

Assignment 3

Due date: 17:00 on Thursday, December 4, 2024.

Late assignments will not be accepted without a valid medical certificate or other documentation of an emergency.

For CSC485 students, this assignment is worth 30% of your final grade.

For CSC2501 students, this assignment is worth 25% of your final grade.

- Read the whole assignment carefully.
- What you turn in must be your own work. You may not work with anyone else on any of the problems in this assignment. If you need assistance, contact the instructor or TA for the assignment.
- Any clarifications to the problems will be posted on the Piazza forum for the class. You will be responsible for taking into account in your solutions any information that is posted there, or discussed in class, so you should check the page regularly between now and the due date.
- The starter code directory for this assignment is distributed via MarkUs. In this handout, code files we refer to are located in that directory.
- When implementing code, make sure to read the docstrings as some of them provide important instructions, implementation details, or hints.
- Fill in your name, student number, and UTORid on the relevant lines at the top of each file that you submit. (Do not add new lines; just replace the NAME, NUMBER, and UTORid placeholders.)

Overview: Symbolic Machine Translation

In this assignment, you will learn how to write phrase structure grammars for some different linguistic phenomena in two different languages: English and Chinese. You can use the two grammars to create an interlingual machine translation system by parsing in one, and generating in the other. Don't panic if you don't speak Chinese, and also don't cheer up yet if you can speak the language — it won't give you much of an advantage over other students. A facility with languages in general will help you, as will the ability to learn and understand the nuances between the grammars of two different languages. In particular, you will start by working on agreement. Then, you will need to analyse the quantifier scoping difference between the two languages.

TRALE Instructions The instructions to setup and launch TRALE on teach.cs can be found in: <https://www.cs.toronto.edu/~niu/csc485/trale/setup/>

1. Agreement: Determiners, Numbers and Classifiers [10 marks]

English expresses subject–verb agreement in person and number. English has two kinds of number: singular and plural. The subject of a clause must agree with its predicate: they should be both singular or both plural. However, the number of a direct object does not need to agree with anything.

- (1) A programmer annoys a dolphin.
- (2) Two programmers annoy a dolphin.
- (3) * Two programmers annoys two dolphins.
- (4) * A programmer annoy two dolphins.

Chinese, on the other hand, does not exhibit subject–verb agreement. As shown in the examples below, most nouns do not inflect at all for plurality. Chinese does, however, have a classifier (CL) part of speech that English does not. Semantically, classifiers are similar to English collective nouns (a bottle of water, a murder of crows), but English collective nouns are only used when describing collectives. With very few exceptions, classifiers are mandatory in complex Chinese noun phrases. Different CLs agree with different classes of nouns that are sorted by mostly semantic criteria. For example, 学生 (xuesheng)¹ student is a person and an occupation, so it should be classified by either 个 (ge) or 名 (ming) and cannot be classified by the animal CL 只 (zhi). However, the rules of determining a noun’s class constitute a formal system that must be followed irrespective of semantic similarity judgements. For example, while wolves and sheep are both animals and can both be classified by the animal CL 只 (zhi), 狼 (lang) wolf can take another classifier, 匹 (pi).

- | | |
|---|---|
| (5) 一个 学生
yi ge xuesheng
one ge-CL student | (10) 一只 羊
yi zhi yang
one zhi-CL sheep |
| (6) 两个 学生
liang ge xuesheng
two ge-CL student | (11) 两只 羊
liang zhi yang
two zhi-CL sheep |
| (7) 三个 学生
san ge xuesheng
three ge-CL student | (12) 三只 羊
san zhi yang
three zhi-CL sheep |
| (8) *三 学生
san xuesheng
three student | (13) *三 匹 羊
san pi yang
three pi-CL sheep |
| (9) *三 只 学生
san zhi xuesheng
three zhi-CL student | (14) *三 名 羊
san ming yang
three ming-CL sheep |

You should be familiar by now with the terminology in the English grammar starter code for

¹Use the Romanized form with no spaces or hyphens, e.g., xuesheng, for multi-character lexical entries.

this question. The Chinese grammar is fairly similar, but there is a new phrasal category called a classifier phrase (CLP), formed by a number and a classifier. The classifier phrase serves the same role as a determiner does in English.

The two grammars below don't appropriately constrain the NPs generated. You need to design your own rules and features to properly enforce agreement.

English Grammar:

Rules:

NP → Det N
NP → Num N
VP → V NP
S → NP VP

Lexicon:

a:	Det
one:	Num
two:	Num
three:	Num
student:	N
students:	N
wolf:	N
wolves:	N
sheep:	N
sheep:	N
see:	V
sees:	V
saw:	V
chase:	V
chases:	V
chased:	V

Chinese Grammar:

Rules:

CLP → Num CL
NP → CLP N
VP → V NP
S → NP VP

Lexicon:

一	yi one/a	Num
两	liang two	Num
三	san three	Num
学生	xuesheng student	N
狼	lang wolf	N
羊	yang sheep	N
追逐	zhuizhu chase	V
看见	kanjian see	V
个	ge	CL
名	ming	CL
只	zhi	CL
匹	pi	CL

Here is a list of all of the nouns in this question and their acceptable classifiers:

- 狼 lang wolf: 只 zhi; 匹 pi;
- 羊 yang sheep: 只 zhi
- 学生 xuesheng student: 个 ge, 名 ming.

- (a) (6 marks) Implement one grammar for each language pursuant to the specifications above. English: q1_en.pl and Chinese: q1_zh.pl.

Neither of your grammars need to handle embedded clauses, e.g., a student caught two wolves see a sheep. Similarly for Chinese, your grammar doesn't need to parse sentences like example (15):

- (15) 一名 学生 看见 两 匹狼 追逐 一只 羊
yi ming xuesheng kanjian liang pi lang zhuizhu yi zhi yang
A student saw two wolves chase a sheep.

For the Chinese grammar, code the lexical entries in pinyin (the Romanized transcriptions of the Chinese characters).

- (b) (4 marks) Use your grammars to parse and translate the following sentences. Save and submit all the translation results in the .grale format. The results of sentence (16) should be named q1b_en.grale and the results of sentence (17) should be named q1b_zh.grale.

- (16) Two wolves saw one sheep
(17) 两个学生 追逐 三只羊
liang ge xuesheng zhuizhu san zhi yang

Operational Instructions

- Independently test your grammars in TRALE first, before trying to translate.
- Use the function `translate` to generate a semantic representation of your source sentence. If your sentence can be parsed, the function `translate` should open another `gralej` interface with all of the translation results.

```
| ?- translate([two,wolves,catch,one,sheep]).
```

- To save the translation results, on the top left of the `Gralej` window (the window with the `INITIAL CATEGORY` entry and all of the translated sentences listed), click `File >> Save all >> TRALE format`.
- Don't forget to close all of the windows or kill both of the `Gralej` processes after you finish. Each `Gralej` process will take up one port in the server, and no one can use the server if we run out of ports.

2. Quantifier Scope [30 marks]

For this assignment, we will consider two quantifiers: the universal quantifier (every, 每 mei) and the existential quantifier (a, 一 yi). In English, both quantifiers behave as singular determiners.

- (18) A professor stole every cookie.
- (19) * A professor stole every cookies.
- (20) * A professors stole every cookie.

In Chinese, both of these quantifiers behave more like numerical determiners. In addition, when a universal quantifier modifies an NP that occurs before the verb (such as with a universally quantified subject), the preverbal operator 都 (dou) is required. When a universally quantified NP occurs after the verb, the dou-operator must not appear with it.

- (21) Every professor stole a cookie.
- (22) A professor stole every cookie.
- (23) 每个 教授 都 偷了 一块 饼干
mei ge jiaoshou dou tou-le yi kuai binggan
 \forall ge-CL professor DOU stole \exists kuai-CL cookie
- (24) * 每个 教授 偷了 一块 饼干
mei ge jiaoshou tou-le yi kuai binggan
 \forall ge-CL professor stole \exists kuai-CL cookie
- (25) 一个 教授 偷了 每 块 饼干
yi ge jiaoshou tou-le mei kuai binggan
 \exists ge-CL professor stole \forall kuai-CL cookie
- (26) * 一个 教授 都 偷了 每 块 饼干
yi ge jiaoshou dou tou-le mei kuai binggan
 \exists ge-CL professor dou stole \forall kuai-CL cookie

Quantifier Scope Ambiguity In lecture, we talked about different kinds of ambiguity. In many English sentences, no matter what the order of the quantifiers, there is a quantifier scope ambiguity. For example, the sentence every student takes a course has two readings:

- ($\exists > \forall$) Every student takes a course. [The course is Math.]
- ($\forall > \exists$) Every student takes a course. [Some students take Math and some students take History.]

The symbol ($\exists > \forall$) means the existential quantifier outscopes the universal quantifier in a logical form representation of the sentence.




Figure 1: Beta Reduction. What should be the LF of S?

We can write the semantics of the two sentences in their logical forms (LF) to distinguish the two readings:

- $\exists y.(\text{course}(y) \wedge \forall x.(\text{student}(x) \Rightarrow \text{take}(x, y)))$
- $\forall x.(\text{student}(x) \Rightarrow \exists y.(\text{course}(y) \wedge \text{take}(x, y)))$

English sentences (27, 28) are scopally ambiguous no matter what the linear order of the quantifiers is. But in Chinese, a sentence is scopally ambiguous only when the universally quantified NP precedes the existential NP: (29) is ambiguous, but (30) is unambiguous.²

- (27) Every student takes a course
Ambiguous: $\exists > \forall, \forall > \exists$
- (28) A student takes every course
Ambiguous: $\exists > \forall, \forall > \exists$
- (29) 每个 学生 都 会上 一种 课程
mei ge xuesheng dou huishang yi zhong kecheng
 \forall ge-CL student DOU take \exists zhong-CL course
Ambiguous: $\exists > \forall, \forall > \exists$
- (30) 一个 学生 会上 每种 课程
yi ge xuesheng huishang mei zhong kecheng
 \exists ge-CL student take \forall zhong-CL course
Unambiguous: $\exists > \forall$

How can we derive the LF of the two readings? We use a process called beta reduction. Recall the lambda-calculus notation: $\lambda x.x^2$ denotes a function that takes a variable x , and returns the square of its value (x^2). After substituting the value for the bound variable x , we can reduce the function application in the body of the lambda term to a new expression. For example, applying 2 to $\lambda x.x^2$ will get us:

$$\lambda x.x^2(2) = 2^2$$

²The actual principles that determine the scopal readings of a Chinese sentence are still an active area of research, but this is obviously beyond the scope of this assignment. You only need to construct an analysis that explains every example mentioned in this assignment.




Figure 2: Beta reduction analysis of the sentence every student takes a course.

This process is also known as beta reduction (denoted as \Leftrightarrow_β). Note that beta reduction itself does not tell us that this equals 4. That is obtained by a subsequent process of arithmetic evaluation. But we can use beta reduction even if we don't evaluate.

We can also perform beta reduction on variables for functions. For example, applying in $\lambda F.F(2)$ to $\lambda x.x^2$ will yield:

$$\lambda F.F(2) (\lambda x.x^2) = (\lambda x.x^2)(2) = 2^2$$

Now, let's look at an example that uses beta reduction to compute the LF of a sentence. For example, as shown in figure 1, we know that the LF of the NP every student is $\lambda P.\forall x.(\text{student}(x) \Rightarrow P(x))$ and the LF of the VP takes a course is $\lambda z.\exists y.(\text{course}(y) \wedge \text{take}(z, y))$. What is the LF of every student takes a course?

$$\begin{aligned} & \lambda P.\forall x.(\text{student}(x) \Rightarrow \lambda z.\exists y.(\text{course}(y) \wedge \text{take}(z, y))) \\ \Leftrightarrow_\beta & \forall x.(\text{student}(x) \Rightarrow \lambda z.\exists y.(\text{course}(y) \wedge \text{take}(z, y))(x)) \\ \Leftrightarrow_\beta & \forall x.(\text{student}(x) \Rightarrow \exists y.(\text{course}(y) \wedge \text{take}(x, y))) \end{aligned}$$

Each step of repeatedly applying beta reduction to every subterm until we reach an irreducible statement is called beta normalisation.

Figure 2 shows the complete analysis of the sentence every student takes a course. Familiarize yourself with every part of the analysis. But this only generates one of the two readings – the surface reading ($\forall > \exists$). We will use a technique called quantifier storage to capture the scopal ambiguity and make both readings available.

Quantifier Storage If quantifier scoping is a semantic effect, how do we represent it in syntax? When there is no ambiguity, keeping track of the quantifier scope is pretty straightforward. To keep track of and resolve scope ambiguities, we will use a list called a quantifier store. The idea behind QSTORE is that, instead of consuming all of the LF components right away, we can choose to keep them in QSTORE and apply them later.

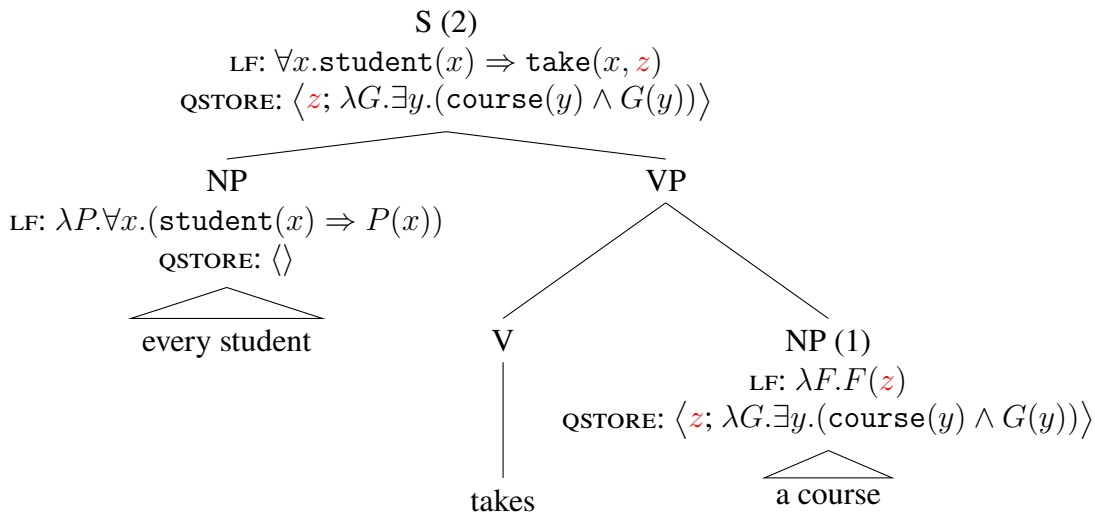


Figure 3: Quantifier Storage. Storing the quantifier at (1), and retrieve it later at (2).

Let's go back to the example, every student takes a course (figure 3). We first store the LF of the NP a course at (1) and replace the LF of the NP with a placeholder $\lambda F.F(z)$. The variable z in this expression is a free occurrence, and it is the same variable as the z in the store and in the LF of the sentence (the free occurrences of z are highlighted in red). We retrieve the logical form from the store at (2). The retrieval process consists of three steps:

1. First, we construct a function $\lambda z.L_S$, where L_S is the current LF, and z is the variable paired in the QSTORE entry. In our particular case, this will yield $\lambda z.(\forall x, \text{student}(x) \Rightarrow \text{take}(x, z))$.
2. Then, we apply this function to the LF from the QSTORE entry.
3. Finally, we beta normalise. Using beta normalisation, we obtain the second reading of the sentence.

$$\begin{aligned}
& \lambda G.\exists y.(\text{course}(y) \wedge G(y))(\lambda z.(\forall x, \text{student}(x) \Rightarrow \text{take}(x, z))) \\
& \Leftrightarrow_{\beta} \exists y.(\text{course}(y) \wedge (\lambda z.(\forall x.\text{student}(x) \Rightarrow \text{take}(x, z))(y)) \\
& \Leftrightarrow_{\beta} \exists y.(\text{course}(y) \wedge (\forall x.\text{student}(x) \Rightarrow \text{take}(x, y))
\end{aligned}$$

Topicalization and Movement Topicalization is a linguistic phenomenon in which an NP appears at the beginning of a sentence in order to establish it as the topic of discussion in a sentence or to emphasize it in some other way. It plays an important role in the syntax of fixed-word-order courses because grammatical function is mainly determined by word order. Both Chinese and English exhibit topicalization. The entire object NP, for example, can be moved to the beginning of the sentence in either course. But in Chinese, object topicalization is more restricted when the subject is quantified: it can happen when the subject is universally quantified, but not when it is existentially quantified (33-36).

- (31) A course, every student takes.
 \exists course \forall student take
 Ambiguous: $\forall > \exists$ and $\exists > \forall$
- (32) Every course, a student takes.
 \forall course \exists student take
 Ambiguous: $\forall > \exists$ and $\exists > \forall$
- (33) 一种 课程 每个 学生 都 会上
 yi zhong kecheng mei ge xuesheng dou huishang
 \exists zhong-CL course \forall ge-CL student dou take
 Unambiguous: $\exists > \forall$
- (34) 每种 课程 每个 学生 都 会上
 mei zhong kecheng mei ge xuesheng dou huishang
 \forall zhong-CL course \forall ge-CL student dou take
- (35) *一种 课程 一个 学生 会上
 yi zhong kecheng yi ge xuesheng huishang
 \exists zhong-CL course \exists ge-CL student take

- (36) * 每 种 课 程 一 个 学 生 (都) 会 上
 mei zhong shu yi ge xuesheng (dou) huishang
 \forall ben-CL course \exists ge-CL student dou take




Figure 4: English topicalization parse tree: example (31).

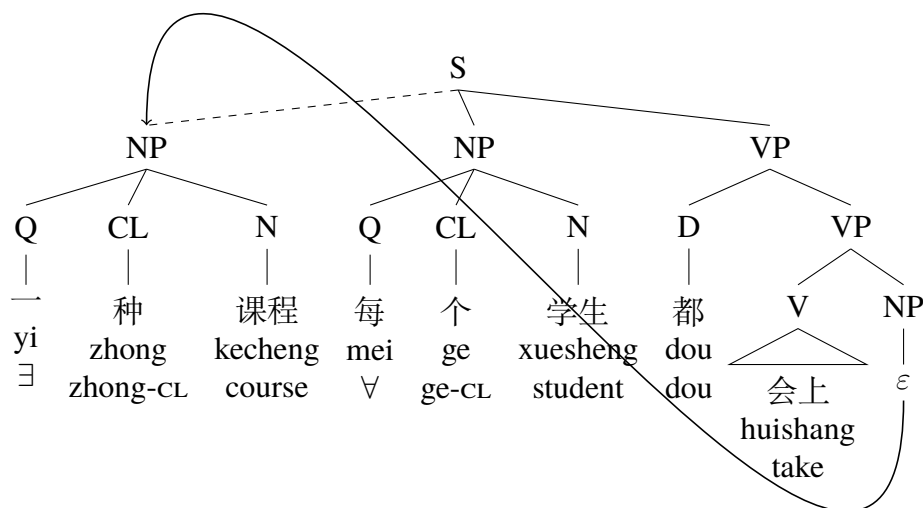


Figure 5: Chinese topicalization parse tree: example (33).

In English, neither subject–verb agreement nor quantifier scope ambiguity is generally affected by movement. In particular, the number and person of the subject should always agree with the predicate no matter where it occurs. Here, you can assume that Chinese also follows subject–verb agreement (regarding the requirement of *dou*) in the same way that English does. But whereas in English, both readings are still available after the sentences are topicalised (31, 32), this is not the case in Chinese. Compared to its untropicalised counterpart (29), the topicalised sentence (33) is no longer ambiguous.

Figures 4 and 5 show the parse trees of sentences (31) and (33). Topicalization is generally analysed with gaps. An empty trace is left in the untropicalized position of the object NP, where the gap is introduced. The gapped NP then percolates up the tree, and is finally unified with the topicalized NP at the left periphery of the sentence.³

³Although Chinese is an SVO (Subject-Verb-Object) course, there is a means of performing “double movement.”

- (a) (2 marks) Manually convert all readings of the sentences (28) and (30) to logical expressions. Put your logical forms in section 2(a) of `analysis.txt`. Use `exists` and `forall` for the quantifiers, and use `=>` and the caret symbol `^` for implication and conjunction.
- (b) (10 marks) Implement grammars for the syntax of quantifier scope ambiguity. You don't need to account for meanings, or for ambiguity in meanings (there should be no syntactic ambiguities). At this point, a correct grammar will produce exactly one parse for every grammatical sentence. Test your implementation before you move on to the next step.
- (c) (10 marks) Augment your grammars to represent meaning and quantifier scope ambiguity. Marks for question 2(b) will be deducted if your work on this part causes errors in the syntactic predictions. Your grammar should generate more than one parse for each ambiguous sentence.
- (d) (4 marks) Translate sentences (28) and (30), as you did in the first question.

Operational Instructions

- Use the function `translate` to generate semantic representations of your source sentences. If your sentences can be parsed, `translate` should open another `gralej` window and with all of the translation results.

```
| ?- translate([a,student,takes,every,course]).
```

- You will be prompted as follows to see the next parse.

```
ANOTHER? y
...
ANOTHER? y
no
```

Answer `y` to see the next parse until you reach the end. Each time `TRALE` will open a new `Gralej` window. You need to store all of your translation results by repeating the previous step. A `no` will be returned when you reach the end of your parses.

- Save your translations of sentence (28) as `q2d_28_1.grale`, `q2d_28_2.grale` ... and your translations of sentence (30) as `q2d_30_1.grale`, `q2d_30_2.grale` ...
- Submit a zip file `q2d.zip` containing all the translation results. You can use this command: `zip -r q2d.zip q2d_*.grale` to create the zip file.
- Again, don't forget to close all the windows and kill your `Gralej` processes after you finish.

-
- (1) 一个 学生 每 种 课程 都 会上
 yi ge xuesheng mei zhong kecheng dou huishang
 \exists ge-CL student \forall zhong-CL course dou take
 A student every course take.

We will ignore these.

- (e) (4 marks) Large language Models (LLMs) have gained quite a lot of popularity recently. In this question, you will explore the question of whether LLMs such as ChatGPT⁴ “understand” quantifier scope.

There are several approaches to this exploration: You can ask ChatGPT to translate sentences and compare its translations to your grammar’s translations. Or, you can directly interrogate ChatGPT about quantifier scope readings and analyse its responses. You can also design some clever linguistic tasks that involve quantifier scope and catch how ChatGPT handles them. Be both creative and precise in your experimentation.

In your write-up, report at least one case where GPT’s behaviour differs from that of your grammar. Document the prompts you used and describe your experimental design. Reflect on the differences observed and share your thoughts on why these differences may have occurred.

Your analysis should be submitted as section 2(e) in `analysis.txt`.

⁴<https://chat.openai.com/>

What to Submit

This assignment is submitted electronically via MarkUs. You should submit a total of eight (8) required files as follows:

- `q1_en.pl` The (entire) `q1_en.pl` with your implementations filled in.
- `q1_zh.pl` The (entire) `q1_zh.pl` with your implementations filled in.
- `q1b_en.grale` The translation result saved in `.grale` format for the English sentence in question 1(b).
- `q1b_zh.grale` The translation result saved in `.grale` format for the Chinese sentence in question 1(b).
- `analysis.txt` Logical expressions for question 2(a).
- `q2_en.pl` The (entire) `q2_en.pl` with your implementations filled in.
- `q2_zh.pl` The (entire) `q2_zh.pl` with your implementations filled in.
- `q2d.zip` The translation results saved in `.grale` format (and compressed into a `.zip` file) for question 2(d).