Expert Finding using Markov Networks in Open Source Communities

Matthieu Vergne[1,2], Angelo Susi[1]
vergne@fbk.eu, susi@fbk.eu

[1]Center for Information and Communication Technology
Fondazione Bruno Kessler

[2]Doctoral School in Information and Communication Technology
University of Trento

CAISE - June 2014

Outline

Some Requirements Elicitation Contexts

The context:

You elicit requirements from:

Some Requirements Elicitation Contexts

The context:

- Personal project

You elicit requirements from:

Some Requirements Elicitation Contexts

The context:

- Personal project

You elicit requirements from:

- Nobody, code on the fly. YEEHA!

Some Requirements Elicitation Contexts

The context:

- Project with friends
- For some users you know

You elicit requirements from:

Some Requirements Elicitation Contexts

The context:

- Project with friends
- For some users you know

You elicit requirements from:

- All the users

Some Requirements Elicitation Contexts

The context:

- Project with friends
- For some users you know

You elicit requirements from:

- All the users
- Your friends (align everyone)

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

Some Requirements Elicitation Contexts

The context:

- Project with some new colleagues
- For dozens of users

You elicit requirements from:

Some Requirements Elicitation Contexts

The context:

- Project with some new colleagues
- For dozens of users

You elicit requirements from:

- All the users (hard, but feasible)

Some Requirements Elicitation Contexts

The context:

- Project with some new colleagues
- For dozens of users

You elicit requirements from:

- All the users (hard, but feasible)
- Your colleagues (fix conflicts + misunderstandings)

Some Requirements Elicitation Contexts

The context:

- Project with hundreds of new colleagues
- For hundreds of users

You elicit requirements from:

Some Requirements Elicitation Contexts

The context:

- Project with hundreds of new colleagues
- For hundreds of users

You elicit requirements from:

- HR

Some Requirements Elicitation Contexts

The context:

- Project with hundreds of new colleagues
- For hundreds of users

You elicit requirements from:

- HR to get a new job...

Some Requirements Elicitation Contexts

The context:

- Project with hundreds of new colleagues
- For hundreds of users

You elicit requirements from:

- HR to get a new job...

Funny but impossible?

Some Requirements Elicitation Contexts

The context:

- Project with hundreds of new colleagues
- For hundreds of users

You elicit requirements from:

- HR to get a new job...

Funny but impossible? This is what you have to face in OSS communities:

- Many anonymous contributors
- Many anonymous users

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

Where do We Start From?

Two main approaches in RE:

Where do We Start From?

Two main approaches in RE:

- Look at people's production (e.g. posts in forums)
  - *Castro-Herrera and Cleland-Huang [2009]*

Where do We Start From?

Two main approaches in RE:

- Look at people's production (e.g. posts in forums)
  - *Castro-Herrera and Cleland-Huang [2009]*
- Look at social relationships (e.g. roles, influence)
  - *Lim et al. [2010]*

Where do We Start From?

Two main approaches in RE:

- Look at people's production (e.g. posts in forums)
  - *Castro-Herrera and Cleland-Huang [2009]*
- Look at social relationships (e.g. roles, influence)
  - *Lim et al. [2010]*

Applicable to OSS communities?

Where do We Start From?

Two main approaches in RE:

- Look at people's production (e.g. posts in forums)
  - *Castro-Herrera and Cleland-Huang [2009]*
- Look at social relationships (e.g. roles, influence)
  - *Lim et al. [2010]*

Applicable to OSS communities?

- production: posts, issues, commits, etc.

Where do We Start From?

Two main approaches in RE:

- Look at people's production (e.g. posts in forums)
  - *Castro-Herrera and Cleland-Huang [2009]*
- Look at social relationships (e.g. roles, influence)
  - *Lim et al. [2010]*

Applicable to OSS communities?

- production: posts, issues, commits, etc.
- social: contributors, translators, companies' jobs, etc.

General Idea

Our approach:

- Goal: improve expert finding in RE by combining production-based and social-based perspectives.

General Idea

Our approach:

- Goal: improve expert finding in RE by combining production-based and social-based perspectives.
- Concepts: stakeholders, roles, topics and terms
  - *Castro-Herrera and Cleland-Huang [2010] + Lim et al. [2010]*

General Idea

Our approach:

- Goal: improve expert finding in RE by combining production-based and social-based perspectives.
- Concepts: stakeholders, roles, topics and terms
  - *Castro-Herrera and Cleland-Huang [2010] + Lim et al. [2010]*
- Instances: extracted from available sources of data
  - *forum posts, official documents, organisational models, etc.*

General Idea

Our approach:

- Goal: improve expert finding in RE by combining production-based and social-based perspectives.
- Concepts: stakeholders, roles, topics and terms
  - *Castro-Herrera and Cleland-Huang [2010] + Lim et al. [2010]*
- Instances: extracted from available sources of data
  - *forum posts, official documents, organisational models, etc.*
- Computation: Markov networks to get expertise probabilities

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

General Idea

Our approach:

- Goal: improve expert finding in RE by combining production-based and social-based perspectives.
- Concepts: stakeholders, roles, topics and terms
  - *Castro-Herrera and Cleland-Huang [2010] + Lim et al. [2010]*
- Instances: extracted from available sources of data
  - *forum posts, official documents, organisational models, etc.*
- Computation: Markov networks to get expertise probabilities
- Outcome: ranking of probable experts to recommend

Concepts and Relations

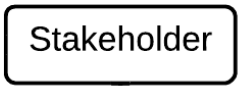

Concepts and Relations

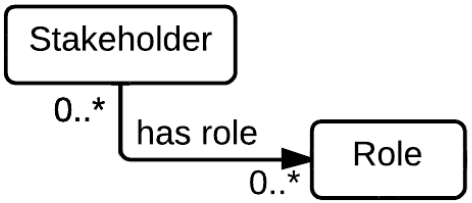

Concepts and Relations

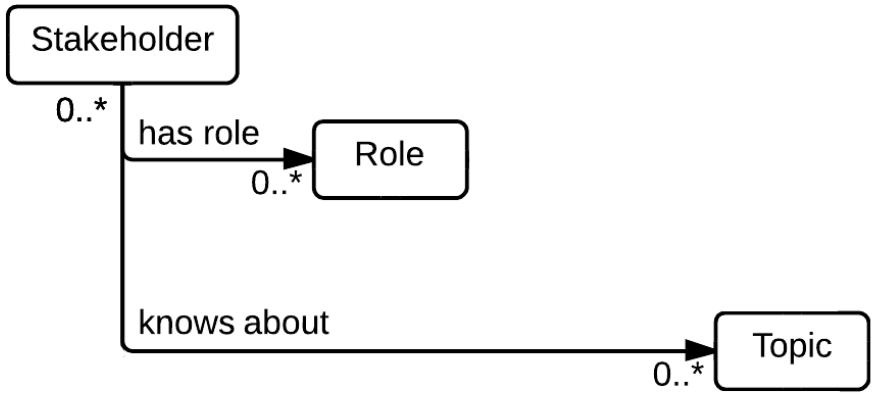

Concepts and Relations

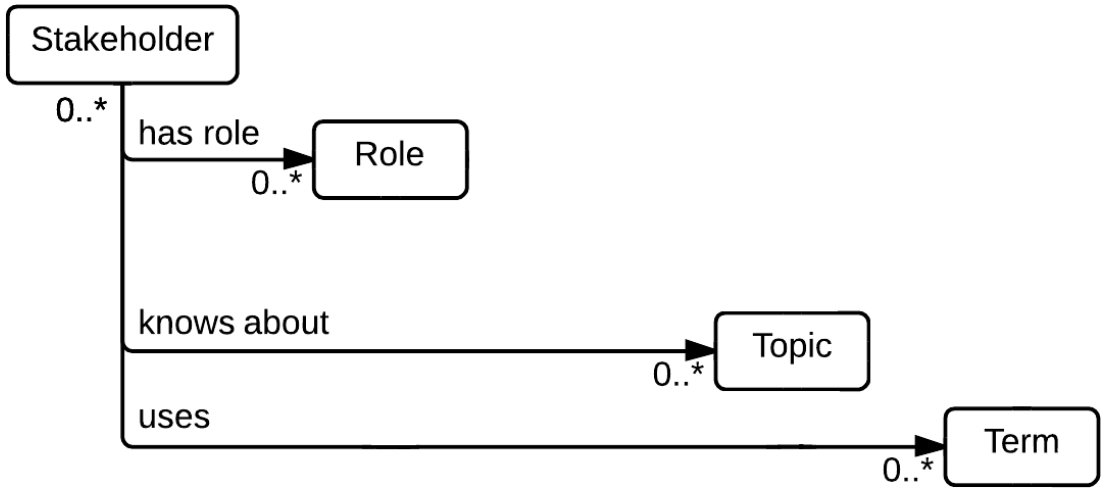

Concepts and Relations

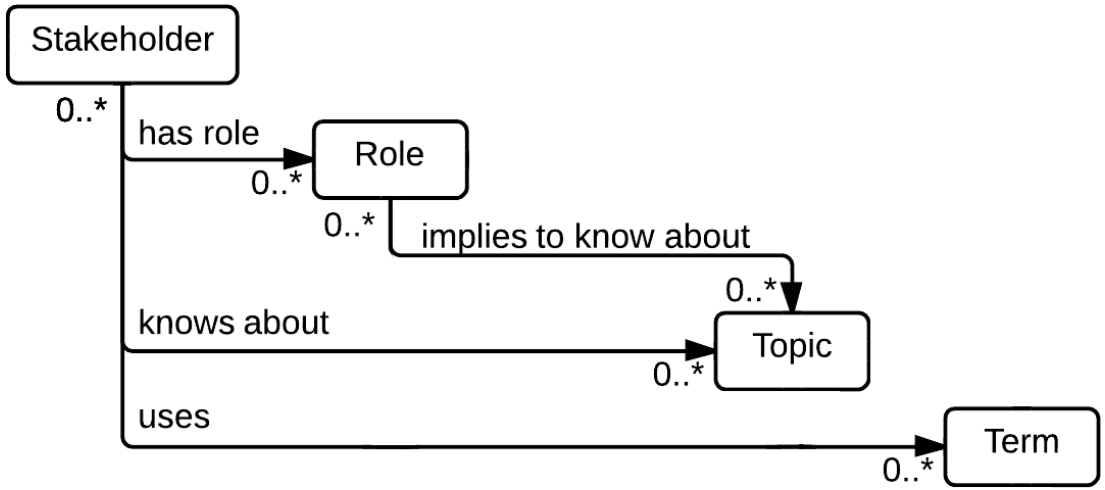

Concepts and Relations

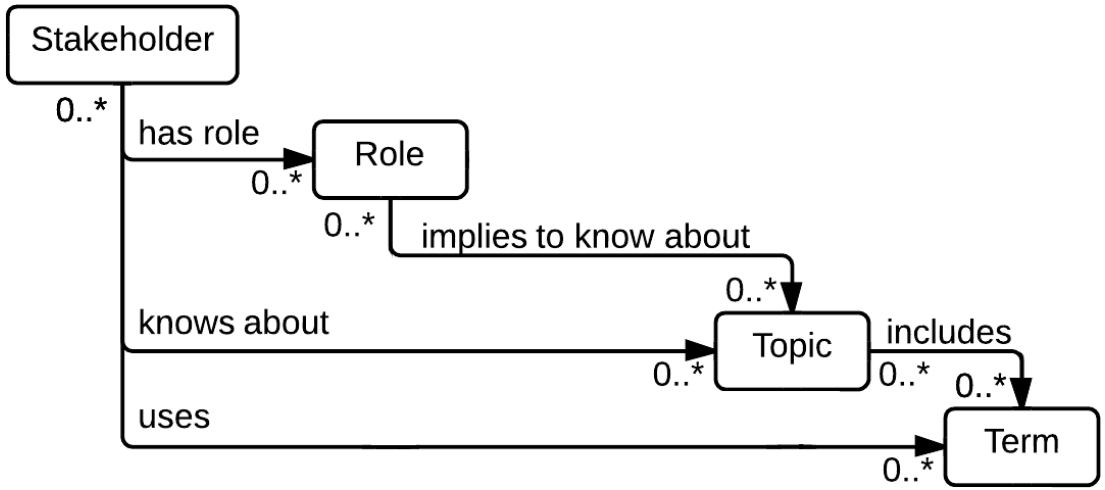

Concepts and Relations

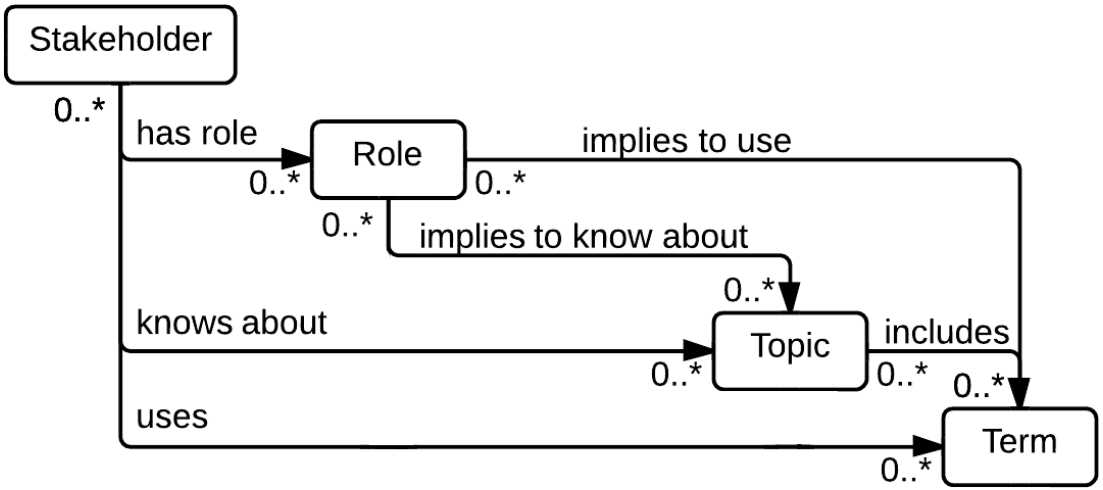

Concepts and Relations

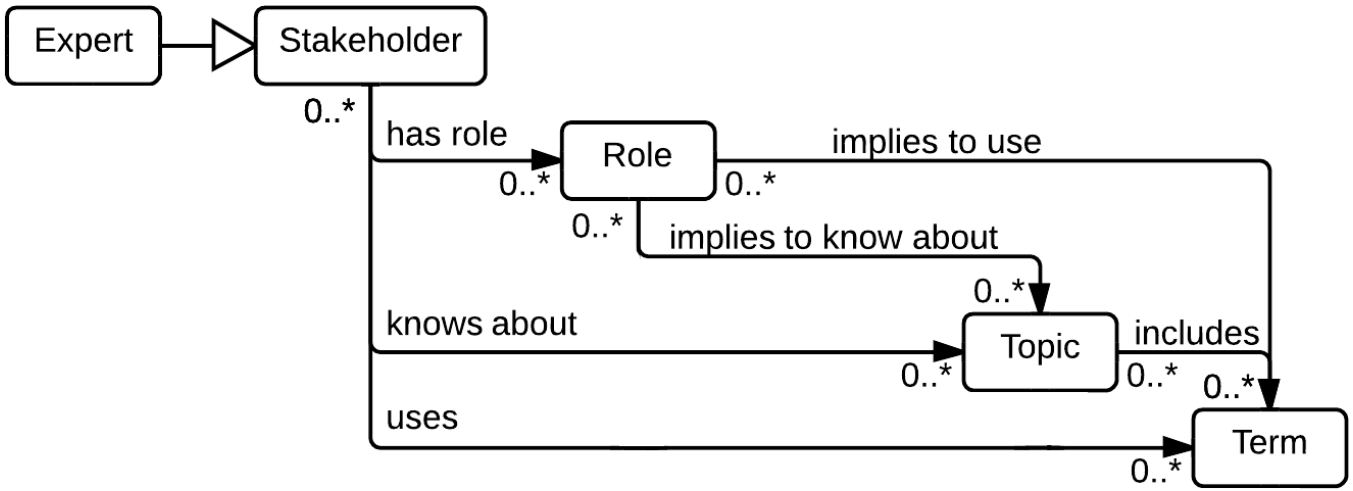

Concepts and Relations

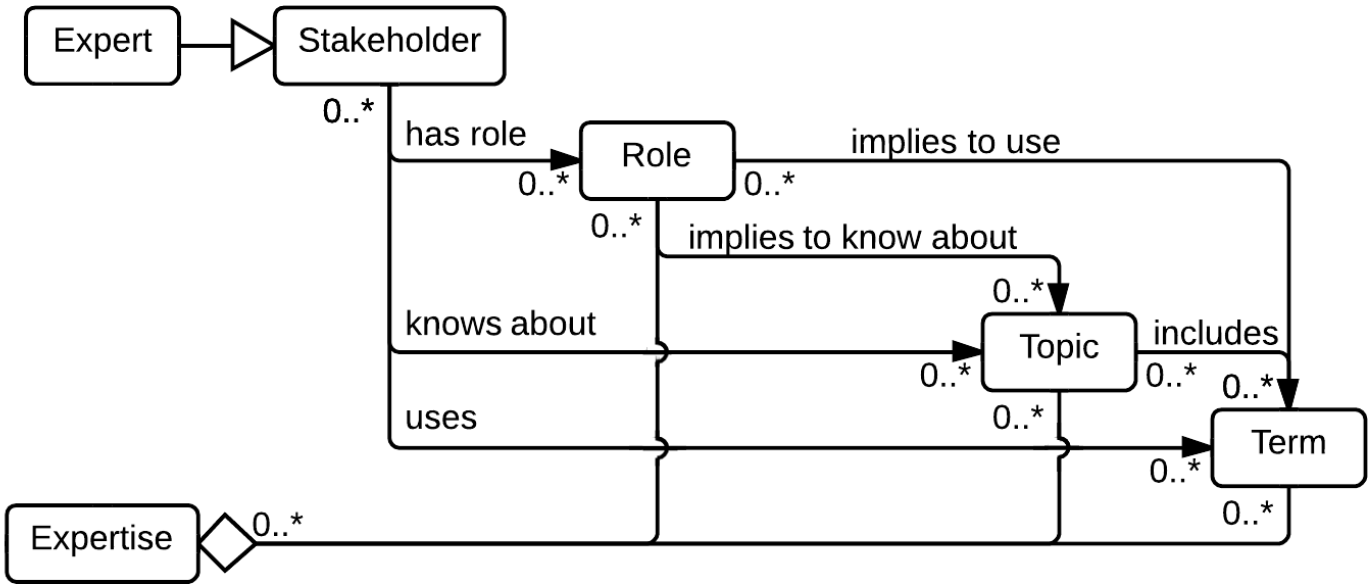

Concepts and Relations

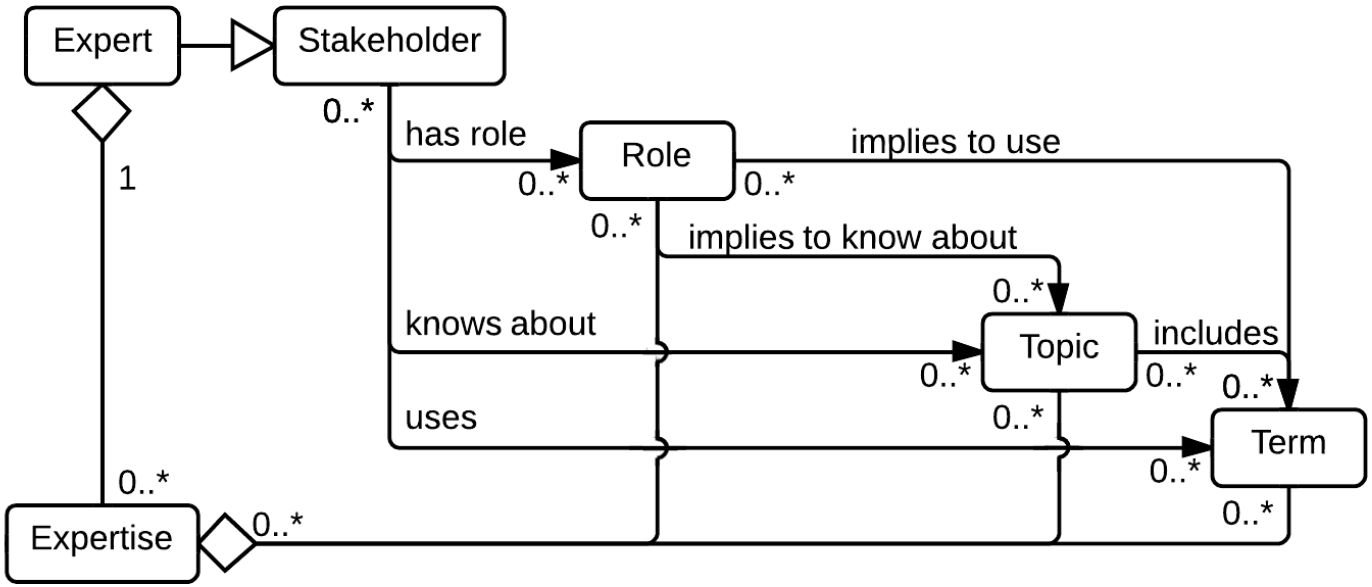

Concepts and Relations

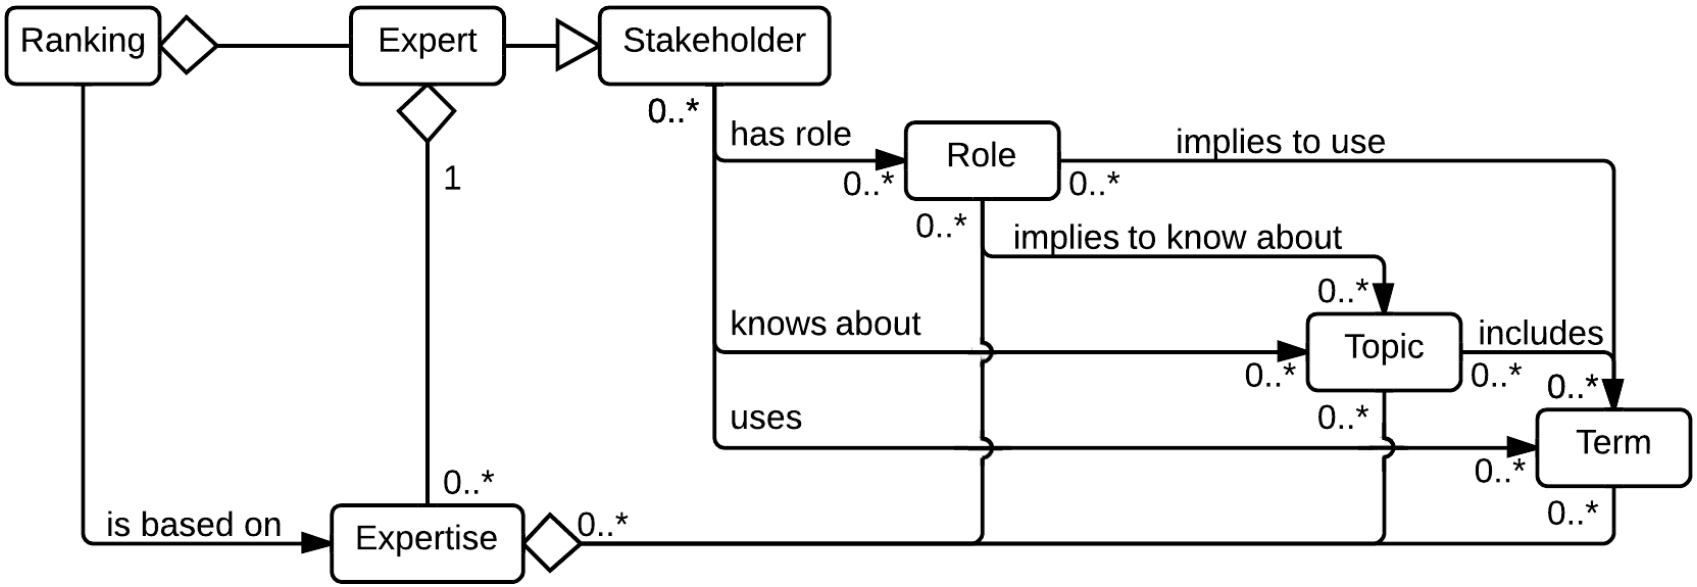

Concepts and Relations

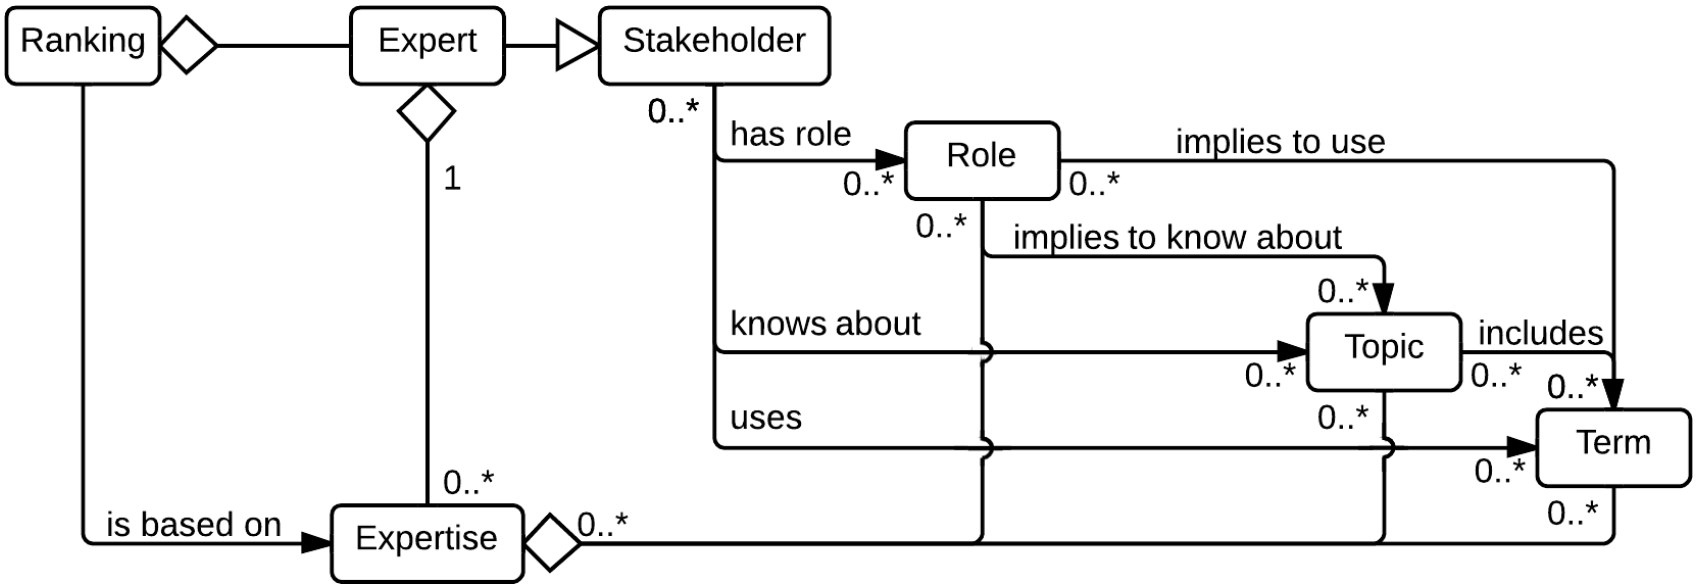

Remark: relative expertise: $A > B$, not $A = \mathit{level}$.

Full Process

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.

Full Process

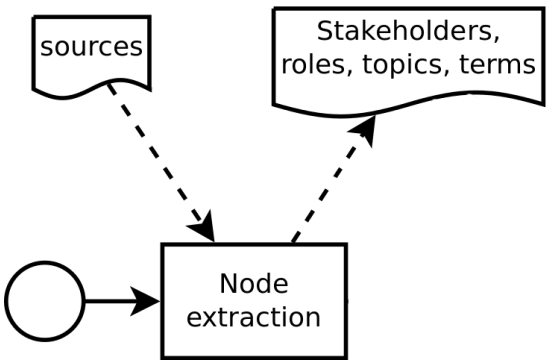

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.

Full Process

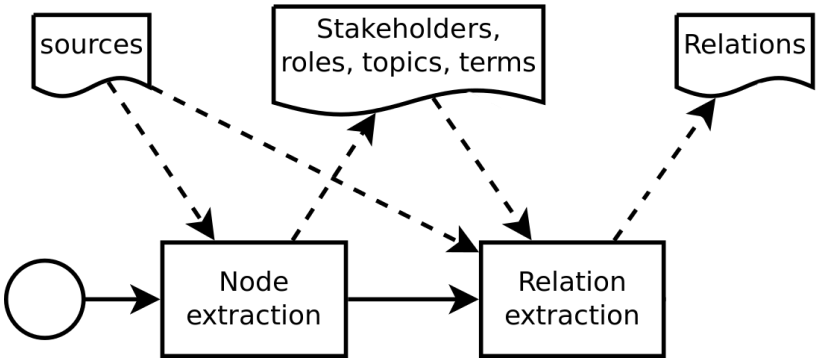

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.

Full Process

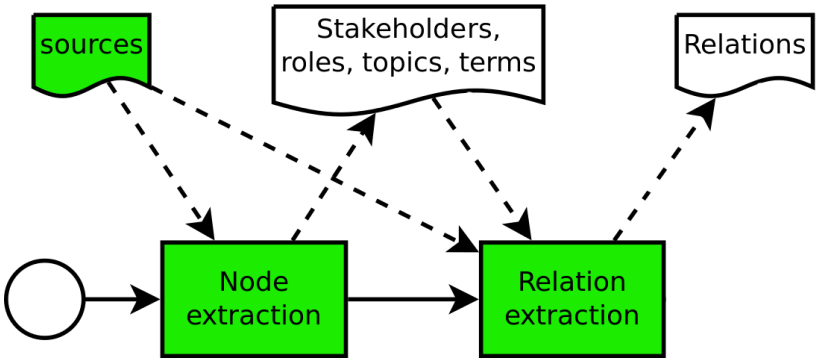

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.
- Context-specific: sources, node/relation extractors

Full Process

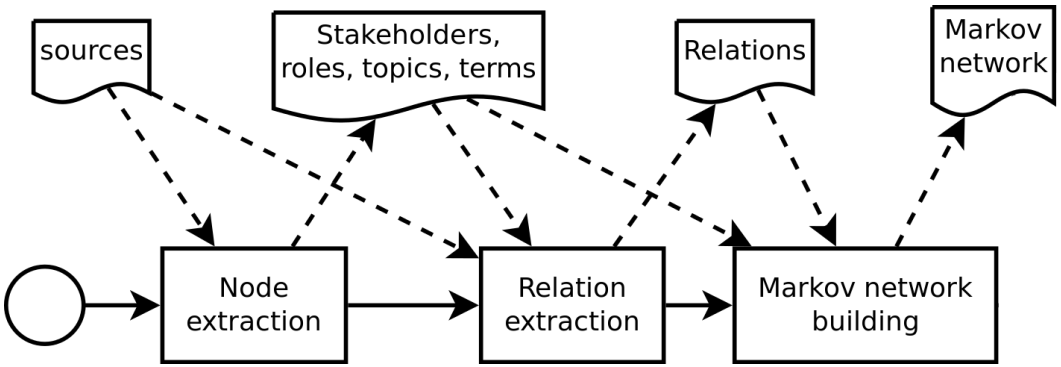

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.
- Context-specific: sources, node/relation extractors

Full Process

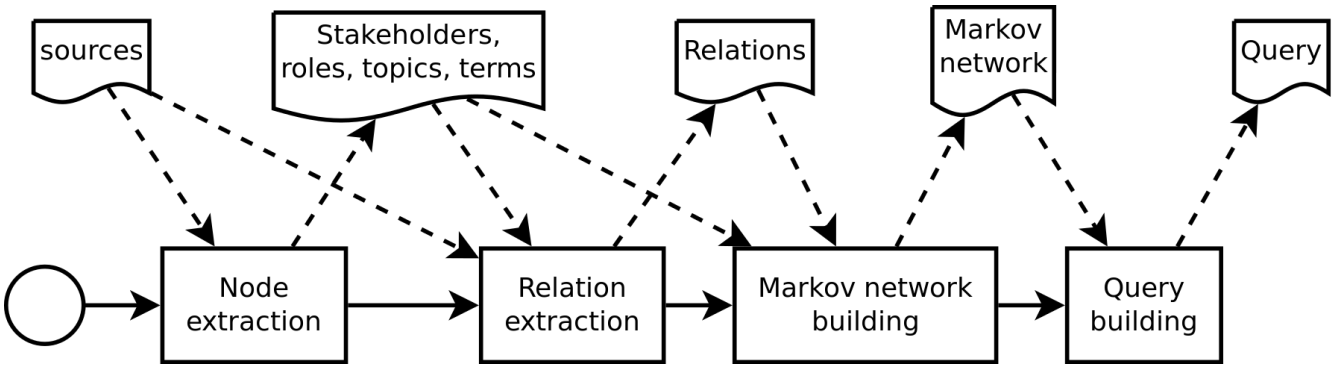

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.
- Context-specific: sources, node/relation extractors

Full Process

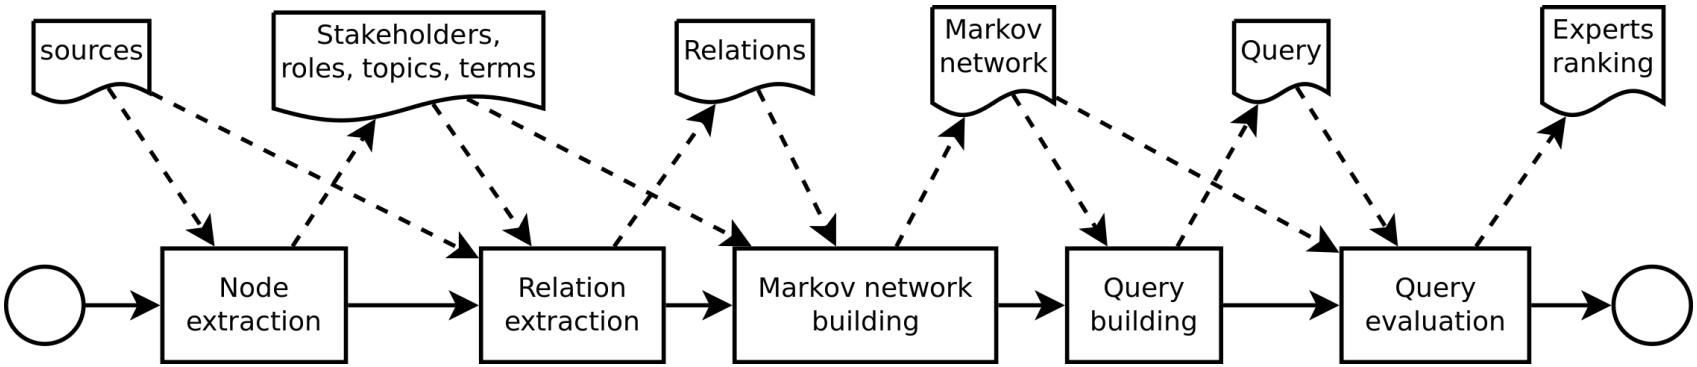

- Sources: forum posts, e-mails, reports, goal-models, social networks, etc.
- Context-specific: sources, node/relation extractors

Model: Complete 4-Partite Weighted Graph

Conceptual level

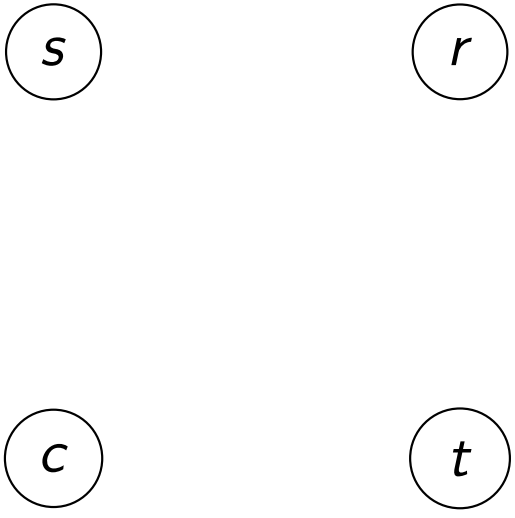

Model: Complete 4-Partite Weighted Graph

Conceptual level

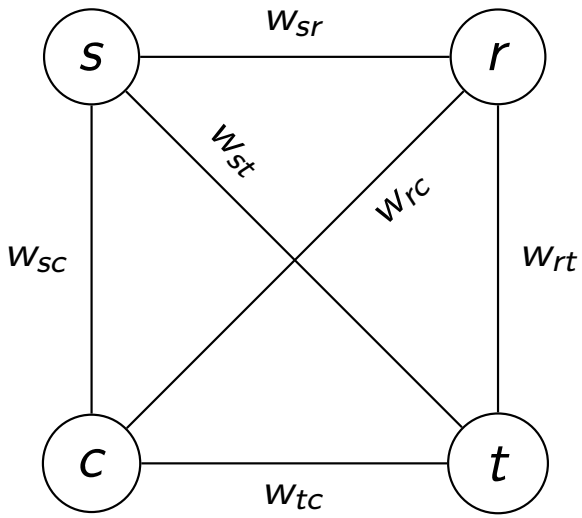

Model: Complete 4-Partite Weighted Graph

Conceptual level and instance level (2 nodes per type):

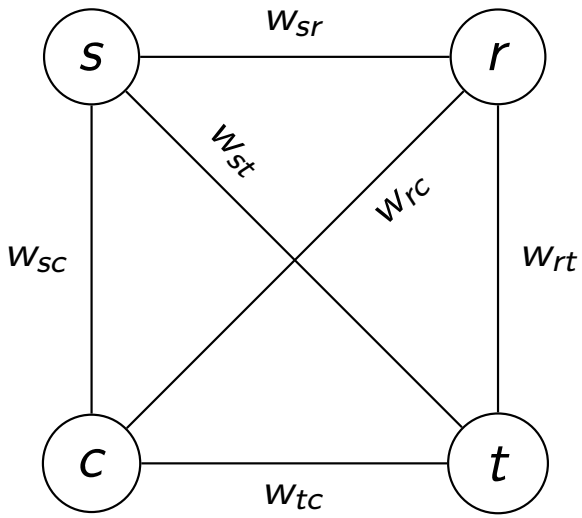

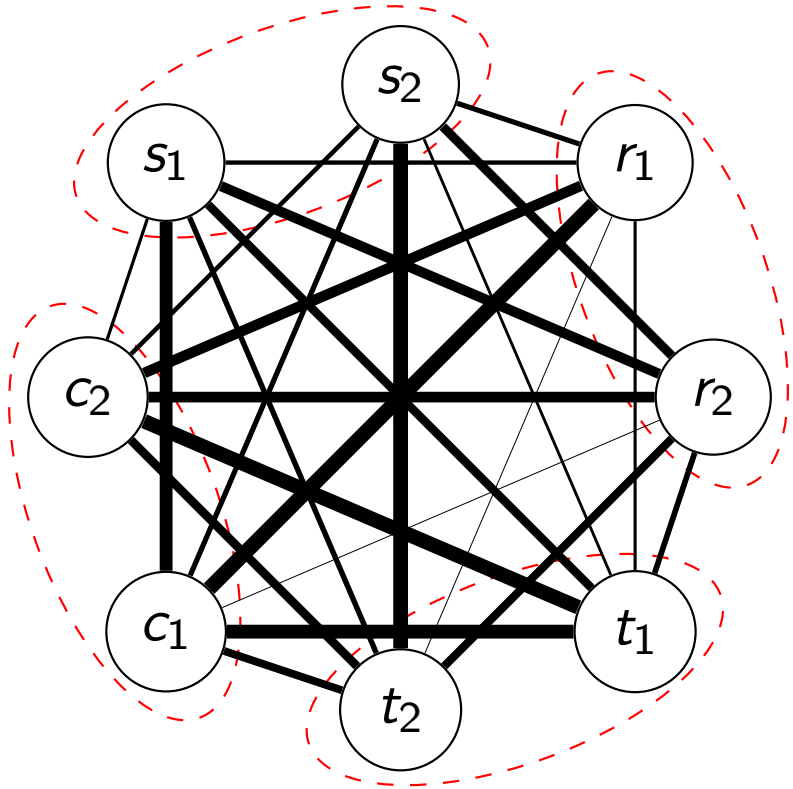

Weighting Policies

Which value for the weights?

- Amount of evidence: $w_{xy} \in \mathbb{R}^+$

Weighting Policies

Which value for the weights?

- Amount of evidence: $w_{xy} \in \mathbb{R}^+$
- $w_{xy} = 0 \Rightarrow$ no evidence

Weighting Policies

Which value for the weights?

- Amount of evidence: $w_{xy} \in \mathbb{R}^+$
- $w_{xy} = 0 \Rightarrow$ no evidence
- $w_{xy} = 5, w_{ab} = 10 \Rightarrow$ evidence for ab = 2x evidence for xy

Weighting Policies

Which value for the weights?

- Amount of evidence: $w_{xy} \in \mathbb{R}^+$
- $w_{xy} = 0 \Rightarrow$ no evidence
- $w_{xy} = 5, w_{ab} = 10 \Rightarrow$ evidence for ab = 2x evidence for xy
- Actual value depends on the interpretation of evidence
  - *Lim et al. [2010]: salience elicited from stakeholders*
  - *Castro-Herrera and Cleland-Huang [2010]: normalized term frequencies*

Weighting Policies

Which value for the weights?

- Amount of evidence: $w_{xy} \in \mathbb{R}^+$
- $w_{xy} = 0 \Rightarrow$ no evidence
- $w_{xy} = 5, w_{ab} = 10 \Rightarrow$ evidence for ab = 2x evidence for xy
- Actual value depends on the interpretation of evidence
  - *Lim et al. [2010]: salience elicited from stakeholders*
  - *Castro-Herrera and Cleland-Huang [2010]: normalized term frequencies*
- Challenge: have meaningful weights

Markov Networks (Markov Random Field)

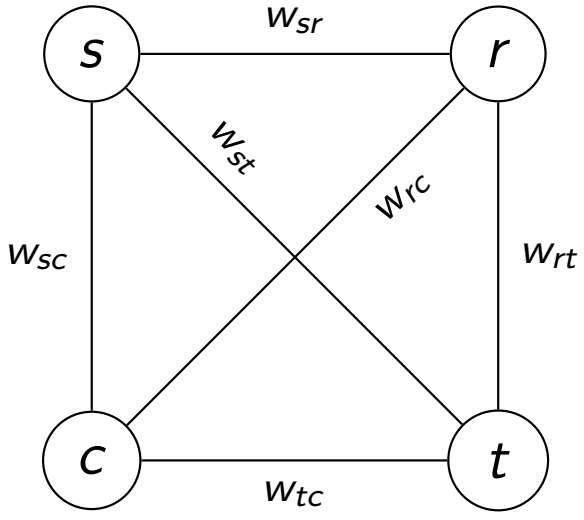

Markov Networks (Markov Random Field)

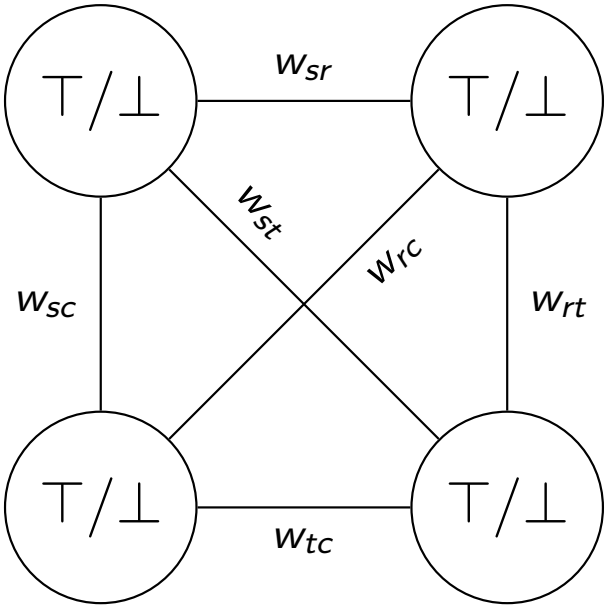

- Node $\rightarrow$ binary state

Markov Networks (Markov Random Field)

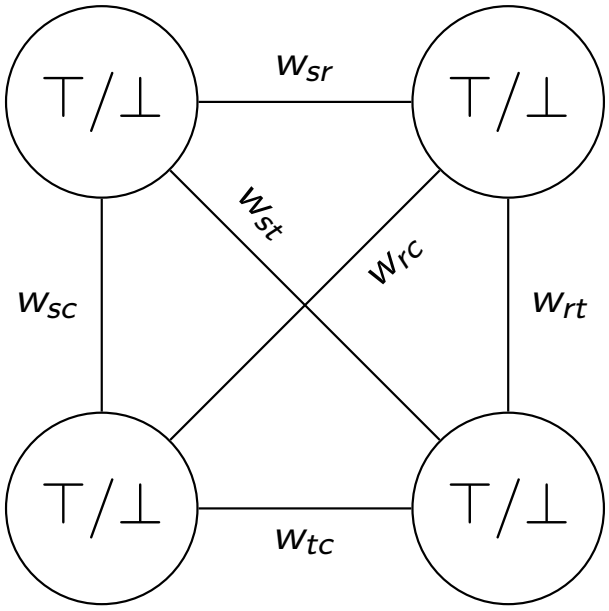

- Node $\rightarrow$ binary state

Interpretation:

- $s = \top \Rightarrow s$ is an expert
- $r/t/c = \top \Rightarrow$ looking for experts in $r/t/c$

Markov Networks (Markov Random Field)

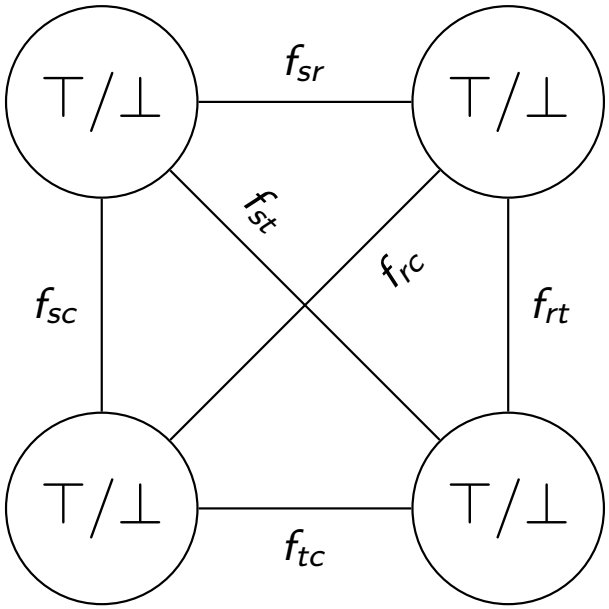

- Node → binary state
- Weight → function (depends on node states)

Interpretation:

- $s = \top \Rightarrow s$ is an expert
- $r/t/c = \top \Rightarrow$ looking for experts in $r/t/c$

Markov Networks (Markov Random Field)

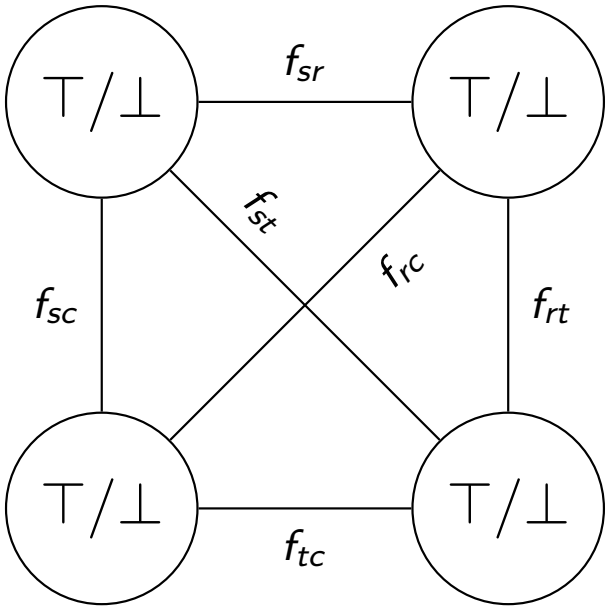

- Node → binary state
- Weight → function (depends on node states)

Interpretation:

- $s = \top \Rightarrow s$ is an expert
- $r/t/c = \top \Rightarrow$ looking for experts in $r/t/c$

Computation:

- Configuration probability: $P(n_1 = \sigma_1, ..., n_N = \sigma_N)$
- Partial + conditional: $P(\{n_i = \sigma_i\}|\{n_j = \sigma_j\})$

Experts Ranking with MN

Use for expert ranking:

- Query: $P(s_i = \top | t_{cryptography} = \top)$

Experts Ranking with MN

Use for expert ranking:

- Query: $P(s_i = \top | t_{cryptography} = \top)$
- Ranking: sort from most to least probable experts.

Experts Ranking with MN

Use for expert ranking:

- Query: $P(s_i = \top | t_{cryptography} = \top)$
- Ranking: sort from most to least probable experts.

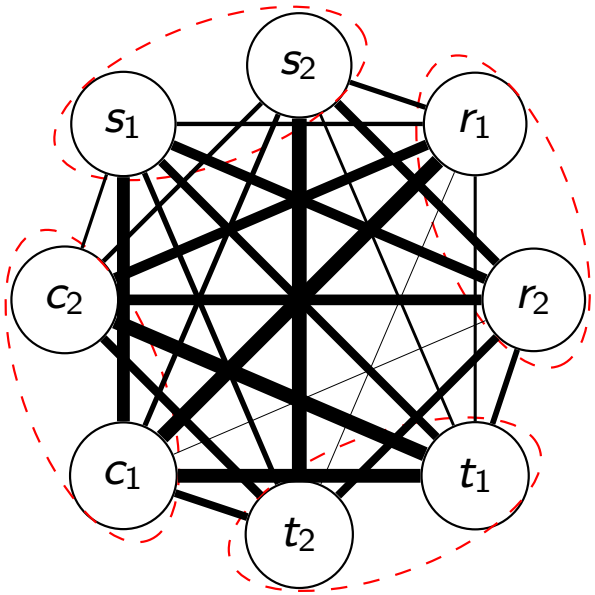

Experts Ranking with MN

Use for expert ranking:

- Query: $P(s_i = \top | t_{cryptography} = \top)$
- Ranking: sort from most to least probable experts.

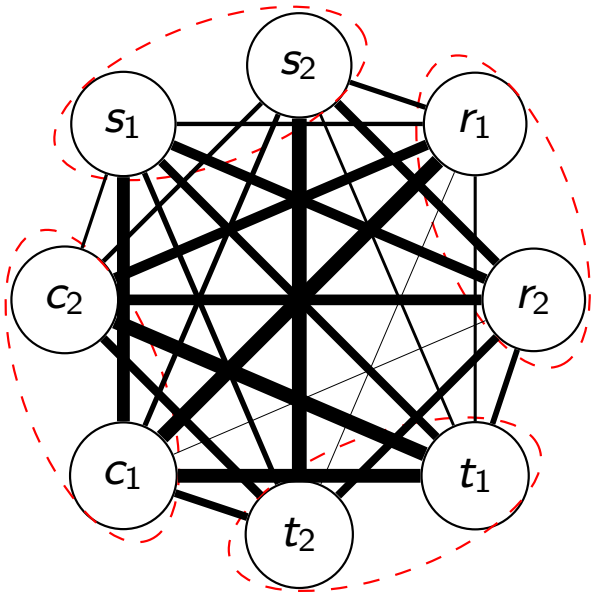

Results for t_1:

- $P(s_1|t_1) = 0.365$
- $P(s_2|t_1) = 0.834$

Experts Ranking with MN

Use for expert ranking:

- Query: $P(s_i = \top | t_{cryptography} = \top)$
- Ranking: sort from most to least probable experts.

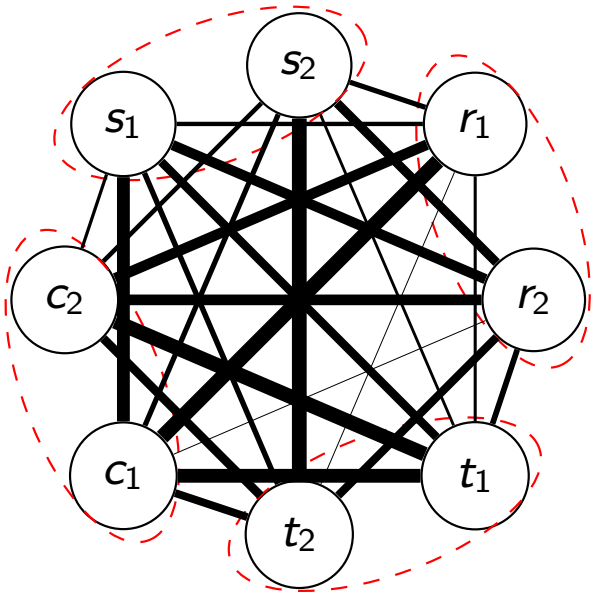

Results for t_1:

- $P(s_1|t_1) = 0.365$
- $P(s_2|t_1) = 0.834$

Ranking:

1. s_2
2. s_1

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla
- 2 topics discussed in 30 e-mails

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla
- 2 topics discussed in 30 e-mails
- Gold Standard (GS): stakeholders' retrospective evaluation

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla
- 2 topics discussed in 30 e-mails
- Gold Standard (GS): stakeholders' retrospective evaluation
- Nodes: authors (stak.), nouns (terms), title nouns (topics)

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla
- 2 topics discussed in 30 e-mails
- Gold Standard (GS): stakeholders' retrospective evaluation
- Nodes: authors (stak.), nouns (terms), title nouns (topics)
- Weights: occurrence counting in each e-mail

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla
- 2 topics discussed in 30 e-mails
- Gold Standard (GS): stakeholders' retrospective evaluation
- Nodes: authors (stak.), nouns (terms), title nouns (topics)
- Weights: occurrence counting in each e-mail

Results:

- Extraction: 3 stakeholders, 4 topics, 293 terms, 2k relations

Simple Case: Cooking in Trento

Experiment

- 3 Stakeholders: Alice, Bob and Carla
- 2 topics discussed in 30 e-mails
- Gold Standard (GS): stakeholders' retrospective evaluation
- Nodes: authors (stak.), nouns (terms), title nouns (topics)
- Weights: occurrence counting in each e-mail

Results:

- Extraction: 3 stakeholders, 4 topics, 293 terms, 2k relations

	$t =$ Europ. dessert			$t =$ Asian food		
Stak.	$P(s\|t)$	Tool	GS	$P(s\|t)$	Tool	GS
Alice	0.501	1	1	0.49975	2	2
Bob	0.500	2	1	0.49946	3	2
Carla	0.499	3	3	0.49978	1	1

Practical Case: XWiki OSS

Experiment

- XWiki mailing list (subset): 805 e-mails in 255 threads

Practical Case: XWiki OSS

Experiment

- XWiki mailing list (subset): 805 e-mails in 255 threads
- No Gold Standard, but obvious experts
  - *XWiki team, topic-specific "heroes"*

Practical Case: XWiki OSS

Experiment

- XWiki mailing list (subset): 805 e-mails in 255 threads
- No Gold Standard, but obvious experts
  - *XWiki team, topic-specific "heroes"*
- Nodes and Relations: same than before

Practical Case: XWiki OSS

Experiment

- XWiki mailing list (subset): 805 e-mails in 255 threads
- No Gold Standard, but obvious experts
  - *XWiki team, topic-specific "heroes"*
- Nodes and Relations: same than before

Results:

- 120 stakeholders, 216 topics, 5k terms, 75k relations

Practical Case: XWiki OSS

Experiment

- XWiki mailing list (subset): 805 e-mails in 255 threads
- No Gold Standard, but obvious experts
  - *XWiki team, topic-specific "heroes"*
- Nodes and Relations: same than before

Results:

- 120 stakeholders, 216 topics, 5k terms, 75k relations
- Scalability issue: good rankings but too reduced

Practical Case: XWiki OSS

Experiment

- XWiki mailing list (subset): 805 e-mails in 255 threads
- No Gold Standard, but obvious experts
  - *XWiki team, topic-specific "heroes"*
- Nodes and Relations: same than before

Results:

- 120 stakeholders, 216 topics, 5k terms, 75k relations
- Scalability issue: good rankings but too reduced
- Not in the paper: approximated computation
  - *Rankings+GS for 2 topics (14 threads)*
  - *18 stakeholders, 42 topics, 880 terms, 7k relations*
  - *Investigating evaluation measures*

Discussion

Results:

- Only preliminary evaluation, but support further research

Discussion

Results:

- Only preliminary evaluation, but support further research
- No roles, but some XWiki's resources can provide them

Discussion

Results:

- Only preliminary evaluation, but support further research
- No roles, but some XWiki's resources can provide them
- Close probabilities, but MN functions can be refined

Discussion

Results:

- Only preliminary evaluation, but support further research
- No roles, but some XWiki's resources can provide them
- Close probabilities, but MN functions can be refined

Approach:

- Similar nodes relations (e.g. synonyms)

Discussion

Results:

- Only preliminary evaluation, but support further research
- No roles, but some XWiki's resources can provide them
- Close probabilities, but MN functions can be refined

Approach:

- Similar nodes relations (e.g. synonyms)
- More structured sources (e.g. taxonomy, ontology)

Discussion

Results:

- Only preliminary evaluation, but support further research
- No roles, but some XWiki's resources can provide them
- Close probabilities, but MN functions can be refined

Approach:

- Similar nodes relations (e.g. synonyms)
- More structured sources (e.g. taxonomy, ontology)
- Challenge of merging different sources (e.g. trust)

Summary

- Goal: expert ranking support for RE in OSS communities

Summary

- Goal: expert ranking support for RE in OSS communities
- Idea: combine production-based and social-based perspectives

Summary

- Goal: expert ranking support for RE in OSS communities
- Idea: combine production-based and social-based perspectives
- Technique: compute expertise probability via MN

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

Summary

- Goal: expert ranking support for RE in OSS communities
- Idea: combine production-based and social-based perspectives
- Technique: compute expertise probability via MN
- Results: good support from preliminary experiment

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

Summary

- Goal: expert ranking support for RE in OSS communities
- Idea: combine production-based and social-based perspectives
- Technique: compute expertise probability via MN
- Results: good support from preliminary experiment
- Future works:
  - *Stronger experiment (more data, better measures)*
  - *Approach refinement (relations, MN functions)*
  - *Sources aggregation (e.g. ontologies and other models)*

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

Thanks for your attention.

Questions?

M. Vergne, A. Susi (vergne@fbk.eu, susi@fbk.eu) FBK - ICT Doctoral School Trento

Expert Finding Examples

production-based:

- Mockus and Herbsleb [2002]: written code evidence to evaluate expertise in software pieces.
- Serdyukov and Hiemstra [2008]: authors' contributions in documents to infer expertise in related topics.

social-based:

- Zhang et al. [2007]: compare algos on social network built from askers/repliers identification in online forums.

Both:

- Karimzadehgan et al. [2009] exploit relationships + e-mails content between employees of a company.

Main differences with Karimzadehgan et al. [2009]

- social relationships not only for post-processing
- inference technique manage multiple topics

Stakeholders Recommendation in RE

Literature review Mohebzada et al. [2012]

- Castro-Herrera and Cleland-Huang [2010]
  - *evaluate stakeholders knowledge through participation in forum*
  - *build abstract topics (term vectors) depending on messages common terms*
  - *relate stakeholders to topics they participate in*
  - *recommend other stakeholders to participate in new, similar topics*
  - *production-based: exploit data provided directly by stakeholders*
- StakeNet Lim et al. [2010]
  - *prioritise requirements depending on stakeholders rating*
  - *core stakeholders suggests others influencing the project*
  - *role and salience describe influence*
  - *built social network + apply measures to evaluate global influence*
  - *social-based: evaluate influence based on other stakeholders suggestions*

Markov Networks (Markov Random Field)

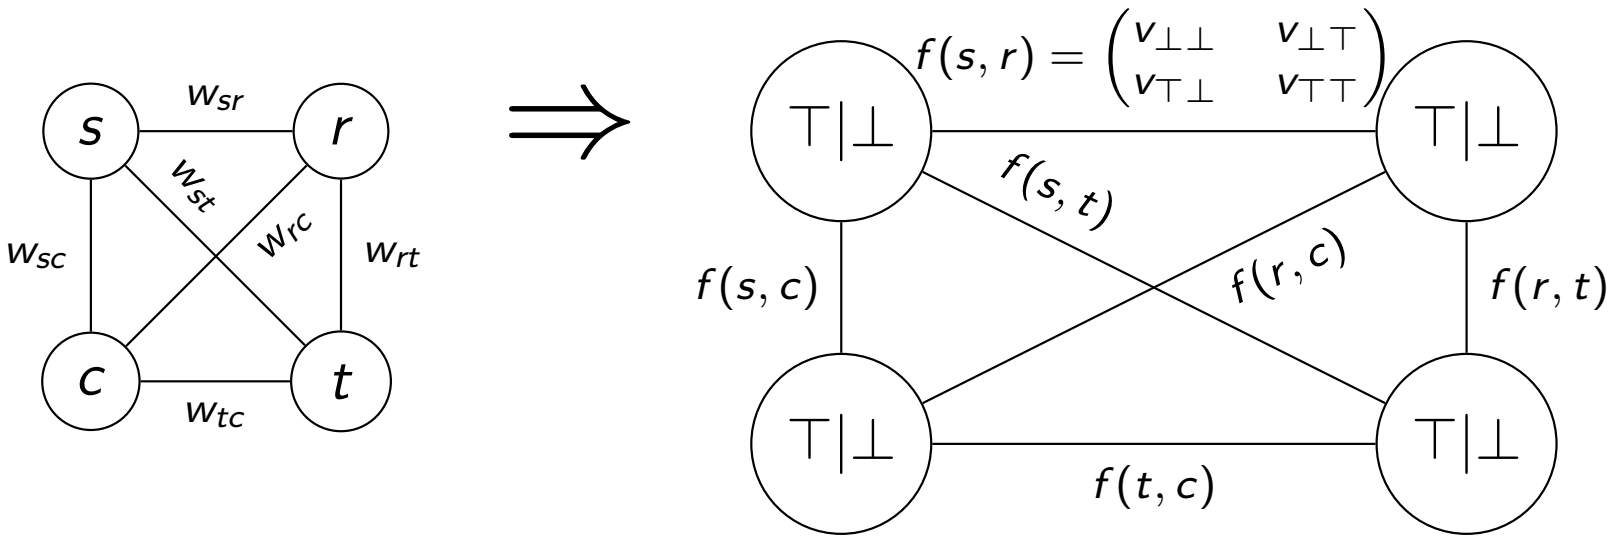

Interpretation:

- $s = \top \Rightarrow s$ is an expert
- $r/t/c = \top \Rightarrow$ looking for experts in $r/t/c$
- $f(x,y) = \begin{cases} 0 & v_{\perp\perp}, v_{\perp\top}, v_{\top\perp} \\ w_{xy} & v_{\top\top} \end{cases}$

Markov Networks (Markov Random Field)

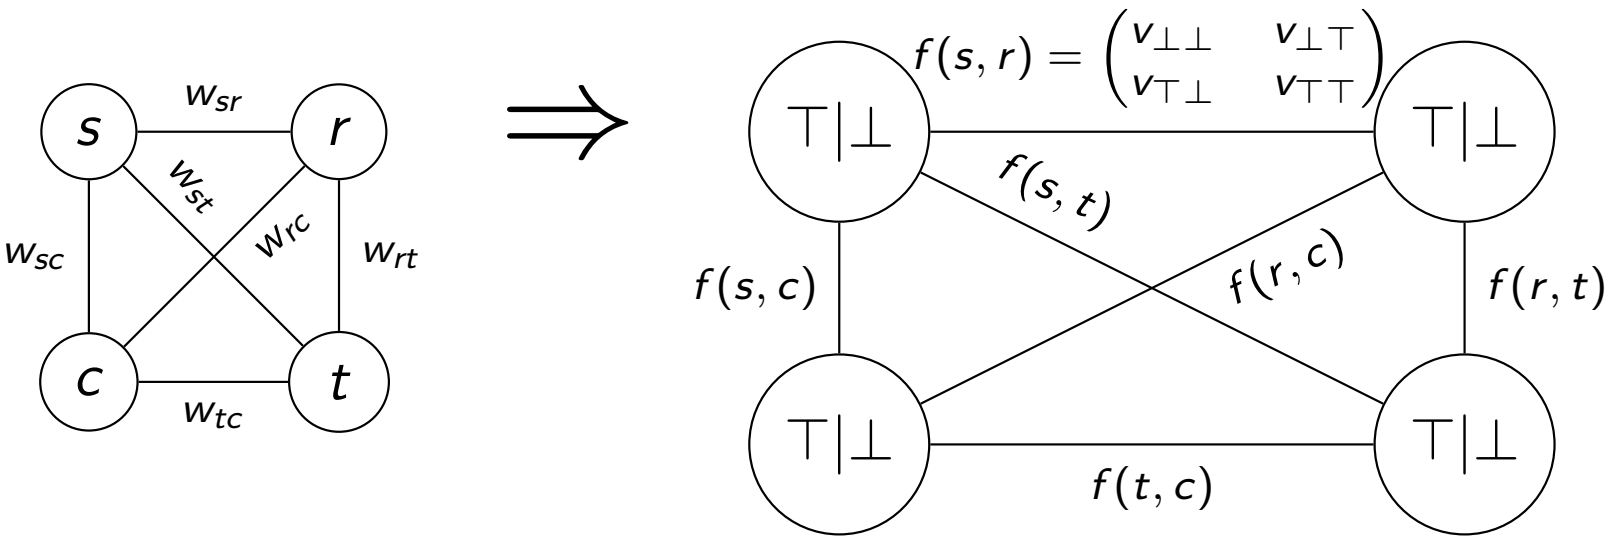

Computation:

- $Network = (N \text{ nodes}, M \text{ functions})$
- $P(n_1 = \perp, ..., n_N = \top) = \frac{\prod_{i=1}^{M} f_i(n_\alpha, n_\beta)}{Z}$ (Z = normalization factor)
- $P(n_1 = \top) = \sum_{\sigma_2,...,\sigma_N \in \{\perp,\top\}} P(n_1 = \top, n_2 = \sigma_2, ...)$
- $P(n_1 = \top | \{n_j = \sigma_j\}) = P(n_1 = \top)$ with Z reduced to sequences where $\{n_j = \sigma_j\}$

Implementation

Coded in Java, uses GATE (NL) and libDai (MN).

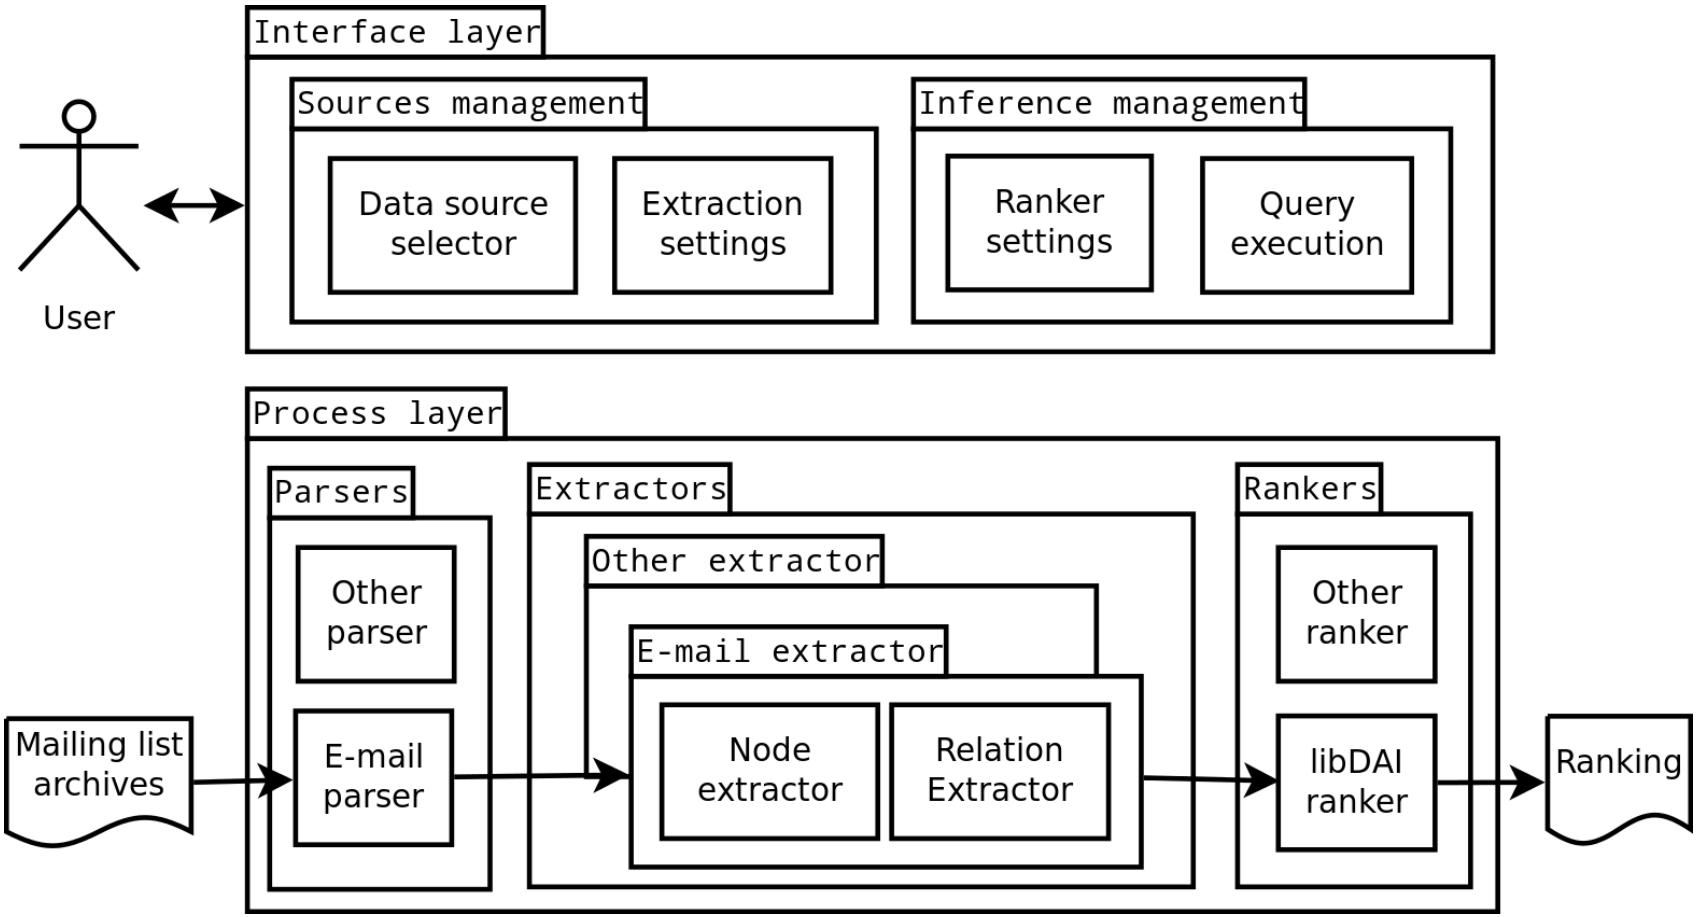

Node extractor for e-mails

Require: *mail*: Natural language e-mail
Ensure: S, R, T, C: Extracted stakeholders, roles, topics and terms

1: $S \leftarrow \{stakeholder(authorOf(mail))\}$
2: $R \leftarrow \emptyset$
3: $T \leftarrow \{topic(x) | x \in nounsOf(subjectOf(mail))\}$
4: $C \leftarrow \{term(x) | x \in nounsOf(bodyOf(mail))\}$

Relation extractor for e-mails

Require: $mail, S, R, T, C$: E-mail, stakeholders, roles, topics, terms
Ensure: L: Weighted relations

$L \leftarrow \emptyset$
$a \leftarrow author(mail)$
if $stakeholder(a) \in S$ **then**
  for all $t \in termsOf(bodyOf(mail))$ **do**
    if $term(t) \in C$ **then**
      $L \leftarrow merge(L, \{\langle stakeholder(a), term(t), 1\rangle\})$
    end if
  end for
end if
for all $topic \in T$ **do**
  if $nounOf(topic) \in nounsOf(subjectOf(mail))$ **then**
    $L \leftarrow merge(L, \{\langle stakeholder(a), topic, 1\rangle\})$
    for all $t \in nounsOf(bodyOf(mail))$ **do**
      if $term(t) \in C$ **then**
        $L \leftarrow merge(L, \{\langle topic, term(t), 1\rangle\})$
      end if
    end for
  end if
end for

C. Castro-Herrera and J. Cleland-Huang. A machine learning approach for identifying expert stakeholders. In 2009 Second International Workshop on Managing Requirements Knowledge (MARK), pages 45 –49, Sept. 2009. doi: 10.1109/MARK.2009.1.

C. Castro-Herrera and J. Cleland-Huang. Utilizing recommender systems to support software requirements elicitation. In Proc. of the 2nd International Workshop on RSSE, pages 6–10, New York, USA, 2010. ACM. ISBN 978-1-60558-974-9. doi: 10.1145/1808920.1808922. URL http://doi.acm.org/10.1145/1808920.1808922.

M. Karimzadehgan, R. W. White, and M. Richardson. Enhancing expert finding using organizational hierarchies. In Advances in Information Retrieval, number 5478 in LNCS, pages 177–188. Springer Berlin Heidelberg, Jan. 2009. ISBN 978-3-642-00957-0, 978-3-642-00958-7. URL http://link.springer.com/

chapter/10.1007/978-3-642-00958-7_18.

S. L. Lim, D. Quercia, and A. Finkelstein. StakeNet: using social networks to analyse the stakeholders of large-scale software projects. In Proc. of the 32nd ACM/IEEE ICSE, volume 1, pages 295–304, New York, USA, 2010. ACM. ISBN 978-1-60558-719-6. doi: http://doi.acm.org/10.1145/1806799.1806844. URL http://doi.acm.org/10.1145/1806799.1806844.

A. Mockus and J. D. Herbsleb. Expertise browser: a quantitative approach to identifying expertise. In Proc. of the 24th ICSE, page 503–512, New York, USA, 2002. ACM. ISBN 1-58113-472-X. doi: 10.1145/581339.581401. URL http://doi.acm.org/10.1145/581339.581401.

J. Mohebzada, G. Ruhe, and A. Eberlein. Systematic mapping of recommendation systems for requirements engineering. In

ICSSP, pages 200 –209, June 2012. doi: 10.1109/ICSSP.2012.6225965.

P. Serdyukov and D. Hiemstra. Modeling documents as mixtures of persons for expert finding. In Proc. of the IR Research, 30th ECIR, page 309–320, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 3-540-78645-7, 978-3-540-78645-0. URL http://dl.acm.org/citation.cfm?id=1793274.1793313.

J. Zhang, M. S. Ackerman, and L. Adamic. Expertise networks in online communities: structure and algorithms. In Proc. of the 16th international conference on WWW, page 221–230, New York, USA, 2007. ACM. ISBN 978-1-59593-654-7. doi: 10.1145/1242572.1242603. URL http://doi.acm.org/10.1145/1242572.1242603.