UNIVERSITY OF OSLO

Faculty of Mathematics and Natural Sciences

Exam in **INF5820 Language technological applications**
Day of exam: **10 June 2015**
Exam hours: **1430-1830**
This examination paper consists of 3 pages including this one
Appendices: **0**
Permitted materials: **None**

Make sure that your copy of this examination paper is complete before answering.

- **You may answer in English, Norwegian, Danish or Swedish.**
- **You should answer all questions. The weight of the various questions are indicated.**
- **You should read through the whole set to see whether anything is unclear so that you can ask your questions to the teachers when they arrive.**
- **If you think some assumptions are missing, make your own and explain them!**

1 The difficult task of Machine Translation (20%)

(a) Machien Translation (MT) is considered a difficult task, and after 65 years of efforts, it is still far from the quality of translations produced by professionally trained human translators. One reason for this is claimed to be typological differences between languages. Give examples of such differences and explain why they are challenging for MT.

(b) Give other reasons for why MT is dfficult and illustrate with examples.

2 The noisy channel model and the IBM models (30%)

(a) One of the underpinnings of statistical MT is the so-called "noisy channel model". The goal for translation is to find the most likely sequence of words E in the target language given a sentence F in the source language, i.e., to find the E which maximizes $P(E|F)$. Explain shortly how the noisy channel model transforms this into several other probabilities.

(b) Consider the so-called IBM1 model for word-based statistical MT. In this model $P(\mathbf{f} \mid \mathbf{e})$ is calculated by the sum:

(1) $$P(\mathbf{f} \mid \mathbf{e}) = \sum_{\mathbf{a}} P(\mathbf{f}, \mathbf{a}|\mathbf{e})$$

Furthermore, each addend is calculated by:

(2) $$P(\mathbf{f}, \mathbf{a}|\mathbf{e}) = \frac{\epsilon}{(k+1)^m} \prod_{j=1}^{m} t(f_j|e_{a_j})$$

Explain what t is in this formula. What is k and what is m?

(c) Without going into the mathematical formulas, explain how the IBM3 model differs from the IBM1 model.

3 Dialogue management (40%)

Imagine you want to integrate a speech interface to a vending machine that can sell three things: biscuits, potato chips, and water bottles. A speech recogniser is used to process the speech signal of the users and generate recognition hypotheses (with confidence scores) . After speech recognition, the dialogue manager decides on how to react (you may assume that the dialogue manager directly select the system responses, without going through a generation module). To keep things simple, you decide to design this dialogue manager using a **finite-state automaton**.

Here are the requirements for your speech interface:

- Upon detecting a new customer, the machine should greet the user and ask whether she/he wants to order something.
- If the user inputs are uncertain (e.g. if the confidence score of the recognition hypothesis is lower than 0.7), the machine should utter a clarification request to make sure the order is correctly understood.
- Once the order is understood, the machine should check whether the request item is available in the machine. If yes, it should deliver the item together with a verbal feedback ("Ok, here is your ..."). Otherwise, it should reply that the item is no more available.
- After delivering an item, the machine should ask whether the customer wants to order anything else. If yes, the same script is repeated. Otherwise, the machine should thank the user and say goodbye.

Questions:

1. Draw the finite-state automaton corresponding to these requirements. For each edge, specify the condition attached to it. For each node, specify its symbol and associated system actions (if any). You may use empty transitions if you want. (25 points)
2. Provide the formal definition of your finite-state automaton as a tuple $\langle S, \Sigma, \delta, s_0, F \rangle$. (7 points)
3. What are the advantages and disadvantages of finite-state automata for this application, compared to other dialogue management approaches? Name at least one advantage and one disadvantage. (8 points)

4 Speech recognition (10%)

What is language model adaptation in speech recognition? Explain how language model adaptation could help improve the speech recognition accuracy in the dialogue domain described in the previous question.