Estimating Relative Query Relevance from Search Engine Logs

Georges Dupret, Carlos Hurtado, Benjamin Piwowarski, Marcelo Mendoza

gdupret@dcc.uchile.cl

Universidad de Chile

Departamento de Ciencias Físicas y Matemáticas

Ciencias de la Computación

Objective

Example of a typical query:

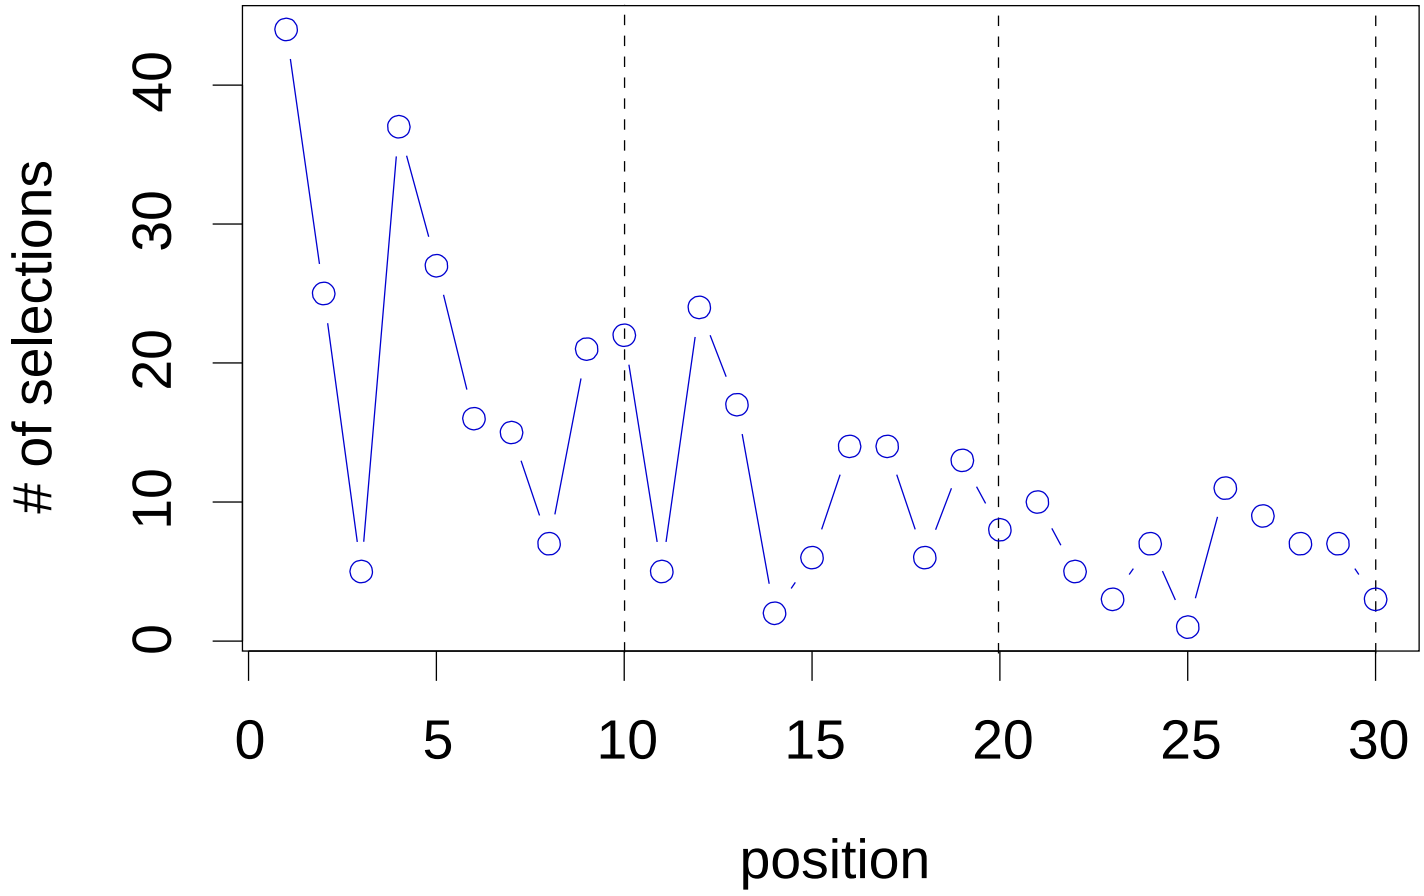

Objective

Example of a typical query:

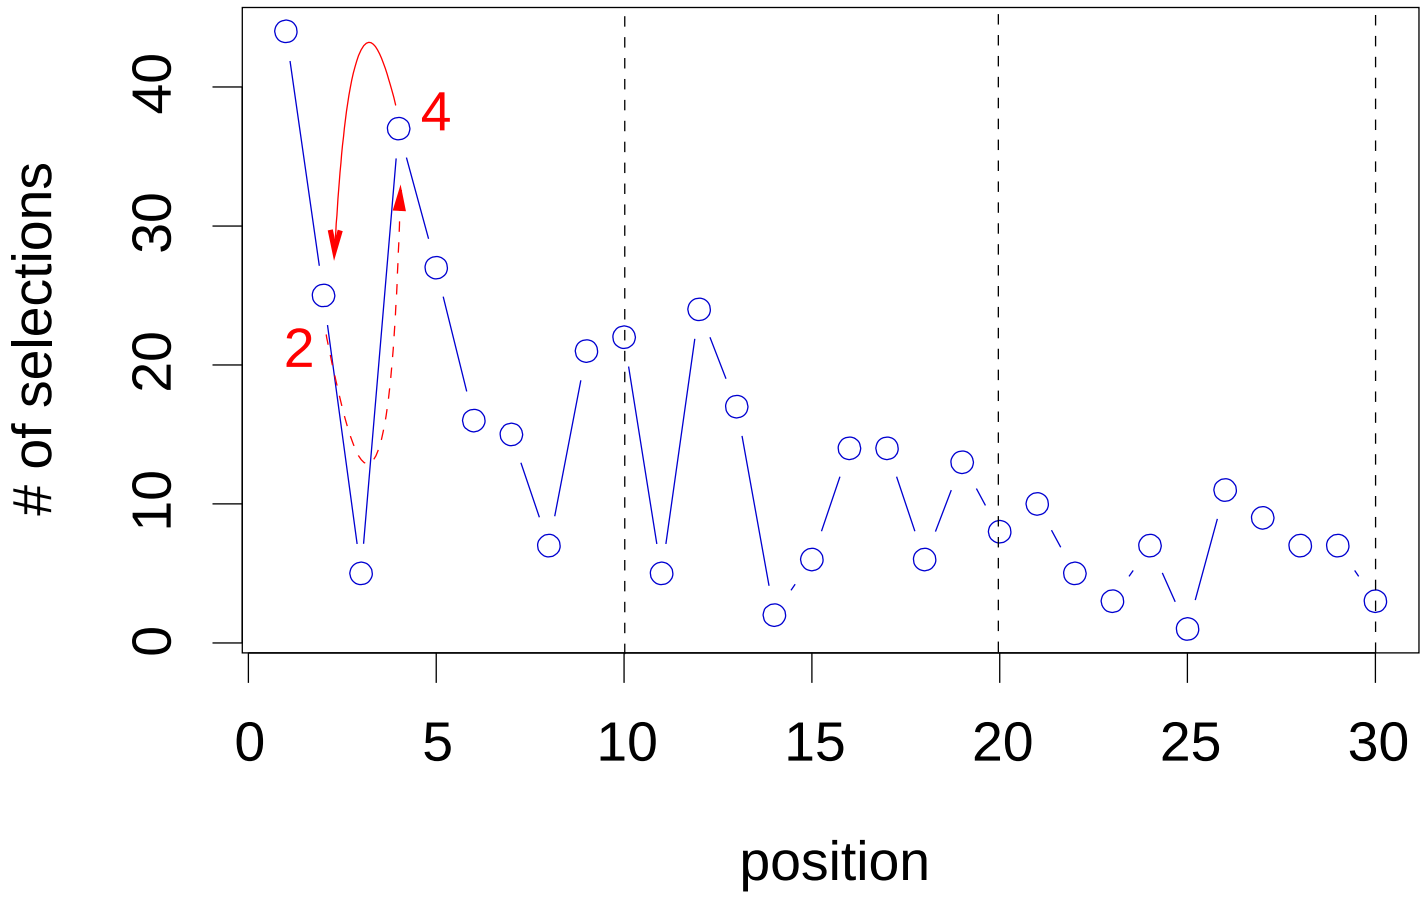

Objective

Example of a typical query:

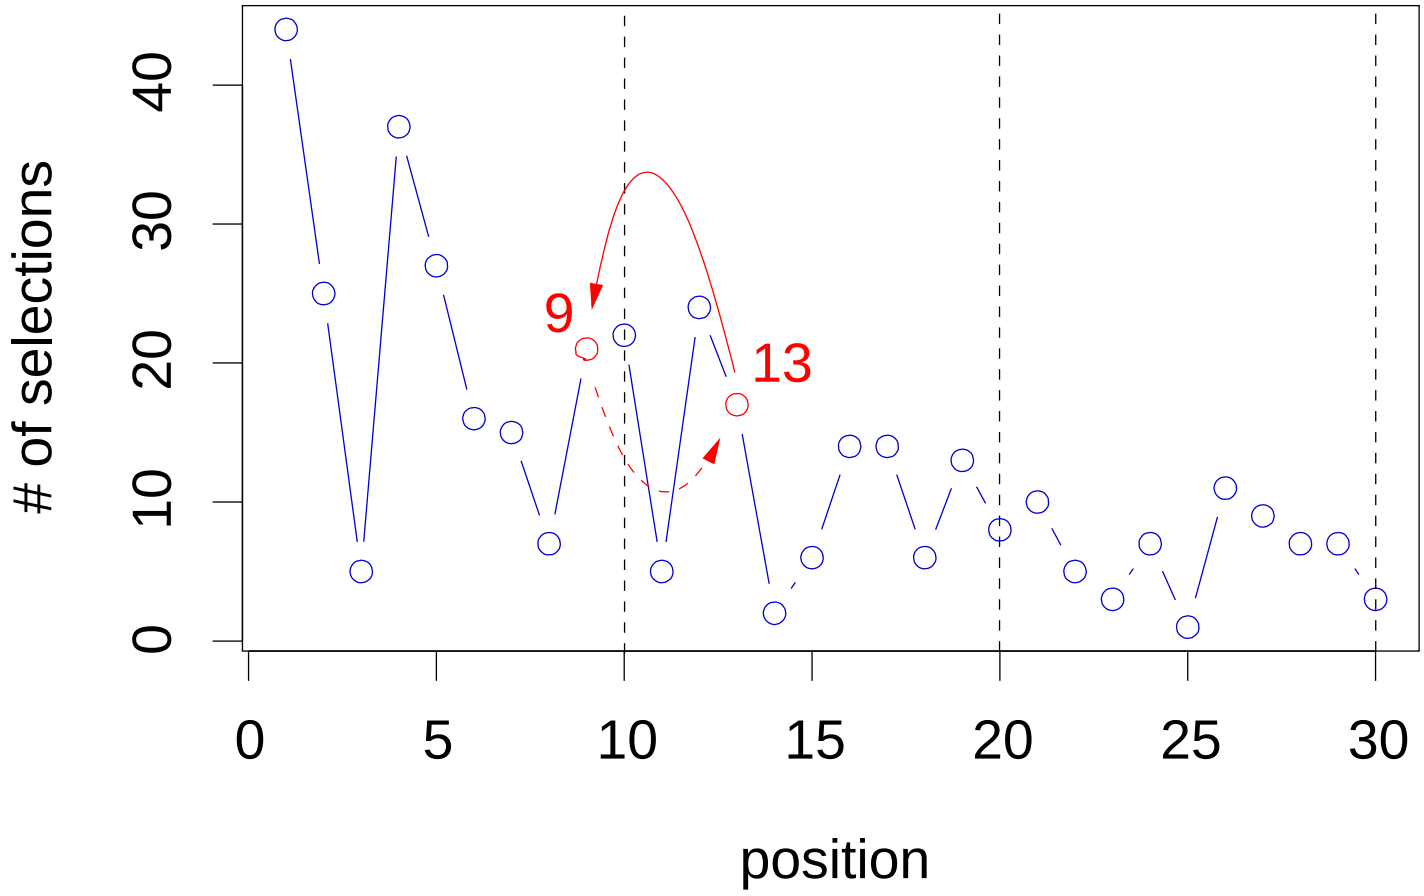

Motivation

- Engine improvement:
  - Re-ranking of precomputed queries,
  - Relevance feedback of precomputed queries,
  - Adjustment of retrieval model parameters,
  - Engine evaluation.
- Query clustering.
- Interface effect evaluation.

Typical Workaround

- Ignore position effect,
- Consider binary variables,
- Compensate by the proportion of total number of selection at the position, etc.

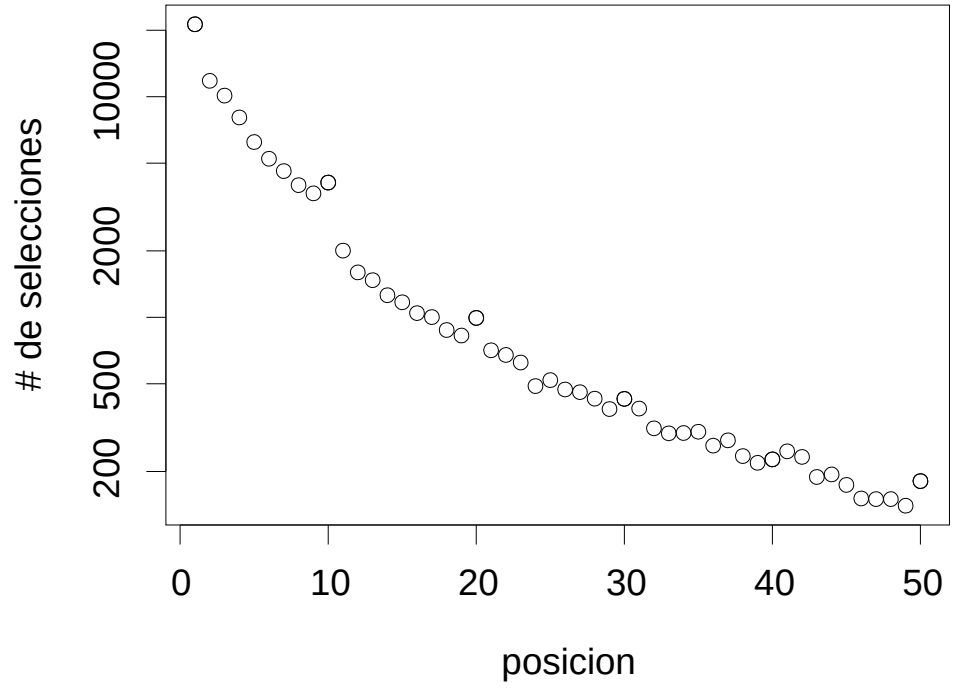

Bayesian Model:
The process of generating the logs

Bayesian Model

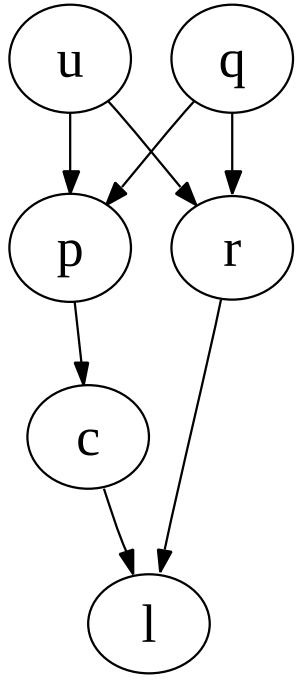

Variables:

- document u
- query q
- relevance r
- position p
- consideration c
- selection l (log)

Bayesian Model

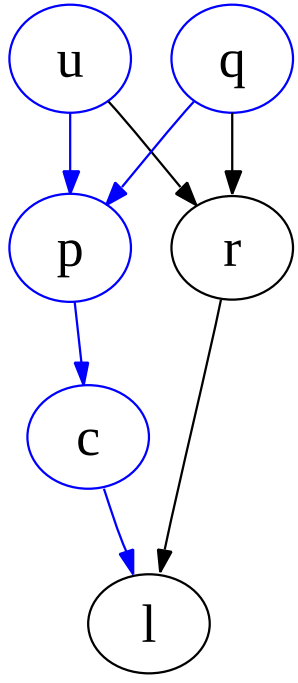

Variables:

- document u
- query q
- relevance r
- position p
- consideration c
- selection l (log)

Bayesian Model

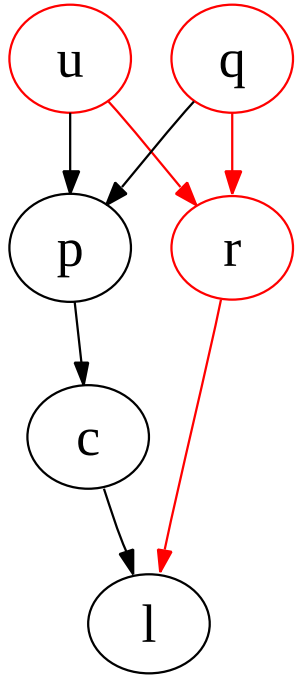

Variables:

- document u
- query q
- relevance r
- position p
- consideration c
- selection l (log)

Position Effect

$$\begin{aligned} \mathsf{p}(\mathbf{l}, p, u, q) &= \mathsf{p}(\mathbf{c}|p)\mathsf{p}(p|u,q)\mathsf{p}(\mathbf{r}|u,q)\mathsf{p}(u,q) \\ &= \mathsf{p}_{\mathcal{L}}(p,u,q) \end{aligned}$$

Hipothesis: consideration depends only on the position:

$$\mathsf{p}(\mathbf{c}|p) \rightarrow \mathsf{p}_{\mathcal{I}}(p)$$

The interface effect comprises:

- position,
- colors,
- page size, etc.

Position Effect

$$\begin{aligned} \mathsf{p}(\mathbf{l}, p, u, q) &= \mathsf{p}(\mathbf{c}|p)\mathsf{p}(p|u, q)\mathsf{p}(\mathbf{r}|u, q)\mathsf{p}(u, q) \\ &= \mathsf{p}_{\mathcal{L}}(p, u, q) \end{aligned}$$

Hipothesis: consideration depends only on the position:

$$\mathsf{p}(\mathbf{c}|p) \rightarrow \mathsf{p}_{\mathcal{I}}(p)$$

$$\mathsf{p}_{\mathcal{L}}(p, u, q) = \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}(p|u, q)\mathsf{p}(\mathbf{r}|u, q)\mathsf{p}(u, q)$$

Relevance

$$\mathsf{p}_{\mathcal{L}}(p,u,q) \;=\; \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}(p|u,q)\mathsf{p}(\mathbf{r}|u,q)\mathsf{p}(u,q)$$

Hipothesis: Relevance depends on the document and the query:

$$\mathsf{p}(\mathbf{r}|u,q)\mathsf{p}(u,q) \;\rightarrow\; \mathsf{p}_{\mathcal{R}}(u,q)$$

Relevance

$$\mathsf{p}_{\mathcal{L}}(p,u,q) \;=\; \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}(p|u,q)\mathsf{p}(\mathbf{r}|u,q)\mathsf{p}(u,q)$$

Hipothesis: Relevance depends on the document and the query:

$$\mathsf{p}(\mathbf{r}|u,q)\mathsf{p}(u,q) \;\rightarrow\; \mathsf{p}_{\mathcal{R}}(u,q)$$

$$\mathsf{p}_{\mathcal{L}}(p,u,q) \;=\; \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}(p|u,q)\mathsf{p}_{\mathcal{R}}(u,q)$$

Search Engine

$$\mathsf{p}_{\mathcal{L}}(p,u,q) \;=\; \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}(p|u,q)\mathsf{p}_{\mathcal{R}}(u,q)$$

The search engine determines the position given the document and the query:

$$\mathsf{p}(p|u,q) \rightarrow \mathsf{p}_{\mathcal{E}}(p|u,q)$$

Definition 1. *A search engine is deterministic if*

$$\mathsf{p}_{\mathcal{E}}(p|u,q) \;=\; \begin{cases} 1 & \text{si } p = p_{uq} \\ 0 & \text{everywhere else} \end{cases}$$

Otherwise, the search engine is stochastic.

Search Engine

$$\mathsf{p}_{\mathcal{L}}(p,u,q) \;=\; \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}(p|u,q)\mathsf{p}_{\mathcal{R}}(u,q)$$

The search engine determines the position given the document and the query:

$$\mathsf{p}(p|u,q) \rightarrow \mathsf{p}_{\mathcal{E}}(p|u,q)$$

$$\mathsf{p}_{\mathcal{L}}(p,u,q) \;=\; \mathsf{p}_{\mathcal{I}}(p)\mathsf{p}_{\mathcal{E}}(p|u,q)\mathsf{p}_{\mathcal{R}}(u,q)$$

Microscopic Law

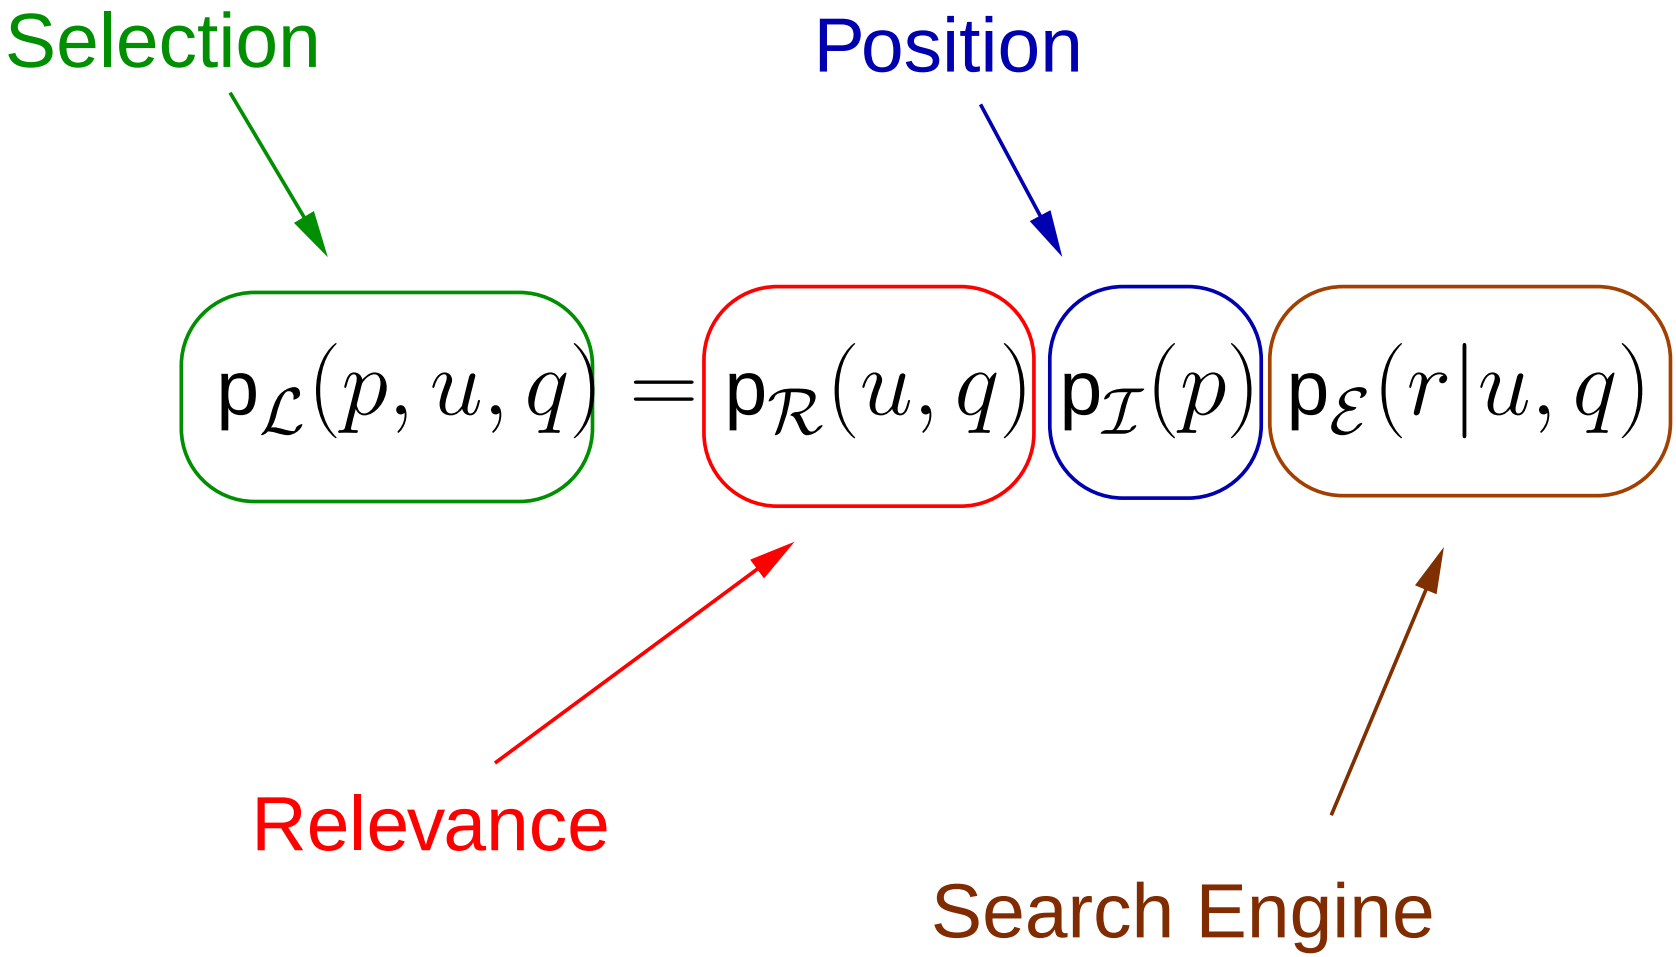

Aggregation:
From microscopic law to influence on user selections

Aggregation

To isolate the relevance effect from the position effect, we need to estimate the accumulated relevance at a position.

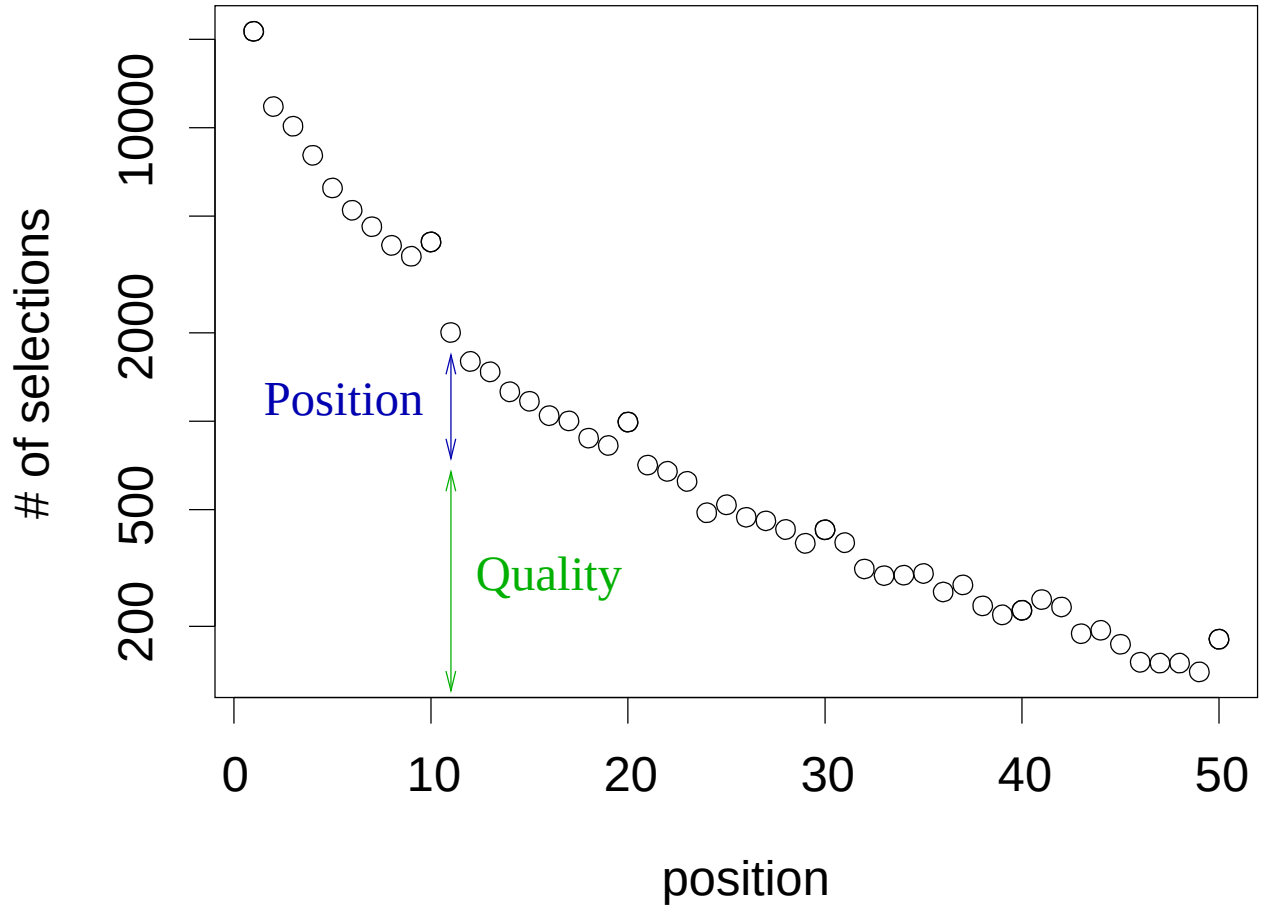

Search Engine Quality

Quality at a given position is the sum of the documents relevance at that position:

$$\mathcal{Q}_\rho = \sum_{\text{documents at } \rho} \mathsf{p}_{\mathcal{R}}(u,q)$$

Definition 2. *Quality at position ρ is*

$$\mathcal{Q}_\rho = \sum_{u,q} \mathsf{p}_{\mathcal{R}}(u,q)\mathsf{p}_{\mathcal{E}}(\rho|u,q)$$

Selections at ρ

$$\mathsf{p}_{\mathcal{L}}(\rho, u, q) = \mathsf{p}_{\mathcal{R}}(u, q)\mathsf{p}_{\mathcal{I}}(\rho)\mathsf{p}_{\mathcal{E}}(\rho|u, q)$$

$$\begin{aligned}
\mathcal{L}_\rho &= \sum_{u,q} \mathsf{p}_{\mathcal{L}}(\rho, u, q) \\
&= \mathsf{p}_{\mathcal{I}}(\rho) \sum_{u,q} \mathsf{p}_{\mathcal{R}}(u, q)\mathsf{p}_{\mathcal{E}}(\rho|u, q) \\
&= \mathsf{p}_{\mathcal{I}}(\rho)\mathcal{Q}_\rho
\end{aligned}$$

Macroscopic Law

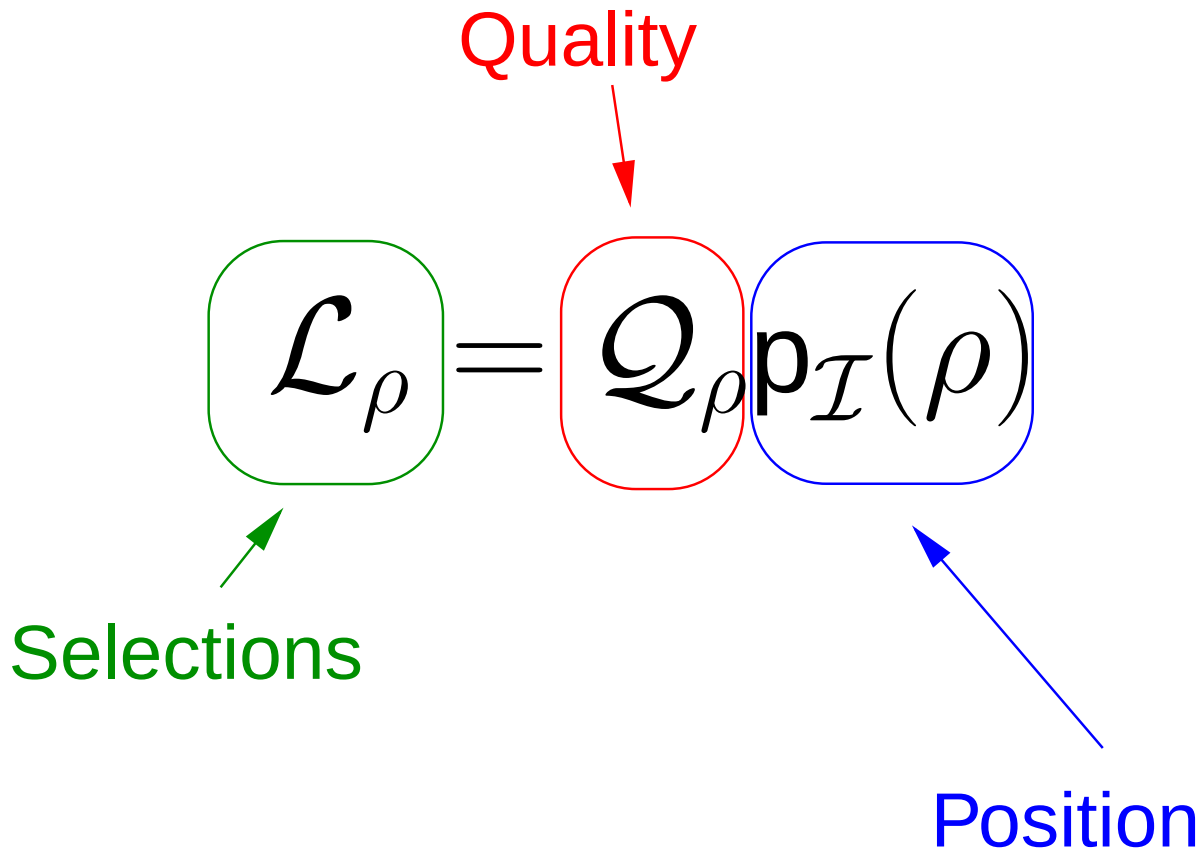

Placebo vs. Deterministic Engines

Deterministic Relevance:

$$\mathsf{p}_{\mathcal{R}}(u,q) \;=\; \frac{\mathcal{Q}_\rho}{\mathcal{L}_\rho}\mathsf{p}_{\mathcal{L}}(\rho,u,q)$$

Placebo Relevance:

$$\mathsf{p}_{\mathcal{R}}(u,q) \;\div\; \frac{1}{\mathcal{L}_\rho}\mathsf{p}_{\mathcal{L}}(\rho,u,q)$$

Model Estimation

$$\mathsf{p}_{\mathcal{L}}(\rho, u, q) = \mathsf{p}_{\mathcal{R}}(u, q)\mathsf{p}_{\mathcal{I}}(\rho)\mathsf{p}_{\mathcal{E}}(\rho|u, q)$$

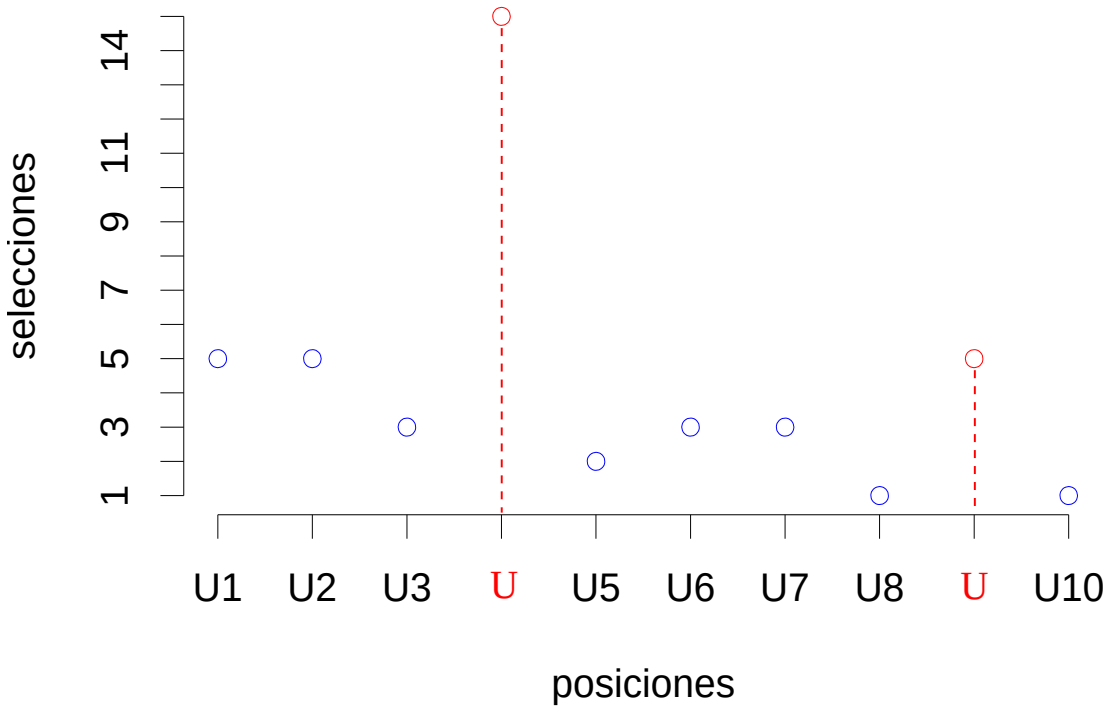

position 4:

$$\mathsf{p}_{\mathcal{L}}(4, u, q) \div 15 = \frac{1}{2}\mathsf{p}_{\mathcal{R}}(u, q)\mathsf{p}_{\mathcal{I}}(4)$$

and position 9:

$$\mathsf{p}_{\mathcal{L}}(9, u, q) \div 5 = \frac{1}{2}\mathsf{p}_{\mathcal{R}}(u, q)\mathsf{p}_{\mathcal{I}}(9)$$

Model Estimation

$$\frac{\mathsf{p}_{\mathcal{I}}(\rho)}{\sum_{r\in\mathcal{R}}\mathsf{p}_{\mathcal{I}}(r)} = \left(\sum_{r\in\mathcal{R}}\frac{A_{r\rho}}{A_{\rho r}}\right)^{-1}$$

where $\mathcal{R}$ is the set of stochastic positions, $\mathcal{D}_{r\rho}$ is the set of documents apearing at positions r and ρ with a positive probability and

$$A_{r\rho} = \sum_{q,u\in\mathcal{D}_{r\rho}}\frac{\mathsf{p}_{\mathcal{L}}(r,u,q)}{\mathsf{p}_{\mathcal{E}}(\rho|u,q)}.$$

Todocl: Document Clustering

Selections $\mathcal{L}_r$

Position Effect $\mathrm{p}_{\mathcal{I}}r$

Quality $\mathcal{Q}_r$

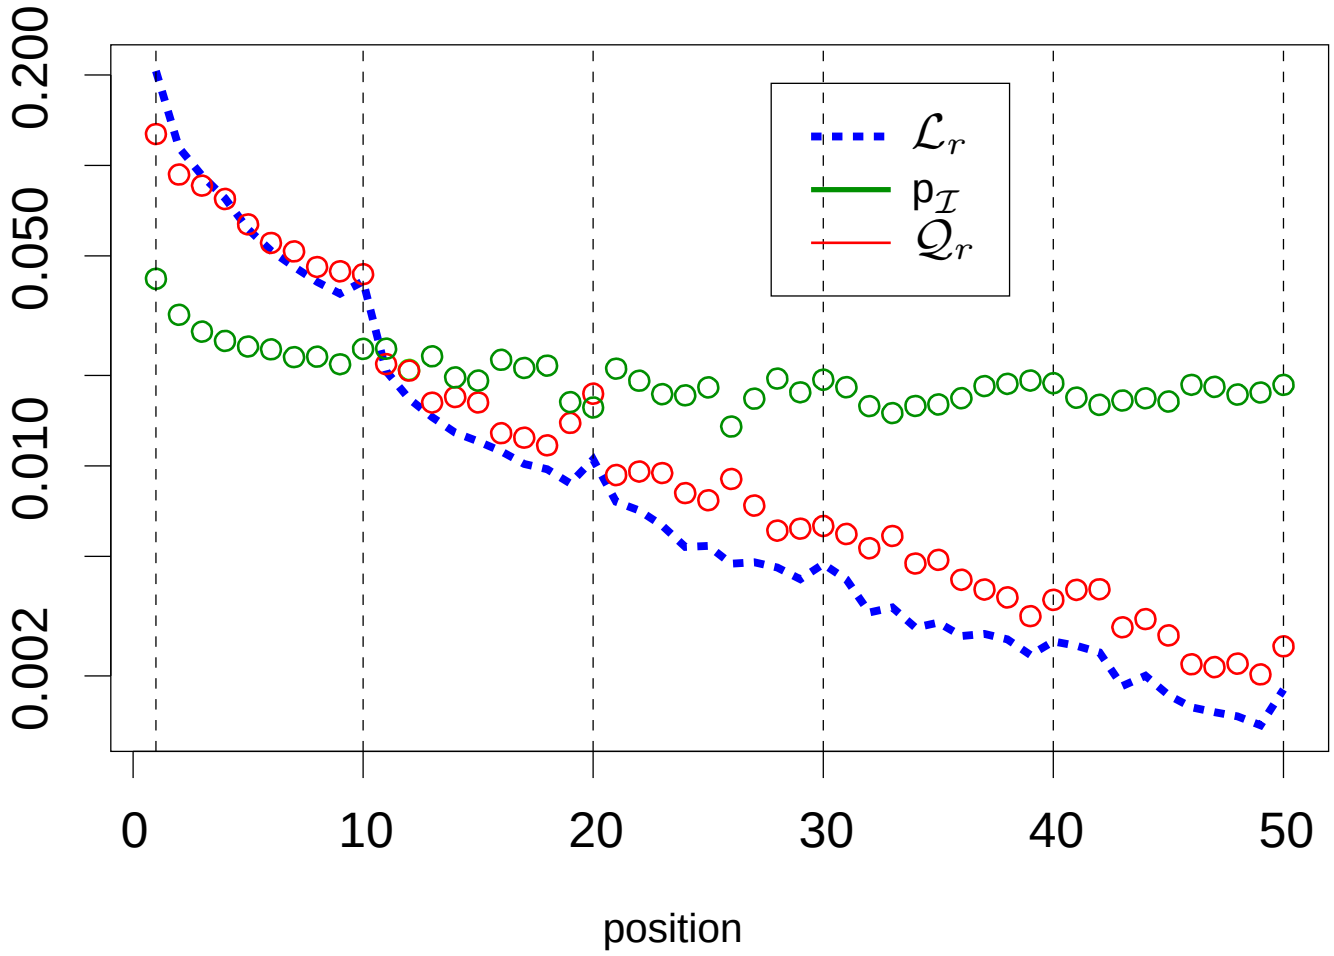

Todocl: Document Clustering

Selections $\mathcal{L}_r$

Position Effect $\mathsf{p}_{\mathcal{I}}r$

Quality $\mathcal{Q}_r$

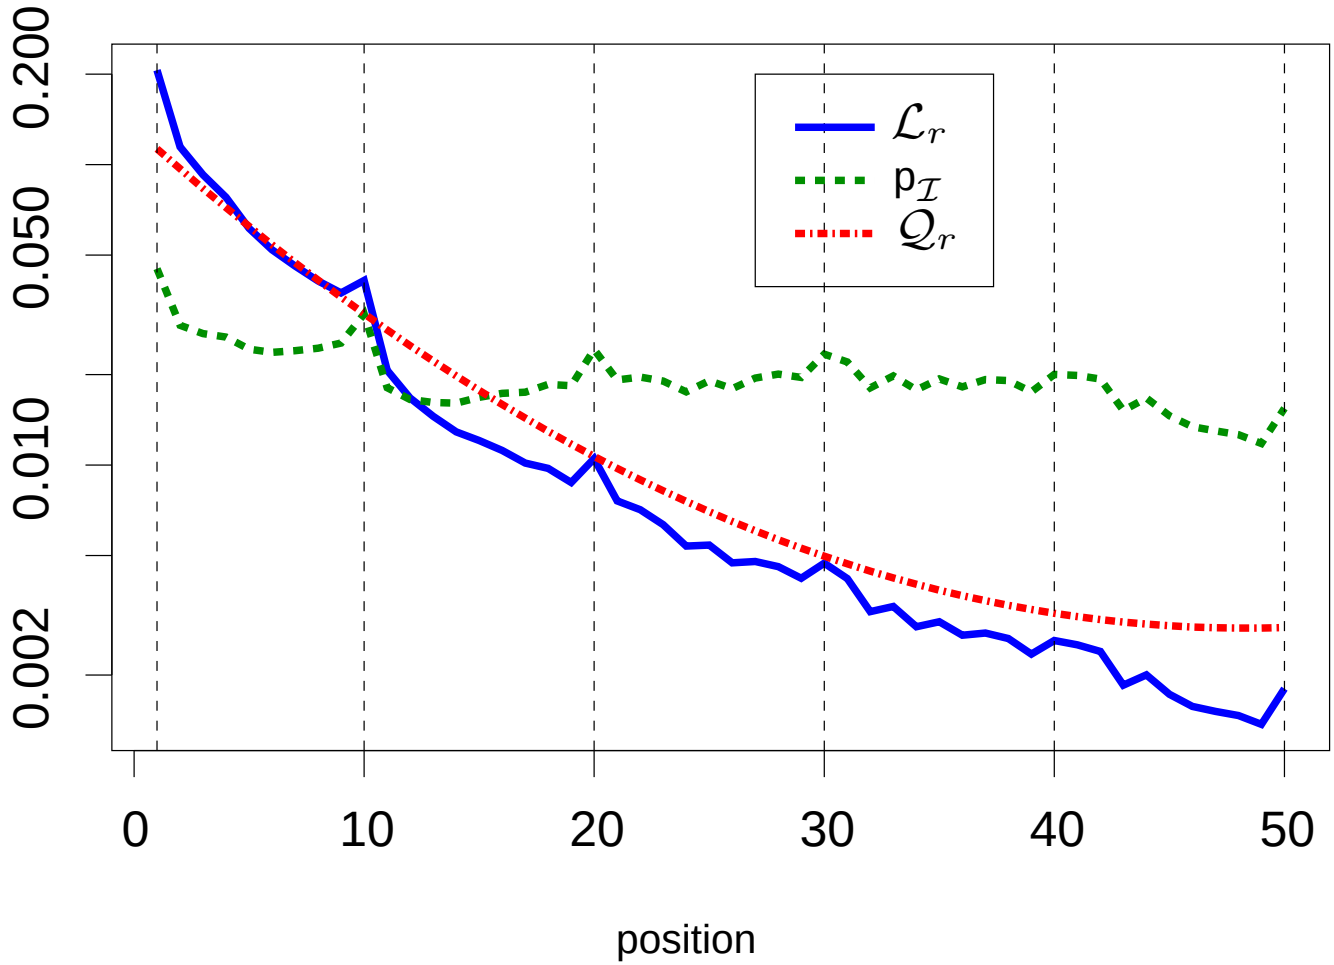

Example: Query for “Planos”

Selections: $p_{\mathcal{L}}(r, u, q)$ Position $p_{\mathcal{I}}(r)$ Relevance $p_{\mathcal{R}}(u, q)$

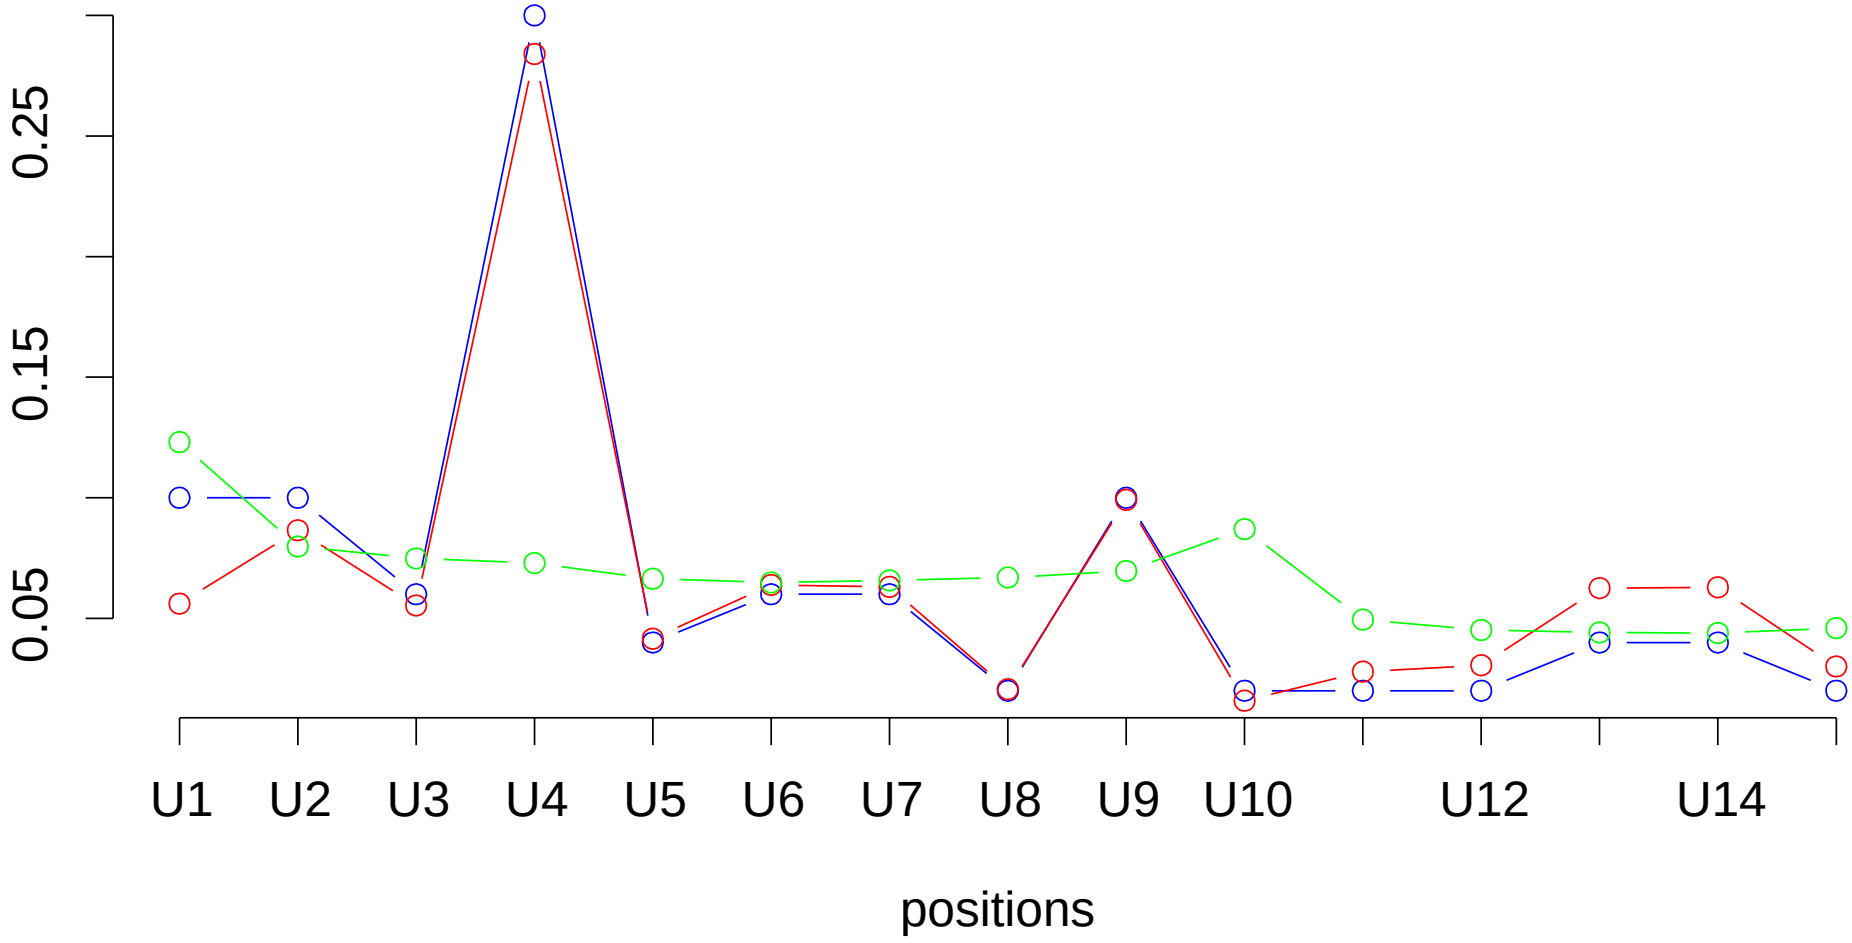

Example: Query for “Planos”

Selections: $p_{\mathcal{L}}(r, u, q)$ Position $p_{\mathcal{I}}(r)$ Relevance $p_{\mathcal{R}}(u, q)$

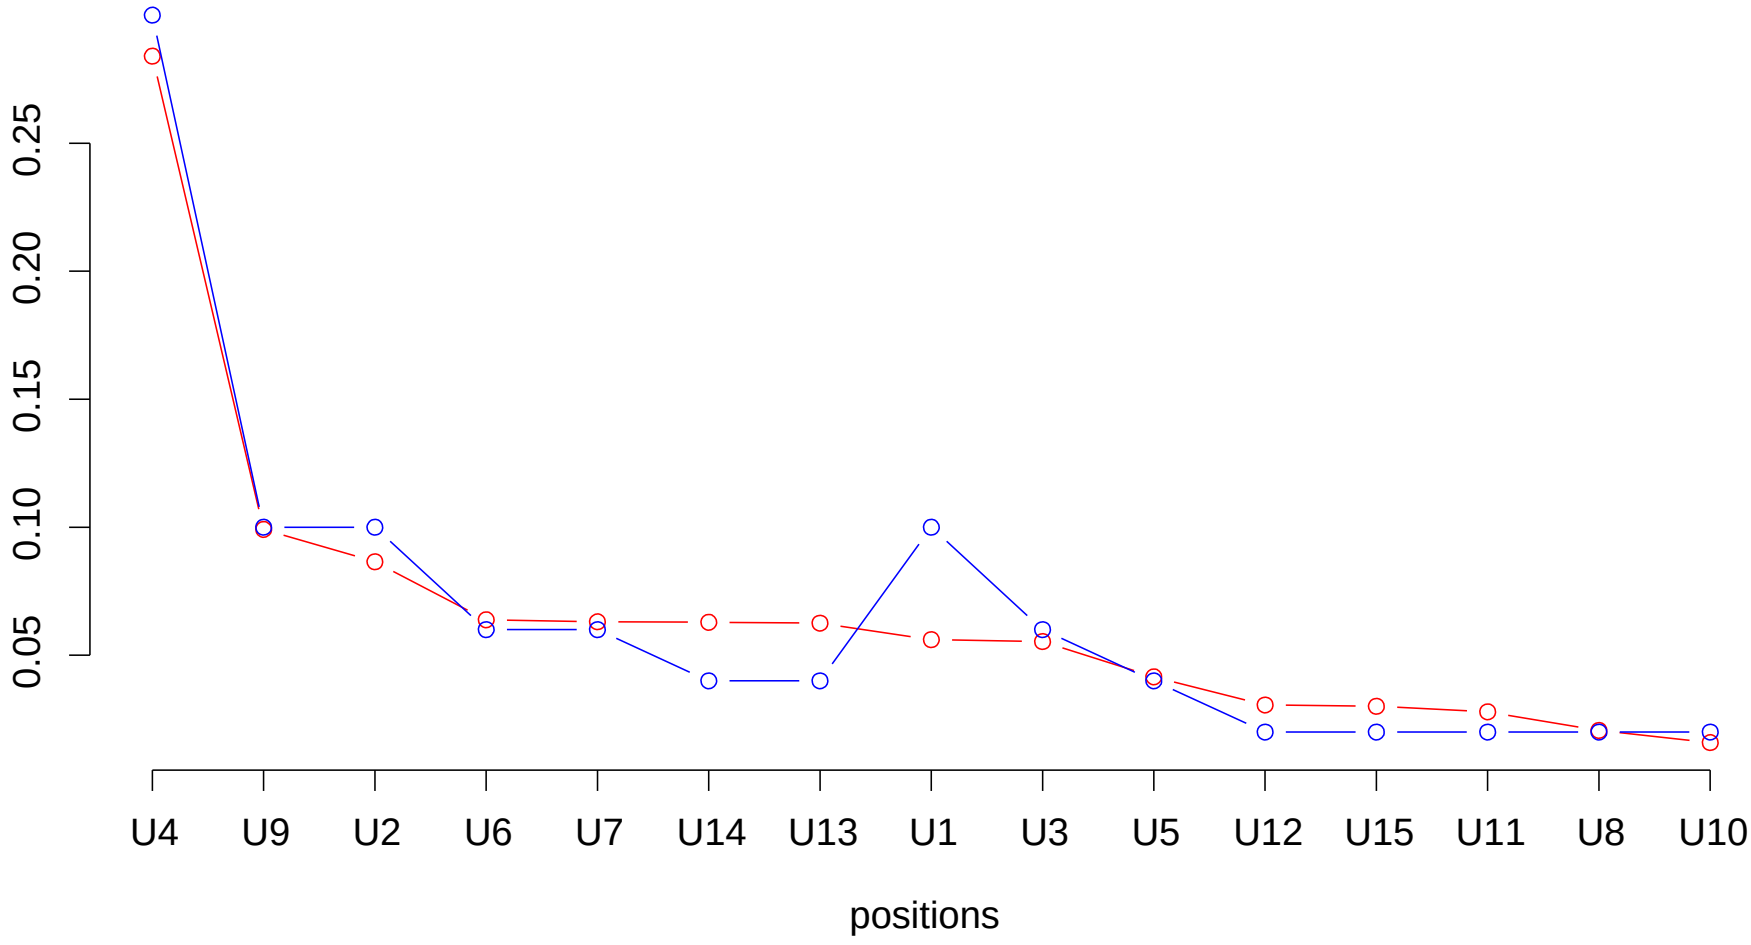

Conclusions

- Achievement:
  - We isolated relevance from position effect,
  - We obtained an empirical measure $\mathcal{Q}_\rho$ of the engine performance.

Conclusions

- Achievement:
  - We isolated relevance from position effect,
  - We obtained an empirical measure $\mathcal{Q}_\rho$ of the engine performance.
- Applications:
  - Engine improvement:
    - Re-ranking of precomputed queries,
    - Relevance feedback,
    - Engine evaluation and comparison.
  - Query clustering.
  - Interface effect evaluation.