The average 20 year old knows between 27,000 and 52,000 different words.
By age 60, that number averages between 35,000 and 56,000.
Spoken out loud, most of these words last less than a second.
So with every word, the brain has a quick decision to make:
which of those thousands of options matches the signal?
About 98% of the time, the brain chooses the correct word.
But how?

Speech comprehension is different from reading comprehension, but it's similar to sign language comprehension— though spoken word recognition has been studied more than sign language.
The key to our ability to understand speech is the brain's role as a parallel processor, meaning that it can do multiple different things at the same time.
Most theories assume that each word we know is represented by a separate processing unit that has just one job:

to assess the likelihood of incoming speech matching that particular word.
In the context of the brain, the processing unit that represents a word is likely a pattern of firing activity across a group of neurons in the brain's cortex. When we hear the beginning of a word, several thousand such units may become active, because with just the beginning of a word, there are many possible matches. Then, as the word goes on, more and more units register that some vital piece of information is missing and lose activity.
Possibly well before the end of the word, just one firing pattern remains active, corresponding to one word. This is called the "recognition point." In the process of honing in on one word, the active units suppress the activity of others, saving vital milliseconds. Most people can comprehend up to about 8 syllables per second. Yet, the goal is not only to recognize the word, but also, to access its stored meaning.

The brain accesses many possible meanings at the same time, before the word has been fully identified.
We know this from studies which show that even upon hearing a word fragment— like "cap"— listeners will start to register multiple possible meanings, like captain or capital, before the full word emerges. This suggests that every time we hear a word there's a brief explosion of meanings in our minds, and by the recognition point the brain has settled on one interpretation.

The recognition process moves more rapidly with a sentence that gives us context than in a random string of words.
Context also helps guide us towards the intended meaning of words with multiple interpretations, like "bat," or "crane," or in cases of homophones like "no" or "know." For multilingual people, the language they are listening to is another cue, used to eliminate potential words that don’t match the language context. So, what about adding completely new words to this system?

Even as adults, we may come across a new word every few days. But if every word is represented as a fine-tuned pattern of activity distributed over many neurons, how do we prevent new words from overwriting old ones?
We think that to avoid this problem, new words are initially stored in a part of the brain called the hippocampus, well away from the main store of words in the cortex, so they don’t share neurons with others words. Then, over multiple nights of sleep, the new words gradually transfer over and interweave with old ones. Researchers think this gradual acquisition process helps avoid disrupting existing words. So, in the daytime, unconscious activity generates explosions of meaning as we chat away. At night, we rest, but our brains are busy integrating new knowledge into the word network.
When we wake up, this process ensures that we’re ready for the ever-changing world of language.

English