Guidelines for Human-AI Interaction

How to use these cards

You can use these cards throughout your design process as you evaluate existing ideas, brainstorm new ones, and collaborate with the multiple disciplines involved in creating AI.

Each card presents a guideline and an example that illustrates the guideline in practice. The guidelines are not rules or patterns and need not be used as a checklist. Not all may apply in every case, and in some specialized cases, such as bots or voice interaction, additional guidelines might be needed. In some cases, you will find you need to make trade-offs between guidelines.

You are using these guidelines “the right way” if you are considering them in your work to engage in dialogue about human-centered AI.

Please let us know how you are using these cards. Send feedback to aiguidelines@microsoft.com.

Learn more: https://aka.ms/aiguidelines.

1

INITIALLY

Make clear what the system can do.

Help users understand what the AI system is capable of doing.

EXAMPLE IN PRACTICE

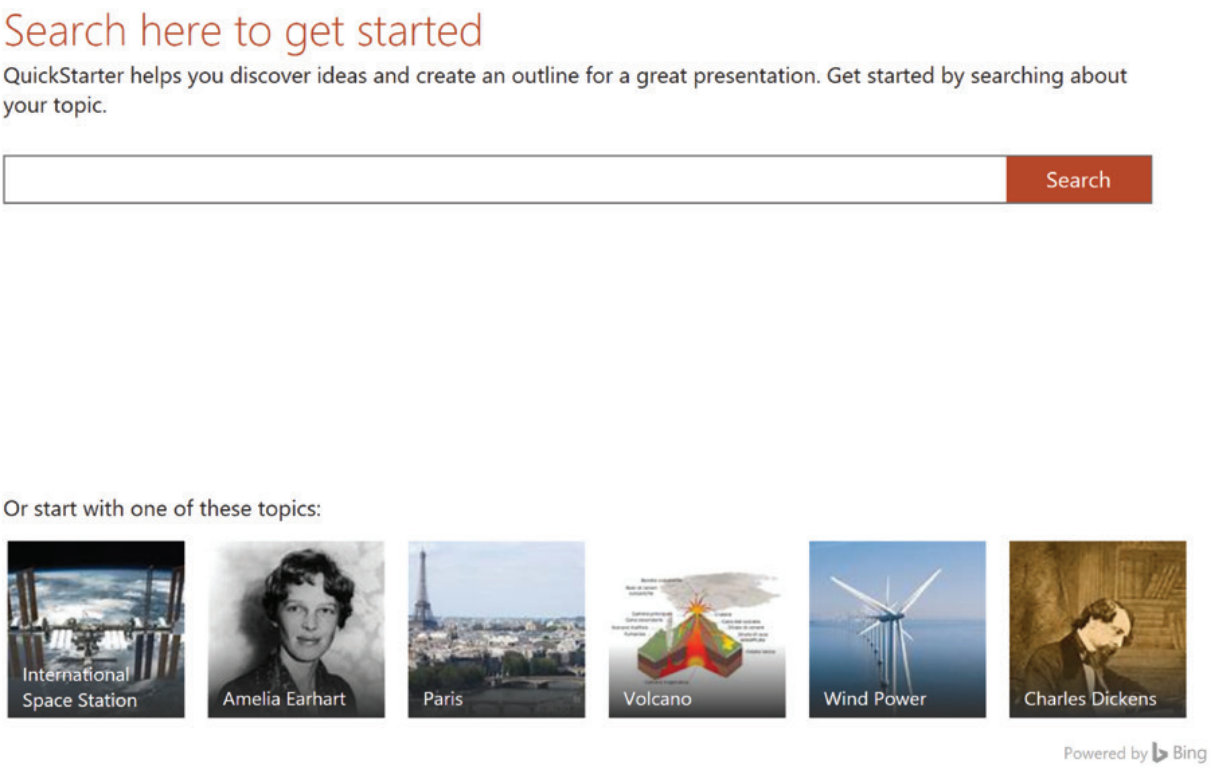

PowerPoint's **QuickStarter** builds an outline to help you get started researching a subject. It displays suggested topics that help you understand the feature's capabilities.

2

INITIALLY

Make clear how well the system can do what it can do.

Help the user understand how often the AI system may make mistakes.

EXAMPLE IN PRACTICE

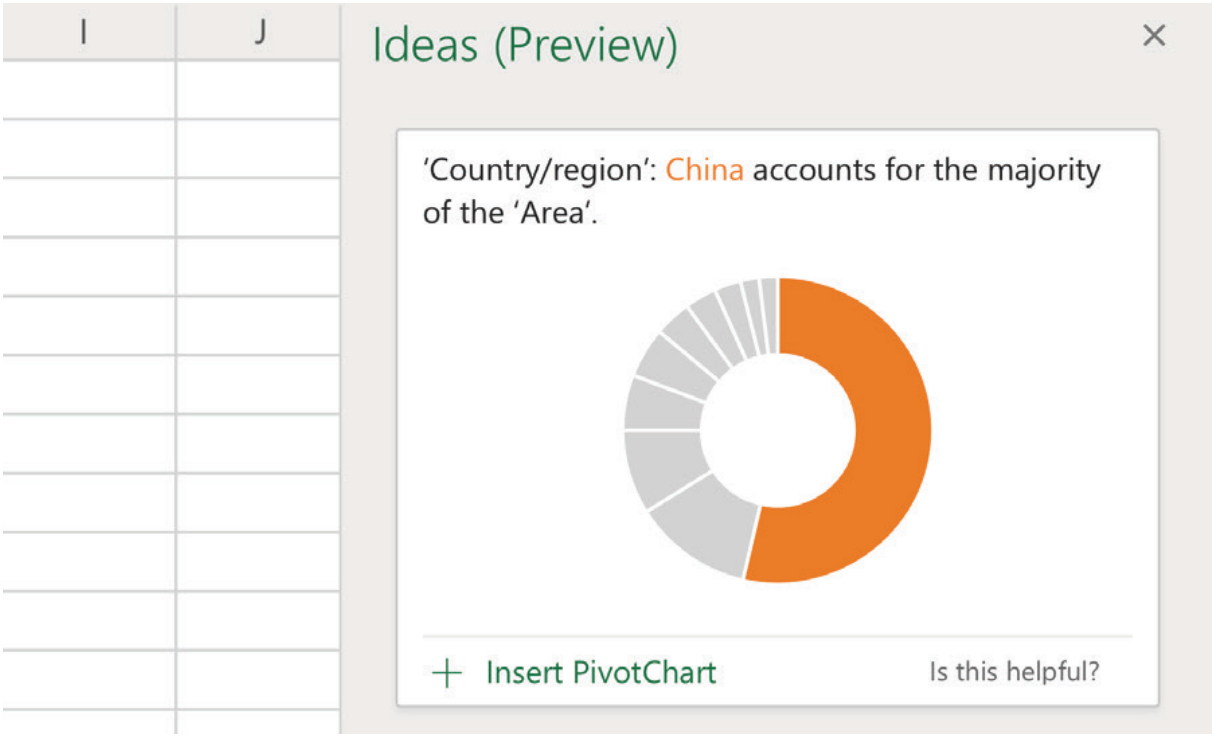

Office's new companion experience, **Ideas**, docks alongside your work and offers one-click assistance with grammar, design, data insights, rich imagery, and more. The unassuming term "Ideas," coupled with the label "Preview," helps set expectations about the presented suggestions.

Make clear how well the system can do what it can do.

3

DURING INTERACTION

Time services based on context.

Time when to act or interrupt based on users' current task and environment.

EXAMPLE IN PRACTICE

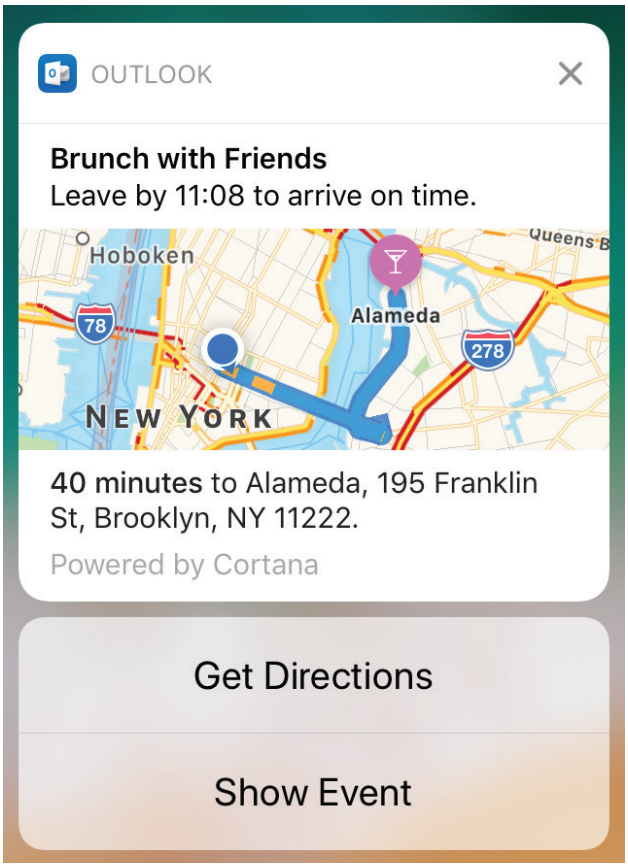

When it is time to leave for appointments, Outlook sends a **Time to Leave** notification—with directions for both driving and public transit—taking into account current location, the event location, and real-time traffic information.

4

DURING INTERACTION

Show contextually relevant information.

Display information relevant to the user's current task and environment.

EXAMPLE IN PRACTICE

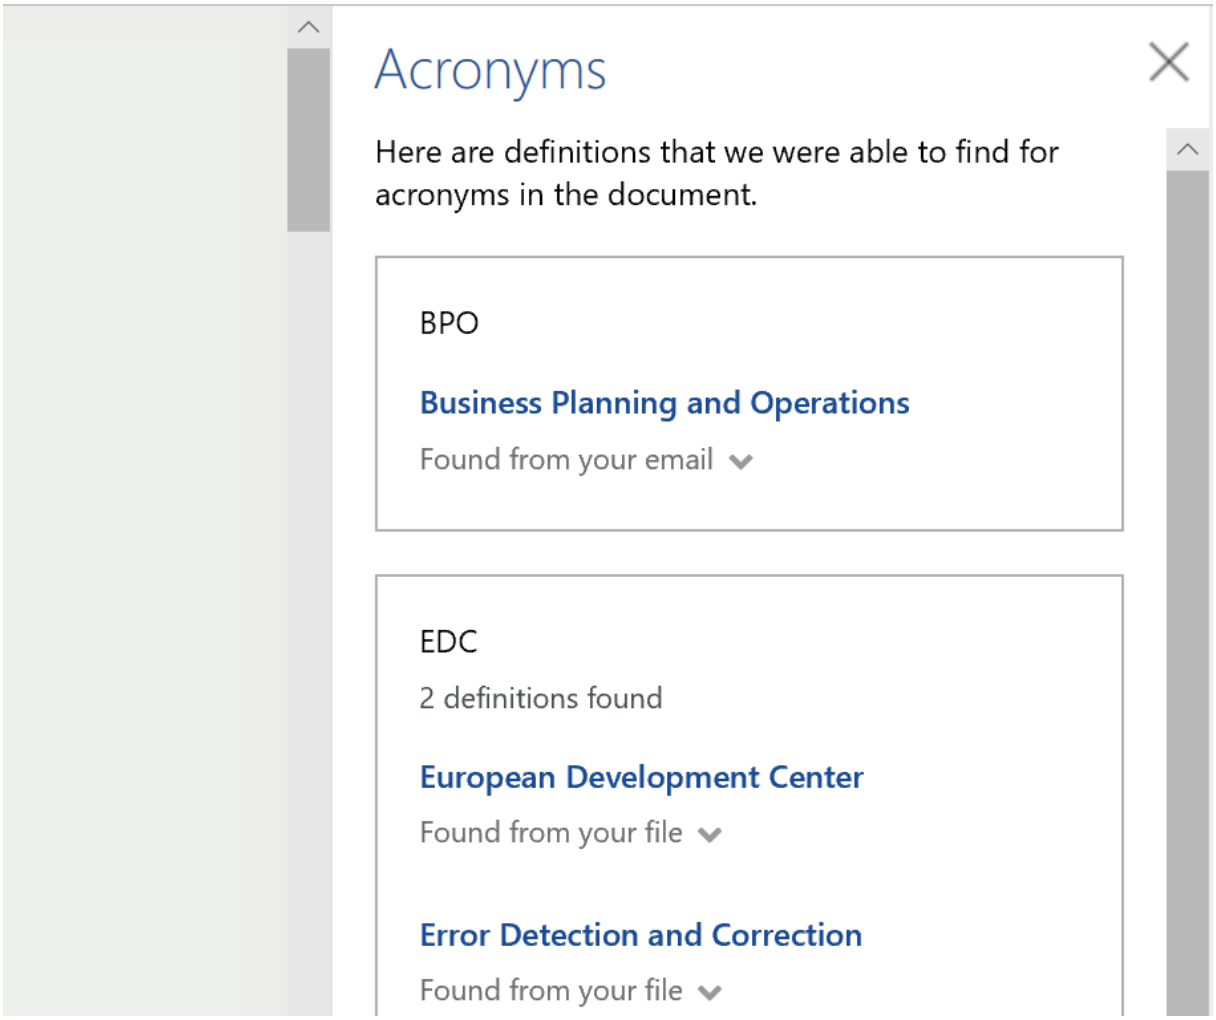

Powered by machine learning, **Acronyms** in Word helps you understand shorthand employed in your own work environment relative to the current open document.

Show contextually relevant information. **4**

5

DURING INTERACTION

Match relevant social norms.

Ensure the experience is delivered in a way that users would expect, given their social and cultural context.

EXAMPLE IN PRACTICE

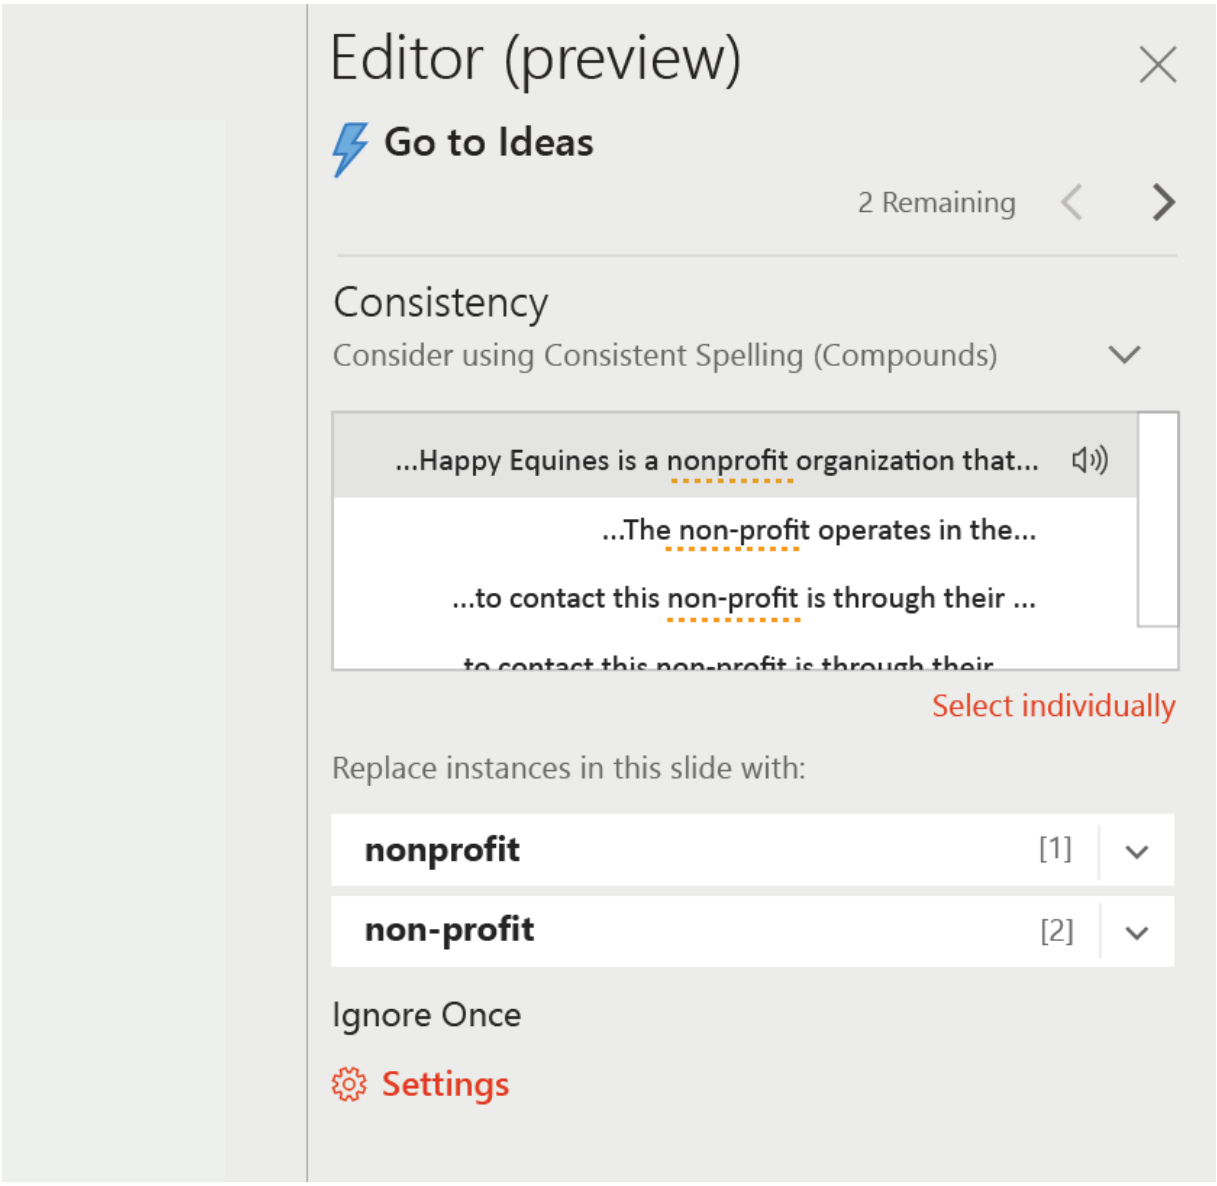

When **Editor** identifies ways to improve writing style, it presents options politely: "Consider using … "

6

DURING INTERACTION

Mitigate social biases.

Ensure the AI system's language and behaviors do not reinforce undesirable and unfair stereotypes and biases.

EXAMPLE IN PRACTICE

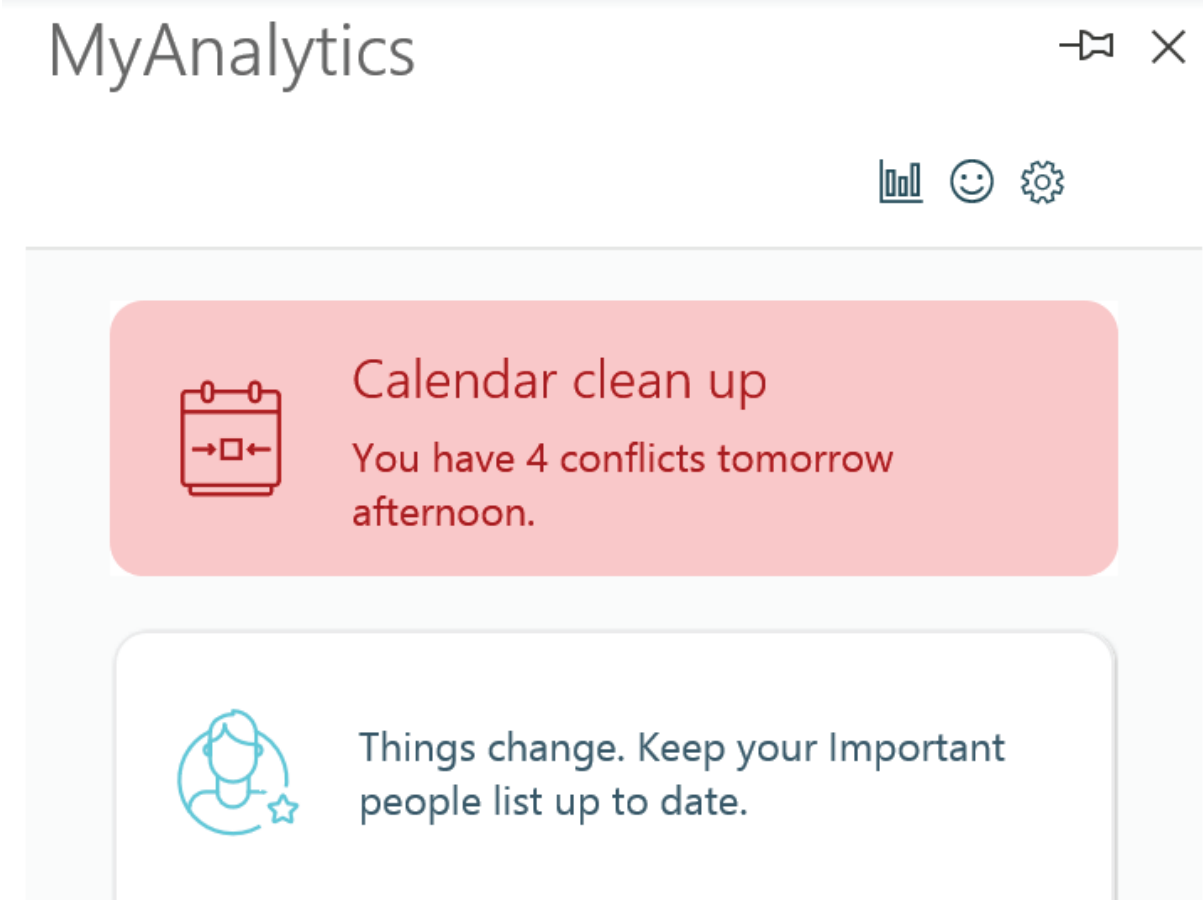

MyAnalytics summarizes how you spend your time at work, then suggests ways to work smarter. One way it mitigates bias is by using gender-neutral icons to represent important people.

7

WHEN WRONG

Support efficient invocation.

Make it easy to invoke or request the AI system's services when needed.

EXAMPLE IN PRACTICE

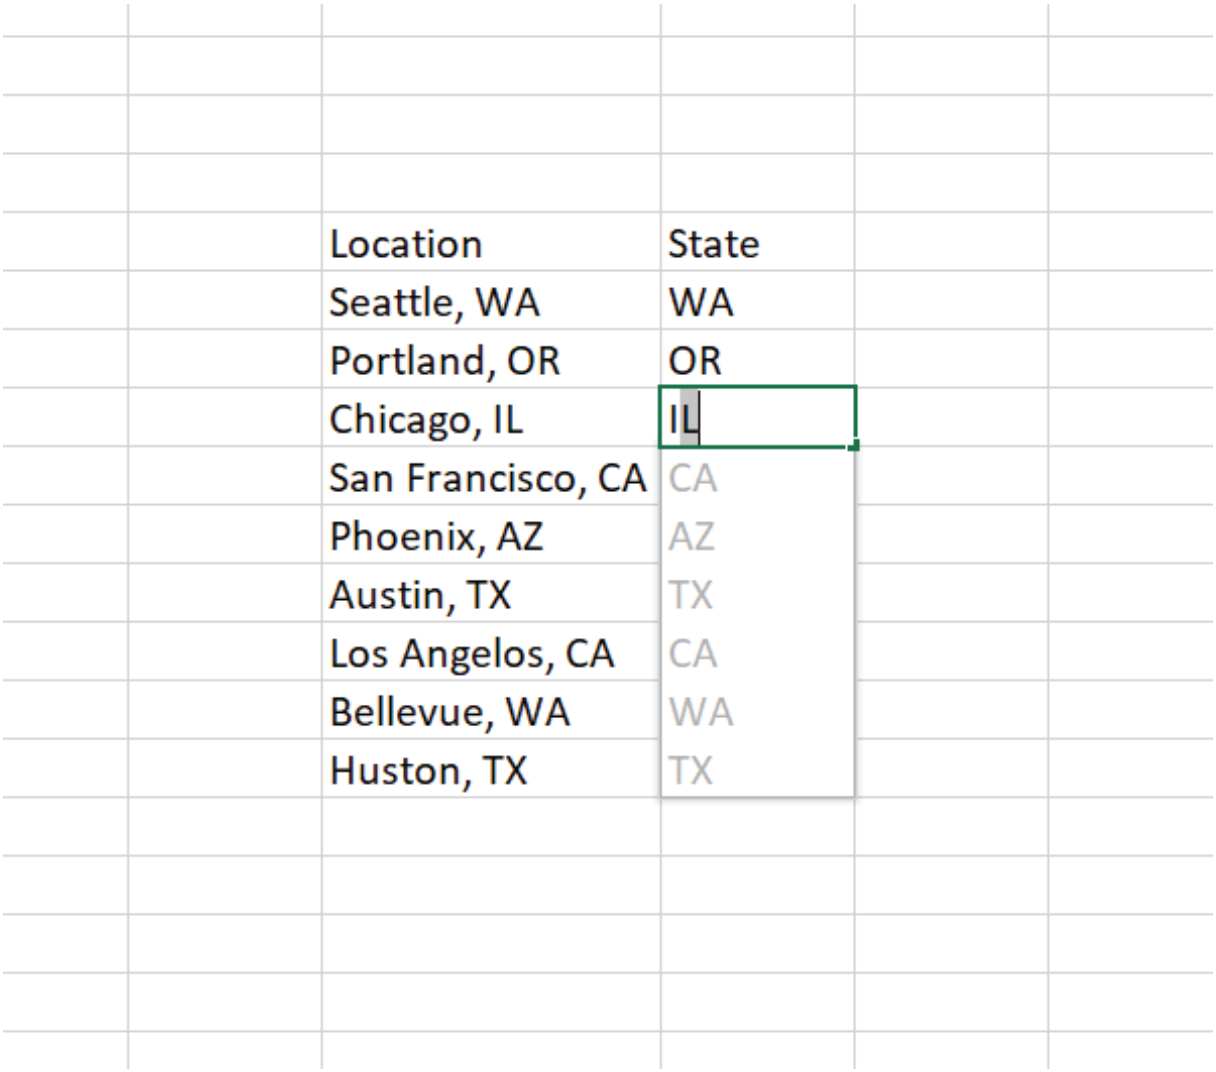

Flash Fill is a helpful time-saver in Excel that can be easily invoked with on-canvas interactions that keep you in flow.

8

WHEN WRONG

Support efficient dismissal.

Make it easy to dismiss or ignore undesired AI system services.

EXAMPLE IN PRACTICE

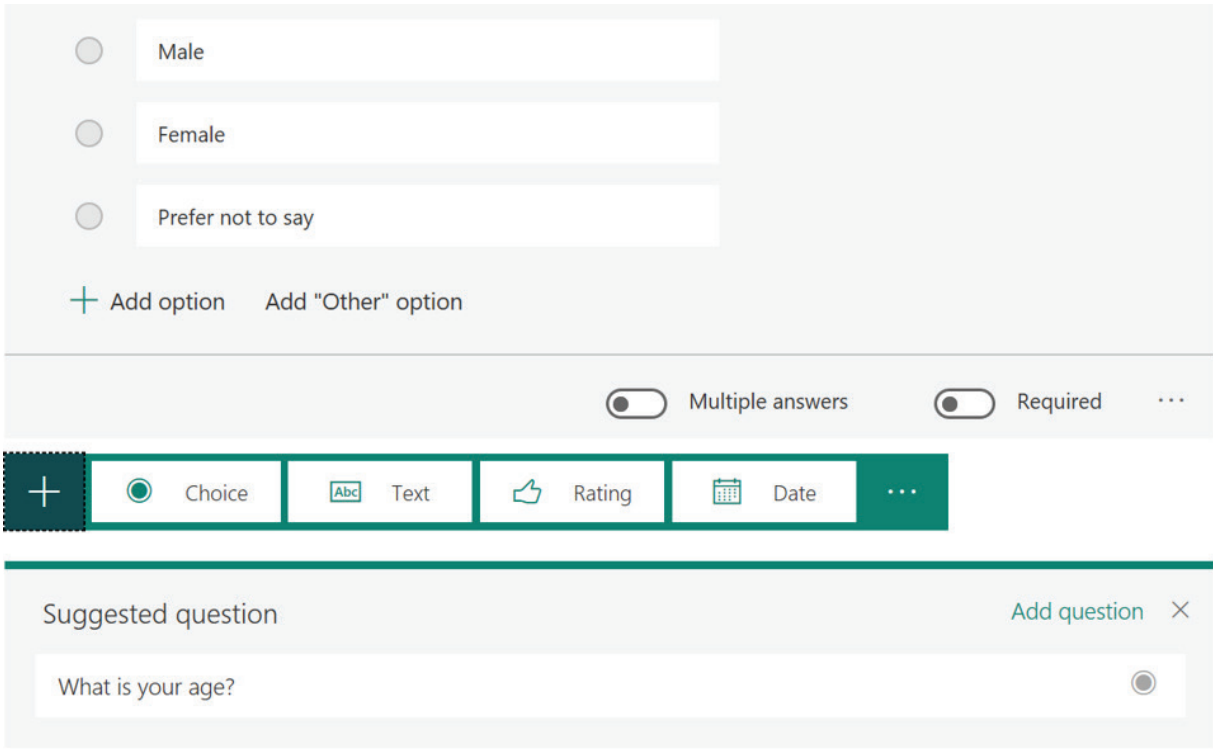

Microsoft Forms allows you to create custom surveys, quizzes, polls, and questionnaires. In Forms, some choice questions trigger **Suggested Options**. Positioned beneath the relevant question, the suggestions can be easily ignored or dismissed by clicking the "X."

Support efficient dismissal.

9

WHEN WRONG

Support efficient correction.

Make it easy to edit, refine, or recover when the AI system is wrong.

EXAMPLE IN PRACTICE

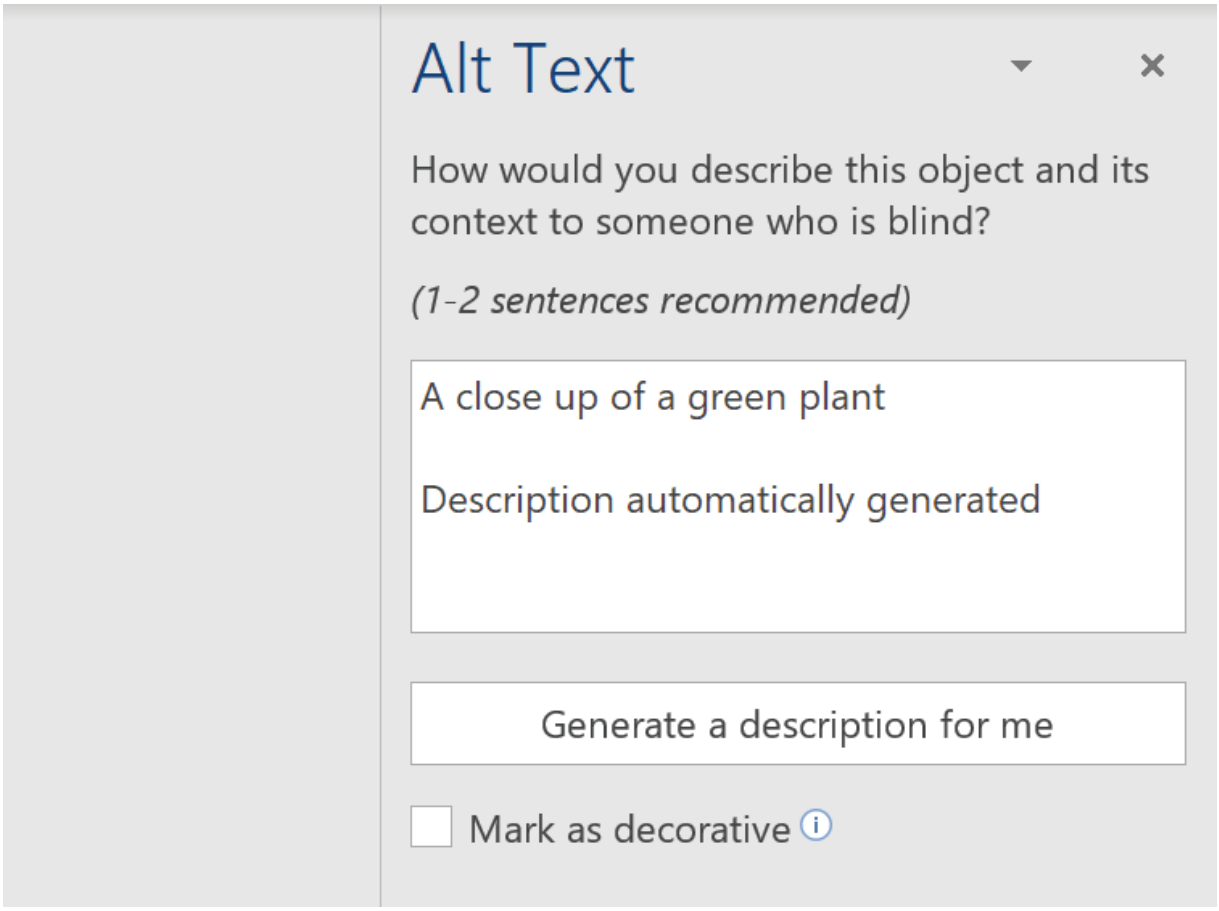

Auto Alt Text automatically generates alt text for photographs by using intelligent services in the cloud. Descriptions can be easily modified by clicking the Alt Text button in the ribbon.

10

WHEN WRONG

Scope services when in doubt.

Engage in disambiguation or gracefully degrade the AI system's services when uncertain about a user's goals.

EXAMPLE IN PRACTICE

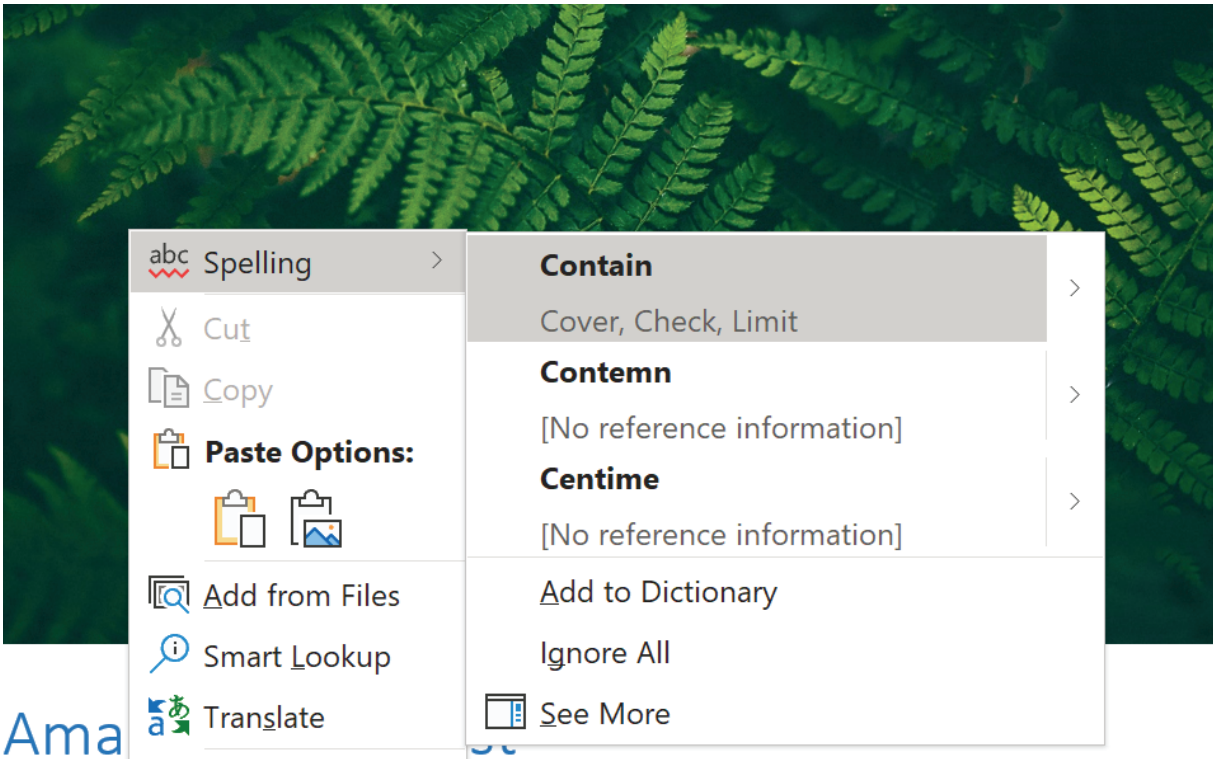

When **AutoReplace** in Word is uncertain of a correction, it engages in disambiguation by displaying multiple options you can select from.

11

WHEN WRONG

Make clear why the system did what it did.

Enable the user to access an explanation of why the AI system behaved as it did.

EXAMPLE IN PRACTICE

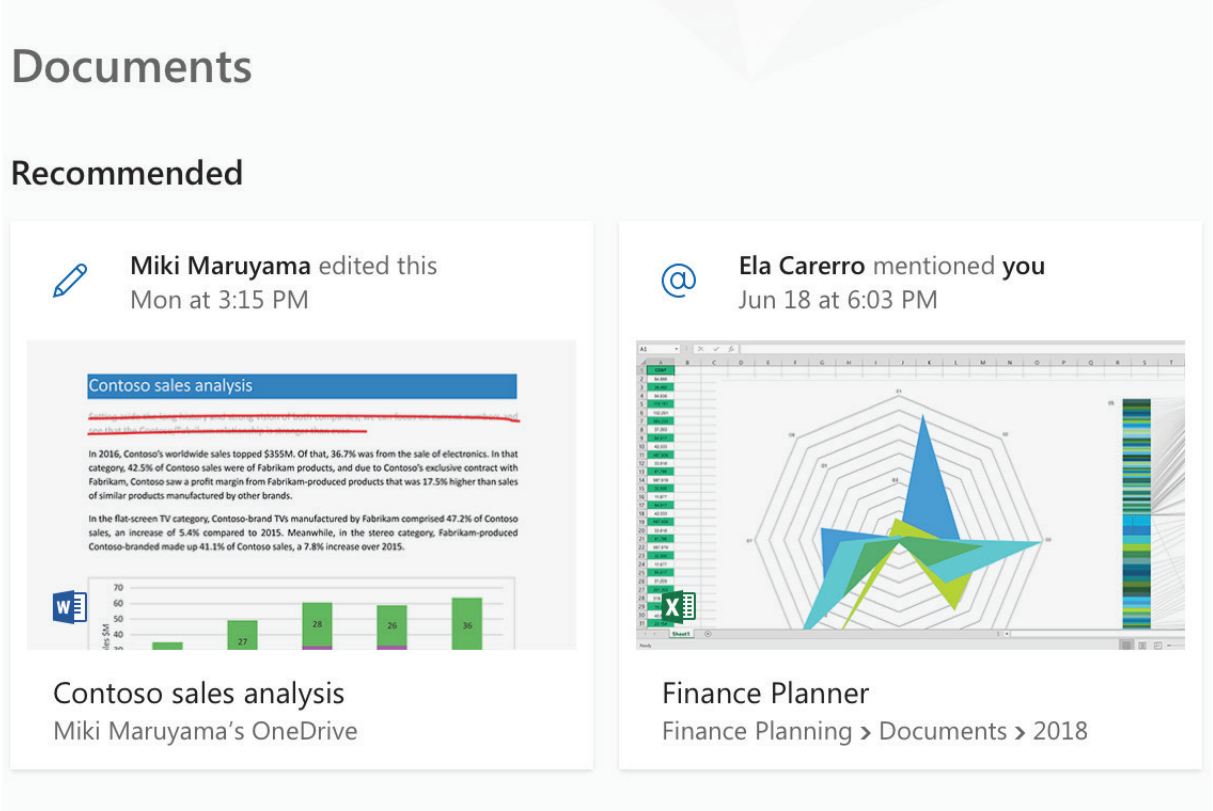

Office online **recommends documents** based on history and activity. Descriptive text above each document makes it clear why the recommendation is shown.

Make clear why the system did what it did.

12

OVER TIME

Remember recent interactions.

Maintain short-term memory and allow the user to make efficient references to that memory.

EXAMPLE IN PRACTICE

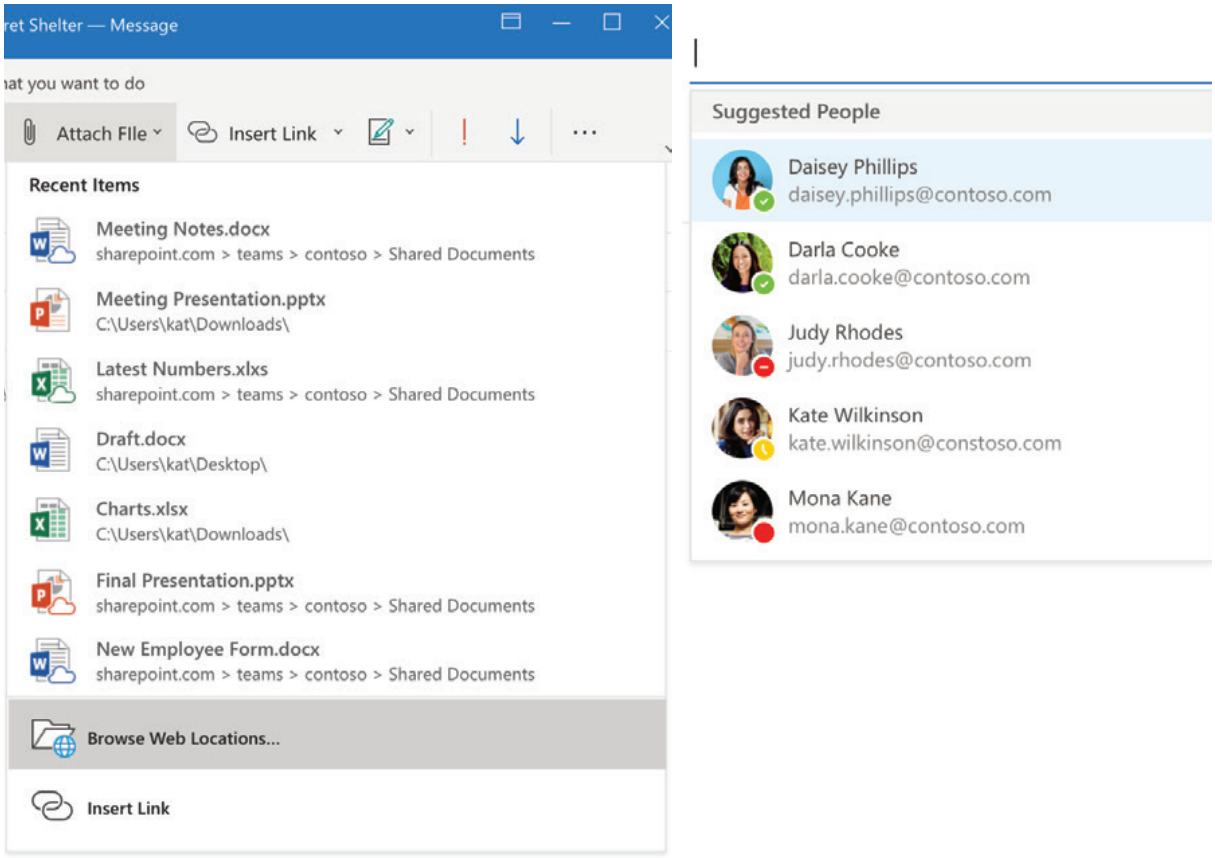

When attaching a file, Outlook offers a list of **recent files**, including recently copied file links. Outlook also remembers people you have interacted with recently and displays them when addressing a new email.

13

OVER TIME

Learn from user behavior.

Personalize the user's experience by learning from their actions over time.

EXAMPLE IN PRACTICE

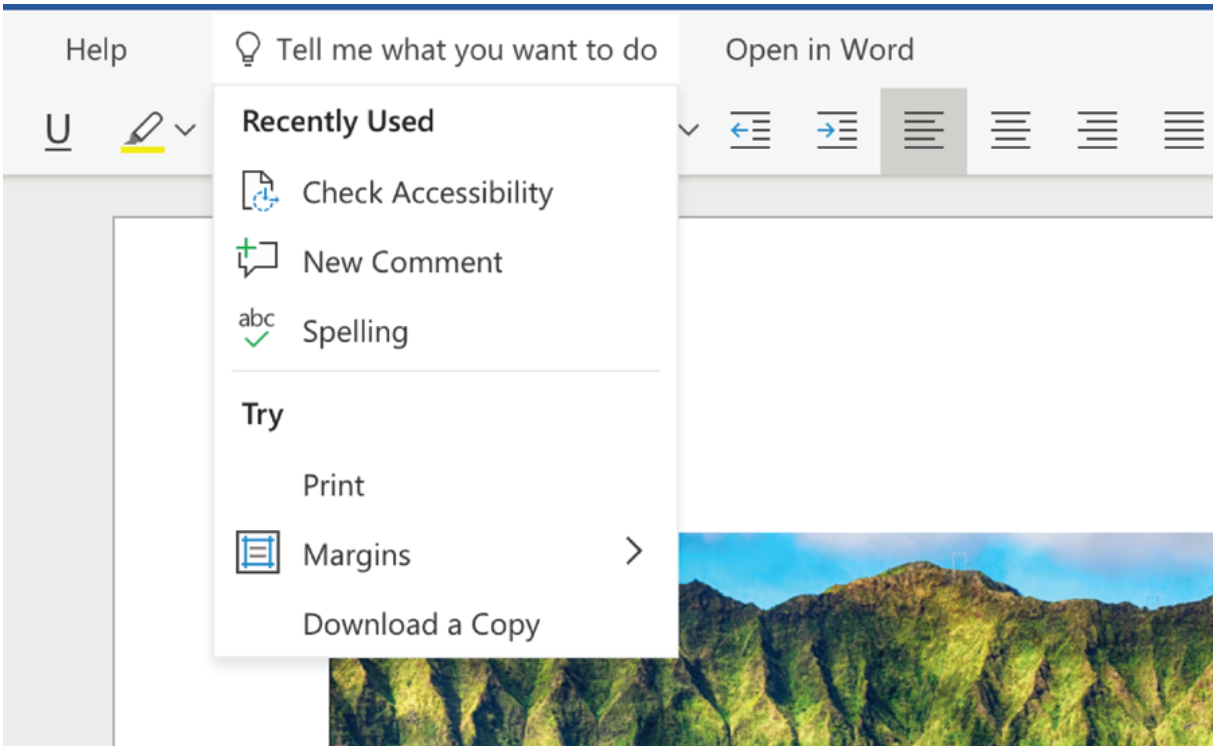

Tap on a Search bar in any Office application and Search lists the top three commands on your screen that you're most likely to need–personalized to you. The technology, called "**0-Query**," doesn't even need you to type in the Search bar to provide a personalized, predictive answer.

14

OVER TIME

Update and adapt cautiously.

Limit disruptive changes when updating and adapting the AI system's behaviors.

EXAMPLE IN PRACTICE

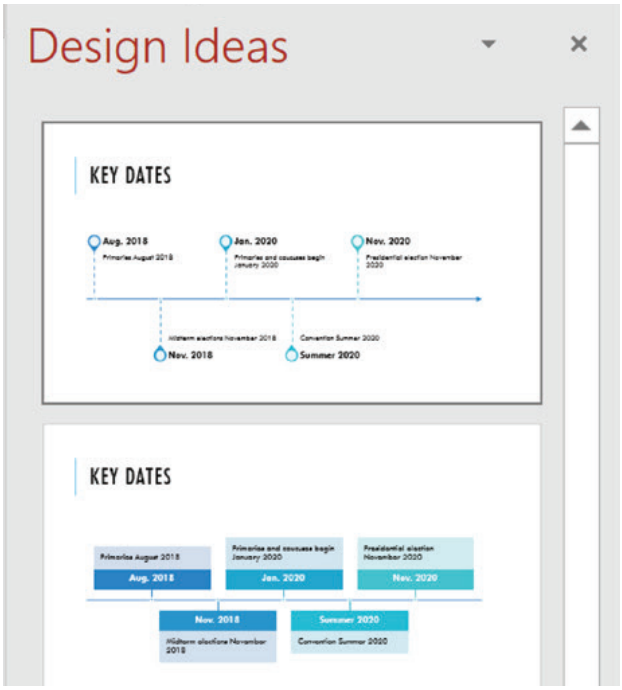

PowerPoint **Designer** improves slides for Office 365 subscribers by automatically generating design ideas to choose from. Designer has integrated new capabilities such as smart graphics and icon suggestions into the existing user experience, ensuring the updates are not disruptive.

15

OVER TIME

Encourage granular feedback.

Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.

EXAMPLE IN PRACTICE

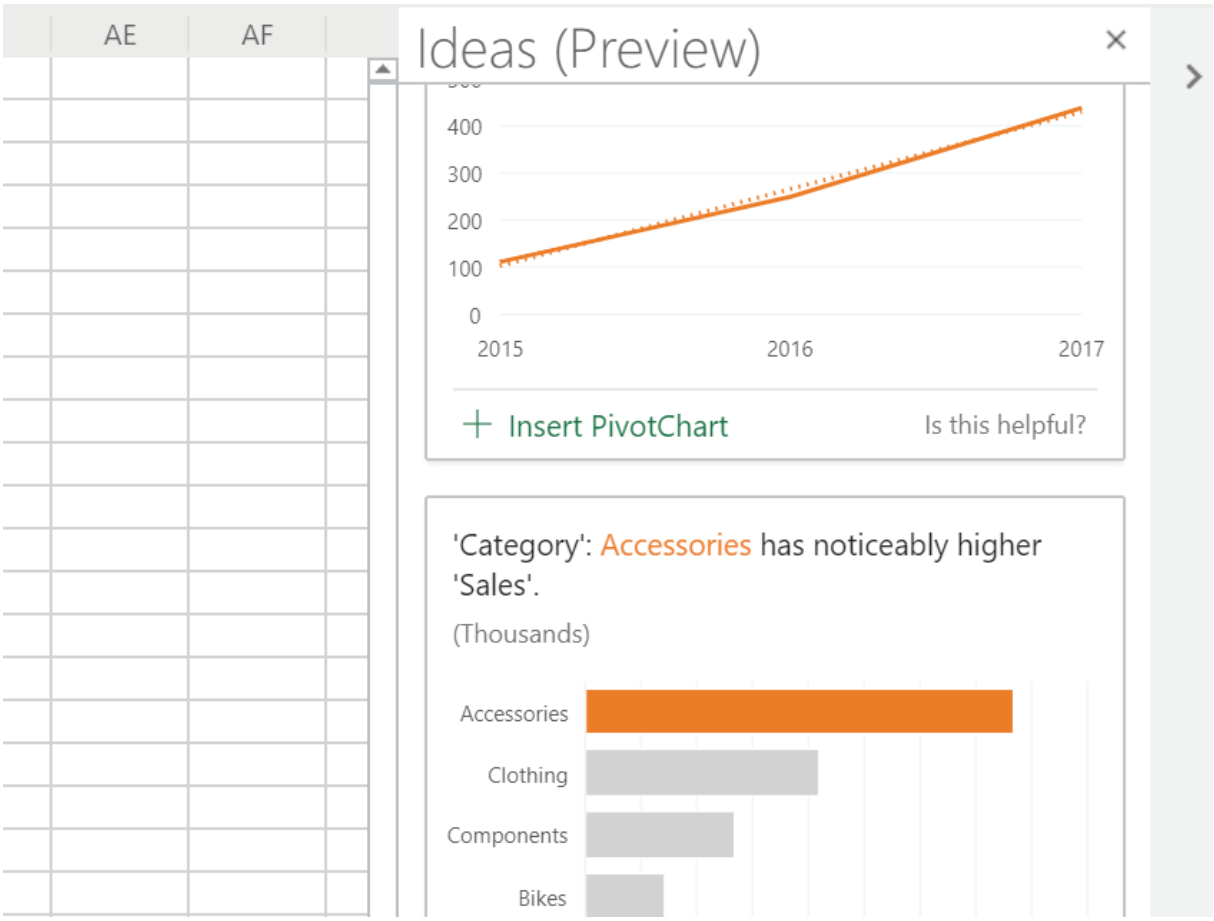

Ideas in Excel empowers you to understand your data through high-level visual summaries, trends, and patterns. It encourages feedback on each suggestion by asking, "Is this helpful?"

Encourage granular feedback.

16

OVER TIME

Convey the consequences of user actions.

Immediately update or convey how user actions will impact future behaviors of the AI system.

EXAMPLE IN PRACTICE

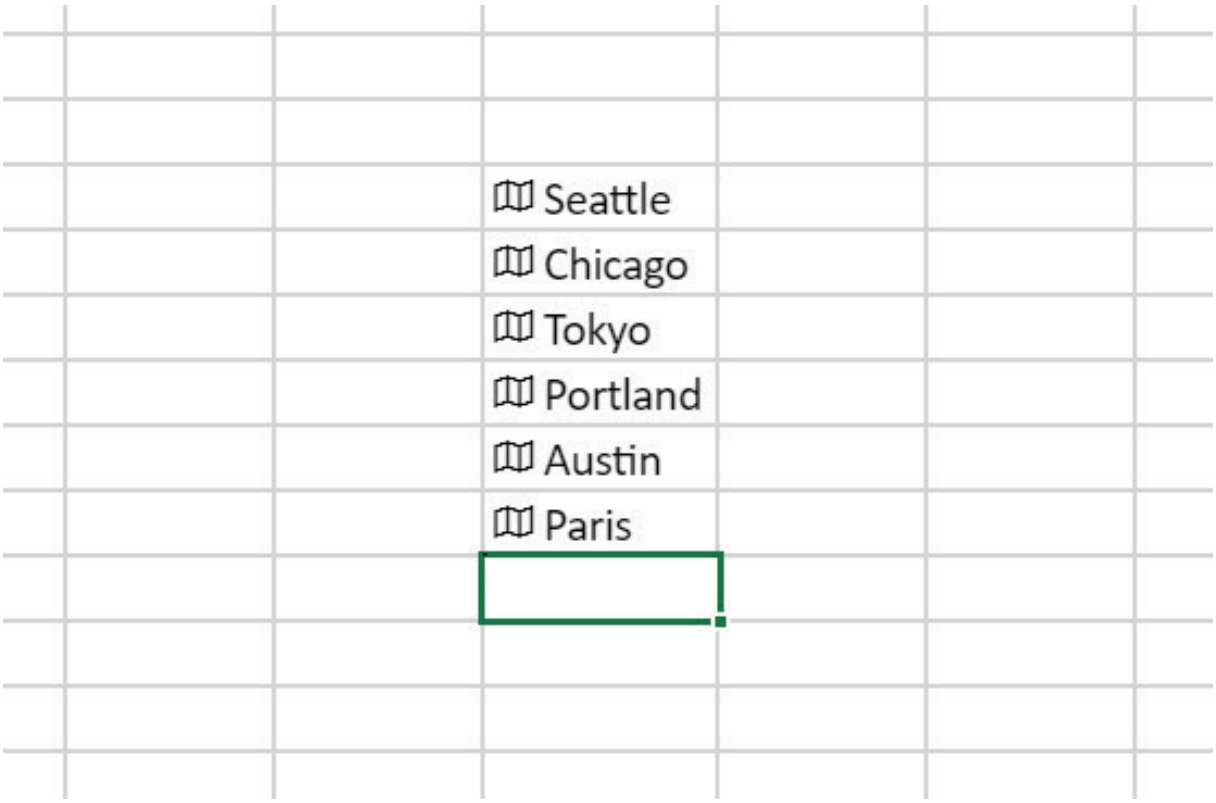

You can get stock and geographic **data types** in Excel. It's as easy as typing text into a cell and converting it to the Stocks data type or the Geography data type. When you perform the conversion action, an icon immediately appears in the converted cells.

17

OVER TIME

Provide global controls.

Allow the user to globally customize what the AI system monitors and how it behaves.

EXAMPLE IN PRACTICE

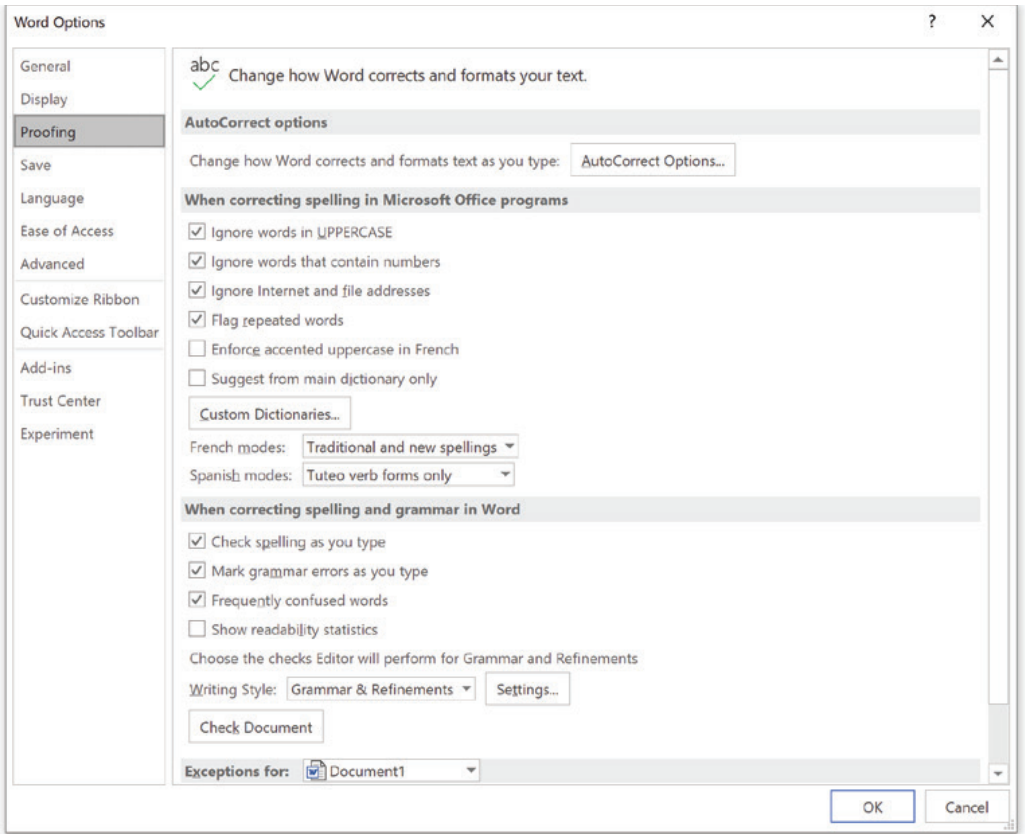

Editor expands on the spelling- and grammar-checking capabilities of Word to include more advanced proofing and editing designed to ensure your document is readable. Editor can flag a range of critique types, and allows you to customize its behavior.

18

OVER TIME

Notify users about changes.

Inform the user when the AI system adds or updates its capabilities.

EXAMPLE IN PRACTICE

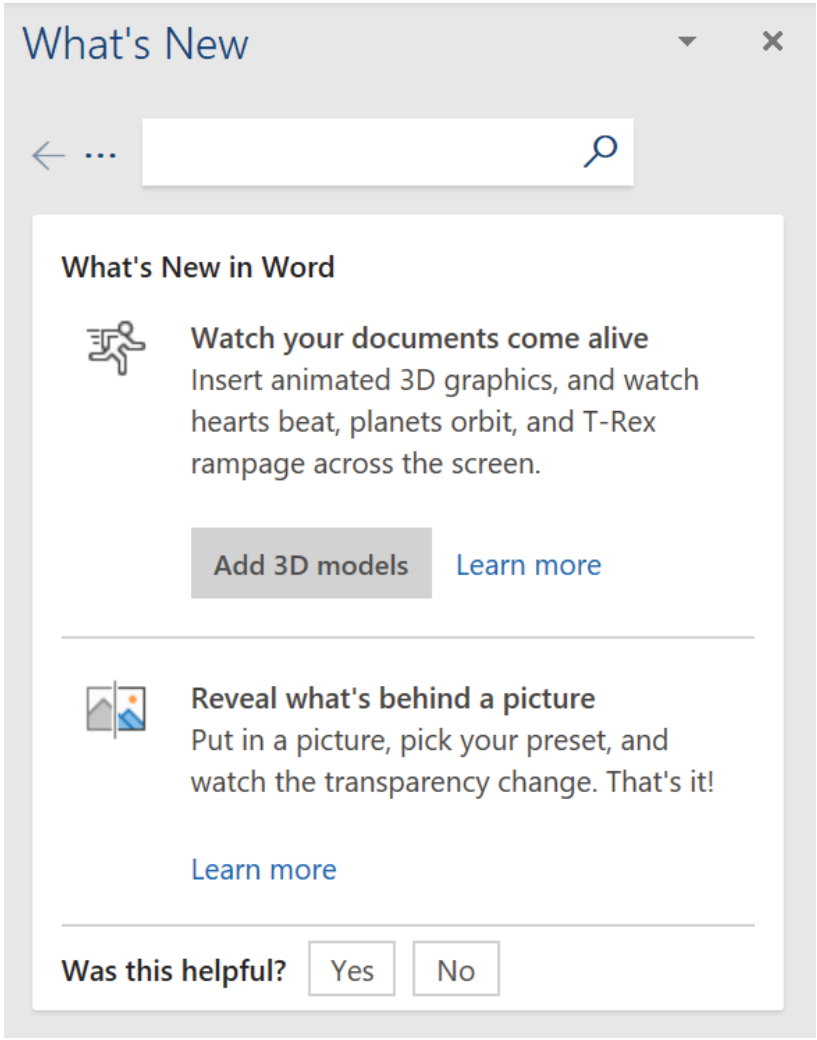

The "**What's New**" dialogue in Office informs you about changes by giving an overview of latest features and updates, including updates to AI features.

INITIALLY

1. Make clear what the system can do.
2. Make clear how well the system can do what it can do.

DURING INTERACTION

3. Time services based on context.
4. Show contextually relevant information.
5. Match relevant social norms.
6. Mitigate social biases.

WHEN WRONG

7. Support efficient invocation.
8. Support efficient dismissal.
9. Support efficient correction.
10. Scope services when in doubt.
11. Make clear why the system did what it did.

OVER TIME

12. Remember recent interactions.
13. Learn from user behavior.
14. Update and adapt cautiously.
15. Encourage granular feedback.
16. Convey the consequences of user actions.
17. Provide global controls.
18. Notify users about changes.


Microsoft

Microsoft

These guidelines are the result of a rigorous synthesis and validation process. We identified more than 150 potential design guidelines in scholarly research, documents across Microsoft, and articles in the public domain. We grouped the guidelines by theme, which resulted in a short list. We then conducted three rounds of validation:

1. We tested the guidelines ourselves, by applying them to popular products' AI features. We revised the list of guidelines, removing redundancies, confusion, and any guidelines that cannot be observed from the UI.

(continued)

2. We tested the revised list of guidelines by having 49 UX designers and researchers across the company apply them to various products. We revised the wording of some guidelines based on that feedback.
3. We asked 11 UX experts to review and validate the revisions.

The research process we used and the guidelines, along with more examples, are presented in a CHI 2019 paper available at https://aka.ms/aiguidelines.