

[I] 次の英文を読んで設問に答えなさい。

Most people say “I’m sorry” many times a day for a host of trivial affronts*—accidentally bumping into someone or failing to hold open a door.

(1) このような謝罪は簡単で、たいていはすんなりと受け入れられ、「大丈夫ですよ」といった返事がくることが多いだろう。

But when “I’m sorry” are the words needed to right truly hurtful words, acts or inaction, they can be the hardest ones to utter. (2) And even when an apology is offered with the best of intentions, it can be seriously undermined by the way in which it is worded. Instead of eradicating the emotional pain the affront caused, a poorly worded apology can result (あ) lasting anger and antagonism, and undermine an important relationship.

I admit to a lifetime of challenges when it comes to apologizing, especially when I thought I was right or misunderstood or that the offended party* was being overly sensitive. But I recently discovered that the need for an apology is less about i than the person who, for whatever reason, is offended by something I said or did or failed to do, regardless of my intentions.

(3) また、誠実な謝罪は、謝罪する側にも謝罪される側にも、思いがけない価値をもたらす強力な特効薬となり得ることも学んだ。 After learning that a neighbor who had assaulted me verbally was furious about an oversight I had not known I committed, I wrote a letter (い) hopes of defusing the hostility.

(4) 私は言い訳を一切せずに、礼儀と敬意を欠いたことを謝罪した。 I said I was not asking for or expecting forgiveness, merely that I hoped we could have a civil, X friendly, relationship going forward, then delivered the letter with a jar of my homemade jam.

Expecting nothing (う) return, I was greatly relieved when my doorbell rang and the neighbor thanked me warmly for what I had said and done. My relief was palpable*. (5) I felt as if I’d not only discarded an enemy but made a new friend, which is indeed how it was realized in the days that followed.

About a week later I learned that, according to the psychologist and author Harriet Lerner, (6) the wording of my apology was just what the “doctor” would have ordered. In the very first chapter of her new book, “Why Won’t You Apologize?,” Dr. Lerner points out that apologies followed by rationalizations are “never satisfying” and can even be harmful. “When ‘but’ is tagged on to an apology,” she wrote, it’s an excuse that counters the sincerity of the original message. The best apologies are short and don’t include explanations that can undo them.

Nor should a request for forgiveness be part of an apology. Forgiveness, Y it come, may depend on a demonstration going forward that the

offense will not be repeated.

“It’s not our place to tell anyone to forgive or not to forgive,” Dr. Lerner said in an interview. ⁽⁷⁾She disputes popular thinking that failing to forgive is bad for one’s health and can lead to a life full of bitterness and hate.

“There is no one path to healing,” she said. “There are many roads (道) letting go of corrosive emotions without forgiving, like therapy, meditation, medication, even swimming.”

The focus of an apology should be on what the offender has said or done, not on the person’s reaction to it. Saying “I’m sorry you feel that way” shifts the focus away from the person who is supposedly apologizing and turns “I’m sorry” into “ii,” the psychologist wrote.

As to why many people find it hard to offer a sincere apology, Dr. Lerner pointed out that “⁽⁸⁾humans are hard-wired for defensiveness. It’s very difficult to take direct, unequivocal* responsibility for our hurtful actions. ⁽⁹⁾It takes a great deal of maturity to put a relationship or another person before our need to be right.”

Offering an apology is an admission of guilt that admittedly leaves people vulnerable. There’s no guarantee (保証) how it will be received. It is the prerogative* of the injured party to reject an apology, even when sincerely offered. In some cases, the person may feel the offense was so enormous that it is impossible to accept a mea culpa*.

Righting a perceived wrong can be especially challenging when it involves family members, who may be inclined to cite history—he was abused by his father, or she was raised by a distant mother—as an excuse (口実) hurtful behavior. “History can be used as an explanation, not an excuse,” the psychologist said. “It should involve a conversation that allows the hurt party to express anger and pain if an apology, Z sincere, is to heal a broken connection.”

As she wrote: “Non-defensive listening to the hurt party is at the heart of offering a sincere apology.” She urges the listener not to “interrupt, argue, refute, or correct facts, or bring (持ち出す) your own criticisms and complaints.” Even when the offended party is largely at fault, she suggests iii.

Dr. Lerner views apology as “central to health, both physical and emotional. “‘I’m sorry’ are the two most healing words in the English language,” she said. “The courage to apologize wisely and well is not just a gift to the injured person, who can then feel soothed and released from obsessive recriminations*, bitterness and corrosive* anger. It’s also ⁽¹⁰⁾a gift to one’s own health, bestowing self-respect, integrity and maturity—an ability to take a clear-eyed look at how our behavior affects others and to assume responsibility for acting at another person’s expense.”

設問

問1 下線部（1）の日本語の内容を英語で言い表しなさい。

問2 下線部（2）を日本語に訳しなさい。

問3 空欄（あ）～（き）に入れるのに最もふさわしい語句を選択肢から選んで、その番号を解答欄に書きなさい。

- | | | | | |
|-----|----------|---------|------------------|-----------|
| （あ） | 1 for | 2 in | 3 of | 4 on |
| （い） | 1 beyond | 2 for | 3 in | 4 without |
| （う） | 1 by | 2 for | 3 from | 4 in |
| （え） | 1 at | 2 in | 3 to | 4 without |
| （お） | 1 as if | 2 as to | 3 in relation to | 4 in that |
| （か） | 1 as | 2 for | 3 of | 4 with |
| （き） | 1 about | 2 along | 3 to | 4 up |

問4 空欄 に入れるのに最もふさわしい表現を選択肢から選んで、その番号を解答欄に書きなさい。

- 1 apologizing
- 2 misunderstanding
- 3 me
- 4 you

問5 下線部（3）の日本語の内容を英語で言い表しなさい。

問6 下線部（4）の日本語の内容を英語で言い表しなさい。

問7 空欄 ～ に入れるのに最もふさわしい語句を選択肢から選んで、その番号を解答欄に書きなさい。解答欄は問4の解答欄の右横にある。

- | | | | | | |
|----|----------|-----------|-----------|------------|---------|
| X: | 1 if any | 2 if ever | 3 if not | 4 if only | 5 if so |
| Y: | 1 could | 2 might | 3 ought | 4 should | 5 would |
| Z: | 1 as | 2 because | 3 however | 4 how much | 5 to be |

問8 下線部（5）を日本語に訳しなさい。

問9 下線部（6）について、この記事の筆者の謝罪の言葉はどのようなものだったのか、具体的に日本語60字以内で簡潔に述べなさい。

問10 下線部（7）を日本語に訳しなさい。

問11 空欄 に入れるのに最もふさわしい表現を選択肢から選んで、その記号を解答欄に書きなさい。

- A I'm not really sorry at all
- B I'm not surprised by your reaction
- C I understand your feelings
- D You shouldn't feel like that

問12 下線部（8）の意味を表す英文を選択肢から選んで、その記号を解答欄に書きなさい。解答欄は問11の解答欄の右横にある。

- A humans always prefer attack to defense
- B humans automatically protect themselves
- C humans hardly ever defend themselves
- D humans sometimes do not like to accept apologies
- E humans think it is strange to defend themselves

問13 下線部（9）を日本語に訳しなさい。

問14 空欄 に入れるのに最もふさわしい表現を選択肢から選んで、その記号を解答欄に書きなさい。解答欄は問11と問12の解答欄の右横にある。

- A accepting full responsibility for causing offence
- B allowing the offended party to interrupt you
- C apologizing for one's own part in the incident
- D only refuting incorrect facts

問15 下線部（10）で、筆者がこのように言う理由は何か。具体的に日本語90字以内で説明しなさい。

[II] Read the passage below and answer the questions that follow it.

A study of newer, bigger versions of three major artificial intelligence (AI) chatbots shows that they are more inclined to generate wrong answers than to admit ignorance. The assessment also found that people aren't great at (a) the bad answers.

Plenty of attention has been (b) to the fact that the large language models (LLMs) used to power chatbots sometimes get things wrong or 'hallucinate*' strange responses to queries. José Hernández-Orallo at the Valencian Research Institute for Artificial Intelligence in Spain and his colleagues (c) such errors to see how they are changing as the models are getting bigger—making use of more training data, involving more parameters* or decision-making nodes* and gobbling up* more computing power. ⁽¹⁾ They also tracked whether the likelihood of errors matches up to human perceptions of question difficulty, and how well people can identify the wrong answers. The study was (d) in *Nature* on 25 September.

The team found that bigger, more refined versions of LLMs are, as expected, more accurate, thanks in large part to having been shaped with fine-tuning methods such as reinforcement learning from human feedback. ⁽²⁾ That is good news. But they are less reliable: among all the non-accurate responses, the fraction of wrong answers has increased, the team reports, because the models are less likely to (e) answering a question—for example, by saying they don't know, or by (f) the subject.

"They are answering almost everything these days. And that (g) more correct, but also more incorrect answers," says Hernández-Orallo. In other words, the chatbots' tendency to offer opinions beyond their own knowledge has increased. "That looks to me like what we would call bluffing*," says Mike Hicks, a philosopher of science and technology at the University of Glasgow, UK, who proposes the term 'ultracrepidarianism' to describe the phenomenon. ⁽³⁾ チャットボットは博識であるふりをするのが上手になってきている。

The result is that everyday users are likely to overestimate the abilities of chatbots and ⁽⁴⁾ that's dangerous, says Hernández-Orallo.

The team looked at three LLM families: OpenAI's GPT, Meta's LLaMA and BLOOM, an open-source* model created by the academic group BigScience. For each, they looked at early, raw versions of models and later, refined versions.

They tested the models on thousands of prompts that included questions on arithmetic, geography, and science, as well as prompts that tested the bots' ability to transform information, such as ア. They also ranked the

human-perceived difficulty of the questions—for example, a question about Toronto, Canada, was ranked as easier than a question about the lesser-known and smaller town of Akil, Mexico.

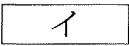
As expected, the accuracy of the answers increased as the refined models became larger and decreased as the questions got harder. And although ⁽⁵⁾it might be prudent for models to avoid answering very difficult questions, the researchers found no strong trend in this direction. Instead, some models, such as GPT-4, answered almost everything. The fraction of wrong answers among those that were either incorrect or avoided rose as the models got bigger, and reached more than 60%, for several refined models.

The team also found that all the models would occasionally get even easy questions wrong, meaning there is no ⁽⁶⁾'safe operating region' in which a user can have high confidence in the answers.

The team then asked volunteers to rank the answers as correct, incorrect or avoidant. ⁽⁷⁾The volunteers incorrectly classified inaccurate answers as being accurate surprisingly often—roughly between 10% and 40% of the time—across easy and difficult questions. “Humans are not able to supervise these models,” says Hernández-Orallo.

Hernández-Orallo thinks that developers should boost AI performance on easy questions, and encourage chatbots to decline to answer hard questions, so that people are able to better gauge the situations in which AIs are likely to be (i). “We need humans to understand: ‘I can use it in this area, and I shouldn’t use it in that area,’” he says.

Making chatbots more inclined to answer tricky questions looks (ii) and does well on leaderboards that rank performance, says Hernández-Orallo, but isn’t always helpful. “I’m still very surprised that with (iii) versions of some of these models, including o1* from OpenAI, you can ask them to multiply two very long numbers, and you get an answer, and the answer is (iv),” he says. That should be fixable, he adds. “You can put a threshold, and when the question is (v), get the chatbot to say, ‘no, I don’t know’.”

“There are some models which will say ‘I don’t know’, or ‘I have (vi) information to answer your question,’” says Vipula Rawte, a computer scientist at the University of South Carolina in Columbia. All AI companies are working hard to reduce hallucinations, and chatbots developed for (vii) purposes, such as medical use, are sometimes refined even further to prevent them from going beyond their knowledge base. But, she adds, for companies trying to sell all-purpose chatbots, “.

Questions

- 1 Select the most suitable verb from the list below to fill each of the spaces marked (a)~(g). Change the form of the verb where necessary, and write your answers on the answer paper. Do not select any of the words more than once.

analyse	avoid	change	demand	give
know	mean	publish	spot	support

- 2 Translate the underlined sentence marked (1) into Japanese.
- 3 What does the underlined word marked (2) refer to? Select the most suitable option from A ~ E and write the letter on your answer paper.
- A Bigger LLMs were expected to be more refined.
 - B Fine-tuning methods are used to create LLMs.
 - C Human learning can be reinforced with feedback from LLMs.
 - D More refined LLMs give more accurate answers.
 - E The team investigated the accuracy of LLMs.
- 4 Translate the underlined sentence marked (3) into English.
- 5 Write about 40 characters in Japanese to explain what the underlined word marked (4) refers to.
- 6 Select the most suitable option from A ~ D to fill the space marked ア. Write the letter on your answer paper. On the answer paper, the space for the answer to Q6 is on the same line as the space for Q3.
- A checking facts in a database
 - B generating random numbers
 - C identifying locations on a map
 - D putting a list in alphabetical order
- 7 Referring to the underlined part marked (5), explain in English why this is so. Write about 25 words. In your answer, do not copy whole sections of text from the passage.
- 8 Write about 20 characters in Japanese to explain what the underlined part marked (6) refers to.
- 9 Translate the underlined part marked (7) into Japanese.

- 10 Select the most suitable adjective from the list below to fill each of the spaces marked (i)～(vii). Write the number corresponding to each of your selections on the answer paper. Do not select any of the words more than once.

1 challenging 2 creative 3 doubtful 4 impressive 5 incorrect
6 insufficient 7 recent 8 reliable 9 specific 10 unknown

- 11 Select the most suitable option from A～D to fill the space marked イ. Write the letter on your answer paper. On the answer paper, the space for the answer to Q11 is on the same line as the space for Q3 and Q6.

- A that is a point that customers do not pay much attention to
B that is not something you typically want to give to your customers
C that is unlikely to affect the overall accuracy of answers generated by chatbots
D that is what users are likely to be most impressed by

- [Ⅲ] In recent years, the number of foreign tourists visiting Japan has increased significantly. Write 100 words or so in English to explain why you think Japan is such a popular destination for foreign tourists. Give at least two reasons.

NOTES

affront	a remark or an action that insults or offends somebody/something
bluff	to try to make somebody believe that you know something that you do not really know
corrosive	tending to damage something gradually
gobble up	to use something very quickly
hallucinate	(of an artificial intelligence system) to create wrong information
mea culpa	used when you are admitting that something is your fault
node	a point at which two lines or systems meet or cross
ol	the name of an AI model made by OpenAI
open-source	used to describe software for which the original source code is made available to anyone
palpable	that is easily noticed by the mind or the senses
parameter	something that decides or limits the way in which something can be done
party	one of the people or groups of people involved in a legal agreement or argument
prerogative	a right or advantage belonging to a particular person or group because of their importance or social position
recrimination	an angry statement that somebody makes accusing somebody else of something, especially in response to a similar statement from them
unequivocal	expressing your opinion or intention very clearly and strongly

Adapted from *Oxford Advanced Learner's Dictionary*.

[出典] 以下の資料に基づく

[Ⅰ] Jane E. Brody, *New York Times*, 30 January 2017. (記事のタイトルは省略)

[Ⅱ] Nicola Jones, *Nature*, 25 September 2024. (記事のタイトルは省略)