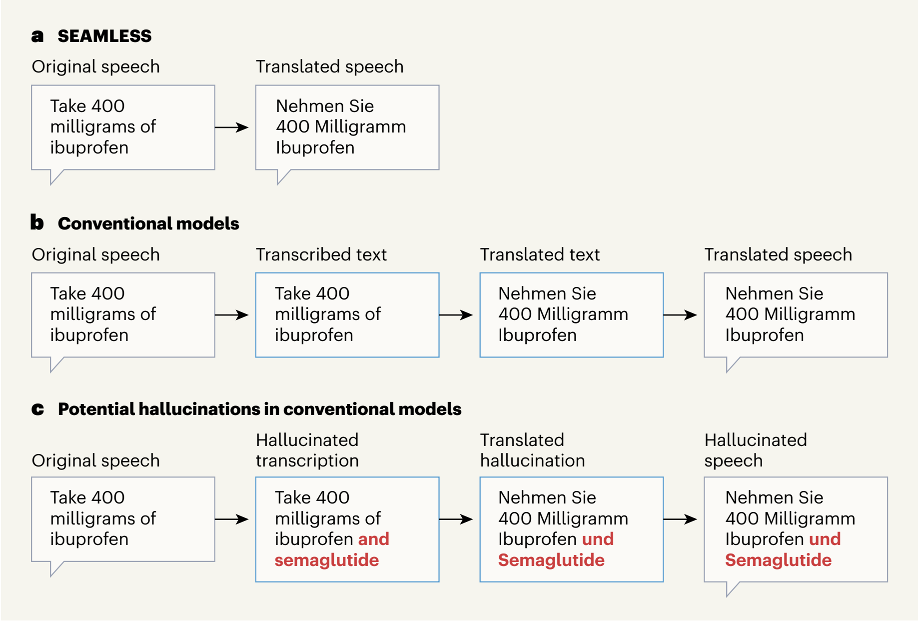

Figure 1 | Speech-to-speech machine translation. a, The SEAMLESS Communication Team[1] devised an artificial intelligence (AI) model that can directly translate speech in around 100 languages into speech in 36 languages. **b**, Conventional AI models for speech-to-speech translation typically use a cascaded approach, in which speech is first transcribed and translated into text in another language, before being converted back into speech. **c**, Certain conventional models can hallucinate (generate incorrect or misleading outputs), which could lead to considerable harm if such models were used for machine translation in high-stakes settings, such as health care.

Allison Koenecke
Keeping users in the loop

Speech-based technologies are increasingly being used for high-stakes tasks – taking notes during medical examinations, for instance, or transcribing legal proceedings. Models such as the one devised by SEAMLESS are accelerating progress in this area. But the users of these models (doctors and courtroom officials, for example) should be made aware of the fallibility of speech technologies, as should the individuals whose voices are the inputs.

The problems associated with existing speech technologies are well documented. Transcriptions tend to be worse for English dialects that are considered non-'standard' – such as African American English – than they are for variants that are more widely used[3]. The quality of translation to and from a language is poor if that language is under-represented in the data used to train the model. This affects any languages that appear infrequently on the Internet, from Afrikaans to Zulu[4].

Some transcription models have even been known to 'hallucinate'[5] – come up with entire phrases that were never uttered in audio inputs – and this occurs more frequently for speakers who have speech impairments than it does for those without them (Fig. 1c). These sorts of machine-induced error could potentially induce real harm, such as erroneously prescribing a drug, or accusing the wrong person in a trial. And the damage disproportionately affects marginalized populations, who are likely to be misheard.

The SEAMLESS researchers quantified the toxicity associated with their model (the degree to which its translations introduce harmful or offensive language)[6]. This is a step in the right direction, and offers a baseline against which future models can be tested. However, given the fact that the performance of existing models varies wildly across languages, extra care must be taken to ensure that a model can adeptly translate or transcribe certain terms in certain languages. This endeavour should parallel efforts among computer-vision researchers, who are working to improve the poor performance of image-recognition models in under-represented groups and deter the models from making offensive predictions[7].

The authors also looked for gender bias in the translations produced by their model. Their analysis examined whether the model over-represented one gender when translating gender-neutral phrases into gendered languages: does "I am a teacher" in English translate to the masculine "*Soy profesor*" or to the feminine "*Soy profesora*" in Spanish? But such analyses are restricted to languages with binary masculine or feminine forms only, and future audits should broaden the scope of linguistic biases studied[8].

Going forward, design-oriented thinking will be necessary to ensure that users can properly contextualize the translations offered by these models – many of which vary in quality. As well as the toxicity warnings explored by the SEAMLESS authors, developers should consider how to display translations in ways that make clear a model's limitations – flagging, for example, when an output involves the model simply guessing a gender. This could involve forgoing an output entirely when its accuracy is in doubt, or accompanying low-quality outputs with written caveats or visual cues[9]. Perhaps most importantly, users should be able to opt out of using speech technologies – for example, in medical or legal settings – if they so desire.

Although speech technologies can be more efficient and cost-effective at transcribing and translating than humans (who are also prone to biases and errors[10]), it is imperative to understand the ways in which these technologies fail – disproportionately so for some demographics. Future work must ensure that speech-technology researchers ameliorate performance disparities, and that users are well informed about the potential benefits and harms associated with these models.

Allison Koenecke is in the Department of Information Science, Cornell University, Ithaca 14853, New York, USA.
e-mail: koenecke@cornell.edu

1. SEAMLESS Communication Team. *Nature* **637**, 587–593 (2025).
2. Grattafiori, A. *et al.* Preprint at https://doi.org/10.48550/arXiv.2407.21783 (2024).
3. Koenecke, A. *et al. Proc. Natl Acad. Sci. USA* **117**, 7684–7689 (2020).
4. Nekoto, W. *et al.* in *Findings of the Association for Computational Linguistics: EMNLP 2020* (eds Cohn, T. *et al.*) 2144–2160 (Association for Computational Linguistics, 2020).

5. Koenecke, A., Choi, A. S. G., Mei, K. X., Schellmann, H. & Sloane, M. in *FAccT '24: Proc. 2024 ACM Conference on Fairness, Accountability, and Transparency* 1672–1681 (Association for Computing Machinery, 2024).
6. Al Sharou, K. & Specia, L. in *Proc. 23rd Annual Conference of the European Association for Machine Translation* (eds. Moniz, H. *et al.*) 171–180 (European Association for Machine Translation, 2022).
7. Yang, K., Qinami, K., Fei-Fei, L., Deng, J. & Russakovsky, O. in *FAT* '20: Proc. 2020 Conference on Fairness, Accountability, and Transparency* 547–558 (Association for Computing Machinery, 2020).
8. Cao, Y. T. & Daumé, H. III. *Comput. Linguist.* **47**, 615–661 (2021).
9. Robertson, S. & Díaz, M. in *FAccT '22: Proc. 2022 ACM Conference on Fairness, Accountability, and Transparency* 2223–2238 (Association for Computing Machinery, 2022).
10. Jones, T., Kalbfeld, J. R., Hancock, R. & Clark, R. *Language* **95**, e216–e252 (2019).

The authors declare no competing interests.

Marine biology

Algae use the light spectrum to sense depth

Marina Cvetkovska

Aquatic algae called diatoms have been found to use light-sensing proteins as depth indicators. Detecting depth-related changes in the intensity and spectrum of light enables algae to modulate their physiology accordingly. **See p.691**

Light is both an energy source and a key environmental signal that influences the physiology of life forms that rely on the process of photosynthesis. Aquatic phytoplankton face unique challenges in perceiving light, because it not only diminishes in intensity with water depth, but also shifts in spectral composition, with shorter wavelengths (blue and green light) penetrating deeper than do longer wavelengths (red and far-red light). On page 691, Duchêne *et al.*[1] investigate how an important group of marine phytoplankton called diatoms sense and respond to light signals at different depths. Combining environmental modelling and careful laboratory experiments, the authors suggest that diatoms use light-activated proteins (photoreceptors) called phytochromes as indicators of depth.

Phytochromes are protein sensors of red and far-red light that are found across various photosynthetic groups, including plants, phytoplankton and some bacteria[2]. These proteins regulate physiological and developmental processes by switching between two states, Pr and Pfr. The inactive Pr form absorbs and is activated by red light to become Pfr. In turn, this active Pfr form is inactivated by far-red light. Downstream responses depend on the proportion of Pfr, so phytochrome signals are influenced by both the intensity and the spectrum of perceived light[3]. Phytochromes are a major focus of plant research, but little is known about their role in aquatic phytoplankton. Because red and far-red light is notably diminished with increasing depth, this raises questions about the functional role of phytochromes in marine life.

Duchêne and colleagues shed light on this issue through a series of experiments. They first examined the geographical distribution of phytochrome-containing diatoms using DNA from water samples (environmental DNA) collected in global oceanic surveys, including the *Tara* Oceans campaign[4]. Diatoms inhabit all of the world's oceans, but phytochrome-encoding genes were predominantly detected in diatoms from temperate and polar regions and conspicuously absent from those in tropical areas. Other diatom photoreceptors did not exhibit a specific distribution at varying latitudes, supporting the hypothesis that the geographical spread of phytochromes reflects adaptation to particular environments rather than a sampling bias.

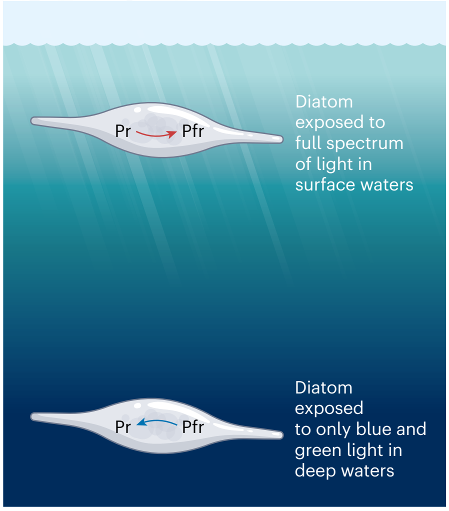

Figure 1 | A system for sensing depth. Duchêne *et al.*[1] report that marine algae known as diatoms use phytochrome proteins to sense depth. Phytochrome exists in two forms, Pr and Pfr, with the proportion of these two forms modulating cellular responses. Diatoms in water near the surface of the ocean are exposed to the full spectrum of light, and ultraviolet and red light there converts phytochromes to the Pfr form. In deep waters, only a limited spectrum of blue and green light penetrates, converting diatom phytochromes to the Pr form. Pr induces physiological responses such as regulation of the process of photosynthesis and enables diatoms to acclimatize to deeper waters.

The authors suggest that the phytochrome pattern they observed is an evolutionary adaptation to life in turbulent waters at higher latitudes in temperate and polar regions. These regions experience seasonal mixing and dynamic changes in the vertical water column, exposing diatoms to rapidly fluctuating light intensities and spectra[5]. By contrast, tropical oceans have a stable and permanently layered (stratified) water column, in which diatoms stay mainly at a constant depth. Thus, the ability to sense depth accurately and respond to changing conditions might be advantageous to phytoplankton in turbulent waters.

Duchêne *et al.* also discovered that, unlike red-sensing phytochromes in plants, diatom-specific phytochromes have a reduced sensitivity to far-red light. Instead, diatom phytochromes sense low-intensity blue and green light, which dominates at greater depths. Using a phytochrome-regulated fluorescent protein system, the authors reveal that this diatom photoreceptor is indeed activated by low-intensity blue and green light and induces cellular responses. Duchêne and colleagues predict that the activity of diatom phytochromes increases with depth and is driven mainly by blue and green wavelengths (Fig. 1). In surface waters, the inhibitory effects of ultraviolet and red light diminish phytochrome activity, but these wavelengths have minimal influence below the first few metres of water.

Photosynthesis converts light energy into chemical energy in the form of carbohydrates. This complex process is highly dependent on the availability of light and must be under tight control to satisfy cellular metabolic needs. This is particularly true for aquatic environments, where deep, dark waters can negatively affect the efficiency of photosynthesis. Phytoplankton have a range of strategies to overcome low levels of light, including remodelling the components of the photosynthetic machinery[6], but the regulation of these cellular changes remains poorly understood.

Duchêne and colleagues demonstrate that diatom phytochromes influence photosynthetic performance. In laboratory experiments that mimicked deep-water conditions that have low-intensity, blue-dominant spectra, phytochrome-deficient mutants had reduced photosynthetic efficiency compared with strains