CLAIMS

What is claimed is:

1. A mobile computing device, comprising:

a touchscreen;

a processor coupled to the touchscreen;

memory coupled to the processor and storing one or more sequences of instructions which, when executed by the processor, cause the processor to perform:

displaying, on the touchscreen, a graphical user interface that includes a virtual keyboard, an input area, and an output area;

receiving, from the touchscreen, first signal data indicative of a location in the input area of a target word within a text sequence;

in response to the first signal data, displaying, in the output area, a subset of candidate synonyms for the target word, the subset of candidate synonyms being created by

inputting the text sequence including the target word into a dependency parser computer program stored on the mobile computing device and receiving, from the dependency parser computer program, output comprising syntactic dependency data for the target word including data indicating a syntactic dependency relationship between the target word and at least one other word of the text sequence;

retrieving a set of candidate synonyms for the target word from a search engine;

inputting the syntactic dependency data into a neural network-based model that has been trained to cause the neural network-based model to output a dependency context vector indicating a semantic meaning of the syntactic dependency data;

for each candidate synonym of the set of candidate synonyms, using a target word vector indicating a semantic meaning of the target word, a candidate word vector indicating a semantic meaning of the candidate synonym, and the dependency context vector, computing a similarity score indicating a semantic similarity of the candidate synonym to the target word;

when the similarity score for a candidate synonym satisfies a criterion, adding the candidate synonym to the subset of candidate synonyms;

receiving, from the touchscreen, second signal data indicating a location, in the output area, of a synonym selected from the subset of candidate synonyms;

using the first signal data, replacing the target word in the input area with the selected synonym.

2. A computer-implemented method comprising:

determining and storing in memory of a mobile computing device, digital data including a target word of a text sequence, a set of candidate synonyms for the target word, a target word vector indicating a semantic meaning of the target word, and, for each candidate synonym of the set of candidate synonyms, a candidate word vector indicating a semantic meaning of the candidate synonym;

inputting the text sequence including the target word into a dependency parser computer program stored on the mobile computing device and receiving, from the dependency parser computer program, output comprising syntactic dependency data for the target word including data indicating a syntactic dependency relationship between the target word and at least one other word of the text sequence;

inputting the syntactic dependency data into a neural network-based model that has been trained to cause the neural network-based model to output a dependency context vector indicating a semantic meaning of the syntactic dependency data;

retrieving the set of candidate synonyms for the target word;

for each candidate synonym of the set of candidate synonyms, using the target word vector, the candidate word vector, and the dependency context vector, computing a similarity score indicating a semantic similarity of the candidate synonym to the target word;

when the similarity score for a candidate synonym satisfies a criterion, adding the candidate synonym to a subset of the set of candidate synonyms;

in response to determining the target word and determining the subset of the set of candidate synonyms satisfies a confidence level value, causing a virtual keyboard software application of the mobile computing device to switch to a synonym mode and programmatically providing the subset of the set of candidate synonyms to the virtual keyboard software application for use by the synonym mode.

3. The method of claim 2, further comprising using cursor position data to determine the target word, the cursor position data obtained from an input area of a graphical user interface of the virtual keyboard software application, the graphical user interface rendered on a display screen of the mobile computing device, the virtual keyboard software application stored in memory of the mobile computing device.

4. The method of claim 3, the cursor position data indicating a time interval during which a position of a cursor in the input area remained unchanged.

5. The method of claim 3, the cursor position data indicating a user selection detected by a touchscreen of the mobile computing device.

6. The method of claim 3, further comprising causing the mobile computing device to display the subset of the set of candidate synonyms in an output area adjacent the input area of the graphical user interface.

7. The method of claim 2, further comprising determining the target word vector by inputting the target word into a second neural network-based model trained to determine a semantic representation of a word, and determining output produced by the second neural network-based model in response to the inputting of the target word into the second neural network-based model.

8. The method of claim 7, further comprising determining the candidate word vector by inputting the candidate synonym into a second neural network-based model and determining output produced by the second neural network-based model in response to the inputting of the candidate word vector into the second neural network-based model.

9. The method of claim 7, the target word vector comprising information about the target word, the information indicating one or more of information about a meaning of the target word, a position of the target word in the text sequence or a word that typically precedes or follows the target word.

10. The method of claim 2, the text sequence including at least three tokens between the target word and the at least one other word of the text sequence.

11. The method of claim 2, further comprising, for each candidate synonym of the set of candidate synonyms, computing the similarity score by computing a dot product of the target word vector and the candidate word vector, computing a dot product of the dependency context vector and the candidate word vector, and adding the dot product of the target word vector and the candidate word vector to the dot product of the dependency context vector and the candidate word vector.

12. The method of claim 2, further comprising determining that the similarity score satisfies a criterion by sorting similarity scores in descending order and adding a candidate synonym that has a top k-ranked similarity score to the subset of the set of candidate synonyms, where k is a positive integer less than or equal to 5.

13. The method of claim 2, further comprising computing a hash of the dependency context vector and storing the hash of the dependency context vector in memory of the mobile computing device.

14. A method, comprising:

displaying, on a touchscreen of a mobile computing device, a graphical user interface comprising a virtual keyboard, an input area, and an output area;

receiving, from the touchscreen, first signal data indicating a location in the input area of a target word;

in response to the first signal data, displaying, in the output area, a subset of candidate synonyms for the target word;

the subset of candidate synonyms being created by determining syntactic dependency data for the target word including data indicating a syntactic dependency relationship between the target word and at least one other word of the text sequence, inputting the syntactic dependency data into a neural network-based model trained to cause the neural network-based model to output a dependency context vector indicating a semantic meaning of the syntactic dependency data, retrieving a set of candidate synonyms for the target word, and using the dependency context vector to compute a similarity score indicating a semantic similarity between a candidate synonym and the target word;

receiving, from the touchscreen, second signal data indicating a location, in the output area, of a selected synonym;

using the first signal data, replacing the target word in the input area with the selected synonym.

15. The method of claim 14, the first signal data indicating a time interval during which a position of a cursor in the input area remained unchanged.

16. The method of claim 14, the first signal data indicating a user selection detected by a touchscreen of the mobile computing device.

17. A mobile computing device, comprising:

a touchscreen;

a processor coupled to the touchscreen;

memory coupled to the processor and storing one or more sequences of instructions which, when executed by the processor, cause the processor to perform:

displaying, on the touchscreen, a graphical user interface that includes a virtual keyboard, an input area, and an output area;

receiving, from the touchscreen, first signal data indicating a location in the input area of a target word within a text sequence;

in response to the first signal data, displaying, in the output area, a subset of candidate synonyms for the target word;

the subset of candidate synonyms being created by determining syntactic dependency data for the target word including data indicating a syntactic dependency relationship between the target word and at least one other word of the text sequence, inputting the syntactic dependency data into a neural network-based model trained to cause the neural network-based model to output a dependency context vector indicating a semantic meaning of the syntactic dependency data, and using the dependency context vector to compute a similarity score indicating a semantic similarity between a candidate synonym and the target word;

the syntactic dependency data produced, by a dependency parser computer program stored in the memory, in response to providing the text sequence including the target word as input to the dependency parser computer program;

receiving, from the touchscreen, second signal data indicating a location, in the output area, of a synonym selected from the subset of candidate synonyms;

using the first signal data, replacing the target word in the input area with the selected synonym.

18. The mobile computing device of claim 17, the virtual keyboard having at least two modes including a synonym mode, the memory coupled to the processor and storing instructions, that when executed by the processor, further cause, in response to the first signal data, the virtual keyboard to switch from another mode to the synonym mode.

19. The mobile computing device of claim 17, the memory coupled to the processor and storing instructions, that when executed by the processor, further cause computing a hash of

the dependency context vector and storing the hash of the dependency context vector in memory of the mobile computing device.

20. The mobile computing device of claim 17, the target word vector comprising information about the target word, the information indicating one or more of information about a meaning of the target word, a position of the target word in the text sequence or a word that typically precedes or follows the target word.