

Foundations of Information Systems

Winter Semester 2025–26, Exercise 4

For discussion on Wednesday, November 12, 2025

1. Describe an algorithm which takes a string containing upper and lower case characters encoded in ASCII and outputs the same string in all lower case. (For simplicity, assume that your input string does not contain any symbols other than characters of the Latin alphabet. You may describe your algorithm in pseudo-code or in Python.)
2. How many characters can we encode with a mixed 8/16 bit representation if we insist that an 8 bit character always has 0 in the first bit, and a 16 bit character always has 1 in the first bit of each of its bytes?
3. List at least three advantages of the UTF-8 encoding over the UTF-16 encoding of Unicode characters.
4. Describe an algorithm that checks whether a UTF-8 encoded string contains any characters that are not part of the ASCII character set.