Non-linear regression on dependency trees

Ramon Ferrer-i-Cancho, Marta Arias, Argimiro Arratia

Complex and Social Networks (2026-2027)
Master in Innovation and Research in Informatics (MIRI)

1 Introduction

In this session, we are going to practice on the fit of a non-linear function to data using collections of syntactic dependency trees from different languages. In a syntactic dependency trees, the vertices are the words (*tokens*) of a sentence and links indicate syntactic dependencies between words [Ferrer-i-Cancho, 2013].

n is defined as the number of vertices of a tree, k_i as the degree of the i-th vertex and d_i as the length of the i-th edge. Notice that, given n, the average degree is constant (it depends only on n. $\langle k^2 \rangle$ is defined as its degree 2nd moment about zero, i.e.

$$\langle k^2 \rangle = \frac{1}{n} \sum_{i=1}^{n} k_i^2$$

$\langle d \rangle$ is defined as the mean length of its edges, i.e.

$$\langle d \rangle = \frac{1}{n-1} \sum_{i=1}^{n-1} d_i.$$

Through some procedure, two groups of students will be formed:

- The degree 2nd moment group. Its teams will have to investigate the scaling of $\langle k^2 \rangle$ as a function of n.
- The mean length group. Its teams will have to investigate the scaling of $\langle d \rangle$ as a function of n.

Each team within a group works independently from other members but is allowed to compare results with other teams within the group.

2 Data preparation

Proceed as follows to generate the data on information of real syntactic dependency trees from different languages:

1. Download the latest version of the Universal Dependencies (UD) treebank collection from `https://universaldependencies.org/download.html`. Use the same list of languages from previous labs.
2. Create a file with name `language_dependency_tree_metrics.txt` for each language, e.g. `English_dependency_tree_metrics.txt`. Within the file, each row corresponds to the syntactic dependency tree of a real sentence. Within each row, the 1st, the 2nd and the 3rd cell contain, respectively, n, $\langle k^2 \rangle$ and $\langle d \rangle$ for a given sentence.
3. Store these files in a directory `data` that you will have to include in your delivery.

Produce an `R`script to check the validity of $\langle k^2 \rangle$ and $\langle d \rangle$. In order to be valid, these metrics must satisfy (at least)[Ferrer-i-Cancho, 2013]

$$4 - 6/n \leq \langle k^2 \rangle \leq n - 1 \tag{1}$$

and also [Ferrer-i-Cancho et al., 2021]

$$\frac{1}{4}\left(n\langle k^2 \rangle + q\right) \leq (n-1)\langle d \rangle \leq \frac{1}{4}\left(3(n-1)^2 + 1 - n \bmod 2\right) \tag{2}$$

where q is the number of vertices of odd degree, i.e.

$$q = \sum_{i=1}^{n} (k_i \bmod 2). \tag{3}$$

You have to produce a table with the format of Table 1. This table might be necessary to interpret the results that are going to be obtained in coming sections. You can adapt the script `summary_table.R` from session 2. The sample size, the mean and then standard deviation can be obtained with the functions `length(...)`, `mean(...)` and `sd(...)`.

3 Data analysis

```
> English = read.table("./data/English_dependency_tree_metrics.txt", header = FALSE)
> colnames(English) = c("vertices","degree_2nd_moment", "mean_length")
> English = English[order(English$vertices), ]
```

loads the information about a collection of dependency trees from sentence in English and sorts the matrix rows by the number of vertices increasingly (that will help later when visualizing the plots).

Table 1: Summary of the properties of the degree sequences. N is the sample size (the number of sentences or dependency trees), μ_n and σ_n are, respectively, the mean and the standard deviation of n, the sentence length (n is the number of vertices of a trees), μ_x and σ_x are the mean and the standard deviation of the target metric ($\langle k^2 \rangle$ in one group; $\langle d \rangle$ in the other group)

Language	N	μ_n	σ_n	μ_x	σ_x
...	...	...	...	...	...
...	...	...	...	...	...
...	...	...	...	...	...

3.1 Preliminary visualization

Note: The following plots are generated using vintage R. Use ggplot2 for plots in your report.

Consider the mean dependency length. A preliminary plot can be obtained with

```
> plot(English$vertices, English$mean_length,
       xlab = "vertices", ylab = "mean dependency length")
```

The same plot taking logs on both axes

```
> plot(English$vertices, English$mean_length,
       xlab = "vertices", ylab = "mean dependency length", log = "xy")
```

suggest a power-law dependency between mean length and number of vertices, in spite of the high dispersion. A clearer intuition about the underlying trend can be obtained by averaging mean lengths for a given number of vertices:

```
> mean_English = aggregate(English, list(English$vertices), mean)
```

Now compare the plot obtained by

```
> plot(mean_English$vertices, mean_English$mean_length,
       xlab = "vertices", ylab = "mean mean dependency length")
```

with

```
> plot(mean_English$vertices, mean_English$mean_length,
       xlab = "vertices", ylab = "mean mean dependency length", log = "xy")
```

An intuition about how far the real scaling of $\langle d \rangle$ is from the a random linear arrangement can be obtained by adding the expected mean length in that case,

which is $E[\langle d \rangle] = (n+1)/3$ [Ferrer-i-Cancho, 2004], to the plots. Consider for instance a plot in double logarithmic scale with the averaged curve and the random linear arrangement expectation

```
> plot(English$vertices, English$mean_length,
       xlab = "vertices", ylab = "mean dependency length", log = "xy")
> lines(mean_English$vertices,mean_English$mean_length, col = "green")
> lines(mean_English$vertices,mean_English$vertices+1)/3, col = "red")
```

As for the scaling of $\langle k^2 \rangle$ a suitable null model are uniformly distributed random undirected trees. $f(n)$ can be estimated numerically by producing many of those trees for a given n. The Aldous-Brother algorithm allows one to generate uniformly random labelled spanning trees from a graph [Aldous, 1990, Broder, 1989]. Here we assume a complete graph as the source for the spanning trees. Indeed, it has been shown that

$$\langle k^2 \rangle = \left(1 - \frac{1}{n}\right)\left(5 - \frac{6}{n}\right) \tag{4}$$

for uniformly random labelled trees [Ferrer-i-Cancho, 2014].

An initial exploration of the scaling of $\langle k^2 \rangle$ can be performed with its expected value in uniformly random trees (Eq. 4) and the theoretical lower and upper bounds (Eq. 1):

```
> plot(English$vertices, English$degree_2nd_moment,
       xlab = "vertices", ylab = "degree 2nd moment")
> lines(mean_English$vertices,mean_English$degree_2nd_moment, col = "green")
> lines(English$vertices,
       (1 - 1/English$vertices)*(5 - 6/English$vertices), col = "red")
> lines(English$vertices,4-6/English$vertices, col = "blue")
> lines(English$vertices,English$vertices-1, col = "blue")
```

The plots suggest that both $\langle d \rangle$ and $\langle k^2 \rangle$ grow sublinearly with n. The plot for $\langle d \rangle$ versus n suggests an almost power-law dependency. The functional dependency between these metrics and n will be investigated next.

3.2 The ensemble of models

We consider an ensemble of models

- $f(n) = (n/2)^b$ (model 1). This models is obtained applying the condition $f(2) = 1$ (satisfied both by $\langle d \rangle$ and $\langle k^2 \rangle$) to a more general model, i.e. $f(n) = an^b$. This leads to $a = 1/2^b$ and finally $f(n) = (n/2)^b$. The motivation of this model is that $\langle k^2 \rangle = \langle d \rangle = 1$ when $n = 2$.

- $f(n) = an^b$ (model 2), a power-law model.
- $f(n) = ae^{cn}$ (model 3), an exponential model.
- $f(n) = a \log n$ (model 4), a logarithmic model.

Model 4 is motivated by the logarithmic growth of $\langle d \rangle$ in minimum linear arrangements of uniformly random labelled trees [Esteban et al., 2016].

Additionally, it is convenient to consider a null model (model 0). Concerning the scaling of $\langle d \rangle$, a random linear arrangement of vertices gives $f(n) = n/3 + 1/3$ as the null model. Concerning the scaling of $\langle k^2 \rangle$, uniformly random labelled trees give $f(n) = \left(1 - \frac{1}{n}\right)\left(5 - \frac{6}{n}\right)$ as the null model. Notice that $f(n)$ depends exclusively on n in both null models (no extra parameters are needed).

The additive term of the random linear arrangement model suggests that the models above should be generalized giving

- $f(n) = (n/2)^b + d$ (model 1+).
- $f(n) = an^b + d$ (model 2+).
- $f(n) = ae^{cn} + d$ (model 3+).
- $f(n) = a \log n + d$ (model 4+).

- Certain + models can be hard to fit.
- The ease of fit of a + model may depend on the group ($\langle k^2 \rangle$ versus $\langle d \rangle$).
- We recommend working in two phases. First, work an all models except the hard-to-fit + models produce an early draft of the report. Second, add the remainder of + models and produce the final report.

4 Non-linear regression with R

We show how to fit model 2 to $\langle d \rangle$ as a function of n. The procedure for $\langle k^2 \rangle$ is the same but replacing `mean_length` by `degree_2nd_moment`. The non-linear regression can be invoked by means of the comment `nls(...)` through,

```
> a_initial = 4
> b_initial = 4
> nonlinear_model = nls(mean_length~a*vertices^b,data=English,
  start = list(a = a_initial, b = b_initial), trace = TRUE)
```

where

- `mean_length~a*vertices^b` is the mathematical definition of the function to fit,
- `data=English` indicates the data source,
- `start = list(a = a_initial, b = b_initial)` defines the initial values of the parameters,
- `trace = TRUE` indicates that the progress of the optimization algorithm must be shown.

We chose initial values for the parameters a and b arbitrarily. The initial values can be crucial to warrant that `nls(...)` is able to find a solution to the optimization problem or simply to increase the time efficiency of the non-linear regression. Try for instance `a_initial = 0` to see an example of failure of `nls(...)`. Good initial values for a and b can be obtained with a double logarithmic transformation, which gives

$$\log \langle d \rangle = b \log n + a', \tag{5}$$

where $a' = \log a$. A linear regression on that transformation allows one to obtain an initial value for b and initial value for a thanks to $a = e^{a'}$. This is what the next code does:

```
> linear_model = lm(log(mean_length)~log(vertices), English)
> a_initial = exp(coef(linear_model)[1])
> b_initial = coef(linear_model)[2]
```

`lm(...)` is used to perform a non linear regression between $\log \langle d^2 \rangle$ and $\log n$. `coef(...)` retrieves the parameters of the linear model: `coef(...)[1]` contains the intercept and `coef(...)[2]` contains the slope.

We run the non-linear regression again with

```
> nonlinear_model = nls(mean_length~a*vertices^b,data=English,
  start = list(a = a_initial, b = b_initial), trace = TRUE)
```

Notice the faster convergence of `nls(...)` with the new initial values of the parameters.

Notice that a hidden assumption of the way we call `nls(...)` and non-linear regression in general is homocesdasticity **of residuals**. You have to check if this assumption holds. Plots of the original data can help. As a first approximation, you can look at the variance of points as a function of n, the number of vertices of the tree. A stronger analysis is the on variance of residuals as function of n. In case that homocesdasticity does not hold, we suggest that `nls(...)` is fed with the output of `aggregate(...)` and not the original data. There are other approaches.

The RSS and the AIC of the nonlinear regression model can be retrieved, respectively, simply with

```
deviance(nonlinear_model)
```

and

```
AIC(nonlinear_model)
```

> It is important to be aware of a crucial difference between this lab session and an Session 2. Here we are following a nonlinear regression approach to curve fitting. Its goal is minimizing the error between the curve and the model (and the AIC is a function of that error). In contrast, session 2 was based on a maximum likelihood approach where there is no direct concern about that error. The AIC is computed differently depending on the approach.

s can be obtained through

```
> sqrt(deviance(nonlinear_model)/df.residual(nonlinear_model))
```

The parameters giving the best fit can be obtained through

```
coef(nonlinear_model)
```

`coef(nonlinear_model)[i]` contains the value of the i-th parameters for the best fit of the model ($i = 1$ for the 1st parameters). The mapping of parameter index to parameter is tricky and thus it is convenient that you use `coef(nonlinear_model)["a"]` to retrieve the value of a giving the best fit and `coef(nonlinear_model)["b"]` to retrieve the value of b giving the best fit.

Notice `nls(...)` cannot be used for models with no parameters. For instance, the RSS and s^2 for $\langle d \rangle = (n+1)/3$ can be calculated with

```
RSS <- sum((English$mean_length-(English$vertices+1)/3)^2)
n <- length(English$vertices)
p <- 0
s <- sqrt(RSS/(n - p))
```

Finally, the corresponding AIC (assuming a non-linear regression model) can be calculated with [Ritz and Streibig, 2008, p. 105; Eq. 7.2]

```
AIC <- n*log(2*pi) + n*log(RSS/n) + n + 2*(p + 1)
```

Language	0	1	2	3	4	1+	2+	3+	4+
...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...

Table 2: Do not forget to include results for model 0 here

	Model															
	1	2		3		4	1+		2+			3+			4+	
Language	*b*	*a*	*b*	*a*	*c*	*a*	*b*	*d*	*a*	*b*	*d*	*a*	*c*	*d*	*a*	*d*
...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...

Table 3: Notice that Model 0 is not included here because it has no parameters.

5 Results

5.1 Model selection

For each metric, you have to prepare:

- Separate tables for s (residual standard error), AIC (Akaike information criterion), Δ (AIC differences) and AIC weights following the format of Table 2.
- Table with the values of the parameters giving the best fit for each model following the format of Table 3.

5.2 Final visualization

For each language, you have to plot the empirical data and the curve for the best fit. Imagine that model 2 is the best model for the English language. Then the data and the best fit can be plotted together with

```
> plot(English$vertices, English$mean_length,
       xlab = "vertices", ylab = "mean dependency length", log = "xy")
> lines(English$vertices, fitted(nonlinear_model), col = "green")
```

6 Deliverables

You have to prepare a report including the following sections (in this order): introduction, results, discussion and methods. Results includes all the tables

and figures (the preliminary plots and the plots of the best model to the real data) and some guiding text. Methods should include any relevant methods not explained in this guide (for instance, decisions that you had to made and might have an influence on the results), initial values of the parameters used to call `nls()`, the techniques used to obtain those initial values and so on... The discussion should include a summary of the results and your interpretation. For instance, you should discuss

- If there is a significant difference between the fit of the functions from null hypotheses and that of alternative hypotheses.
- If the original data satisfy the assumption of homoscedasticity of the non-linear regression methods considered here. In case that it does not hold, you should explain how you have addressed it.
- Discuss if the function giving the best fit gives a reasonably good fit (e.g., checking visually that the best function provides a sufficiently good fit). Remember that the best function of an ensemble is not necessarily the best in absolute terms.
- The extent to which languages resemble or differ.

The discussion section should also include some conclusions.

Important rule: The lab session, and especially the report you have to hand in, are strictly individual work. Plagiarism will be prosecuted. Nevertheless, you are encouraged to ask the teacher as soon as possible if you think you do not understand what you are supposed to do, and also if you feel you are spending much more time than the rest of the group – sometimes a tiny error can be tricky to find and does not add much to your knowledge. Questions can be asked either in person or by email, and you will never be penalized by asking questions, no matter how stupid they look in retrospect.

To deliver: You must deliver

1. The lab report explained above as a PDF.
2. The source code of your implementations, including some minimal comments that can help the reader, in a folder called `code`.
3. Preprocessed data in a folder called `data`.
4. A self-assessment report on your use of AI or other resources (GitHub, or other repositories, etc) and the extend of it (if any) as a PDF.
5. Do not forget to put full names of all members of your team in the lab report and the self-assesment report.

Procedure: Submit your work through the raco platform as a single zipped file.

Deadline: Work must be delivered within 2 weeks from the lab session you attend. Late deliveries risk being penalized or not accepted at all. If you anticipate problems with the deadline, please tell us as soon as possible.

7 Advanced exercises

- Compare the results of model selection using AIC against model selection with BIC (Bayesian Information Criterion) [Wagenmakers and Farrell, 2004].
- Consider adding the following models:
  - $f(n) = an^b e^{cn}$ (model 5), a generalization of model 2 that is widely used in quantitative linguistics.
  - $f(n) = an^b e^{cn} + d$ (model 5+), a generalization of model 5 with an additive term.

The challenge of the advanced exercises may require a careful choice of the initial values of the parameters used to call `nls(...)`, tuning some parameters of `nls(...)` or using `aggregate(...)` to smooth the original dataset.

References

[Aldous, 1990] Aldous, D. (1990). The random walk construction of uniform spanning trees and uniform labelled trees. *SIAM J. Disc. Math.*, 3:450–465.

[Broder, 1989] Broder, A. (1989). Generating random spanning trees. In *Symp. Foundations of Computer Sci., IEEE*, pages 442–447, New York.

[Esteban et al., 2016] Esteban, J. L., Ferrer-i-Cancho, R., and Gómez-Rodríguez, C. (2016). The scaling of the minimum sum of edge lengths in uniformly random trees. *Journal of Statistical Mechanics*, page 063401.

[Ferrer-i-Cancho, 2004] Ferrer-i-Cancho, R. (2004). Euclidean distance between syntactically linked words. *Physical Review E*, 70:056135.

[Ferrer-i-Cancho, 2013] Ferrer-i-Cancho, R. (2013). Hubiness, length, crossings and their relationships in dependency trees. *Glottometrics*, 25:1–21.

[Ferrer-i-Cancho, 2014] Ferrer-i-Cancho, R. (2014). A stronger null hypothesis for crossing dependencies. *Europhysics Letters*, 108(5):58003.

[Ferrer-i-Cancho et al., 2021] Ferrer-i-Cancho, R., Gómez-Rodríguez, C., and Esteban, J. L. (2021). Bounds of the sum of edge lengths in linear arrangements of trees. *Journal of Statistical Mechanics*, page 023403.

[Ritz and Streibig, 2008] Ritz, C. and Streibig, J. C. (2008). *Nonlinear regression with R*. Springer, New York.

[Wagenmakers and Farrell, 2004] Wagenmakers, E.-J. and Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11:192–196.