



Reconstructing the decision-making processes of a city council based on references between documents

Thomas Schoegje
t.m.schoegje@uu.nl
Utrecht University
Utrecht, Netherlands

Arjen de Vries
arjen@acm.org
Radboud University
Nijmegen, Netherlands

Lynda Hardman
lynda.hardman@cwi.nl
Centrum Wiskunde & Informatica (CWI)
Amsterdam, Netherlands

Toine Pieters
t.pieters@uu.nl
Utrecht University
Utrecht, Netherlands

ABSTRACT

Council members and policy workers need to understand (long-term) processes that lead to decisions. Gaining such an overview of a topic through a search engine can be challenging however, as searching a complex topic can result in an overwhelming number of documents and does not show how these documents are interrelated.

This study investigates how to create an overview of a decision-making process, which may be integrated into a search engine. Interviews show that policy workers consider documents relevant to the overview when the document and proposal were both created in response to the same council decision document. We identify such provenance based on the co-citation of documents and textual references between documents. In an exploratory user study, policy workers are tasked to understand the development of policy proposals based on provided timelines. Their relevance assessments show that our approach nearly exclusively finds relevant documents (a precision of 0.97).

Whereas the proposed approach identifies 91% of references made in documents, it only finds an exact target document in 39% of the total references. A further 52% of references finds a subset of documents including the target. A human in the loop can aid in finding the exact documents, and potentially add documents based on their domain expertise. The proposed approach creates an overview of a city council's decision-making process on a given topic with high precision, and might apply to other domains oriented around a decision-making process.

CCS CONCEPTS

- **Applied computing** → Document management and text processing; • **Information systems** → Digital libraries and archives;
- **Human-centered computing** → User studies.



This work is licensed under a Creative Commons Attribution International 4.0 License.

DGO 2024, June 11–14, 2024, Taipei, Taiwan
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0988-3/24/06
<https://doi.org/10.1145/3657054.3657116>

KEYWORDS

decision-making process, timeline, provenance, temporal information retrieval, decision history, policy-making, e-government

ACM Reference Format:

Thomas Schoegje, Lynda Hardman, Arjen de Vries, and Toine Pieters. 2024. Reconstructing the decision-making processes of a city council based on references between documents. In *25th Annual International Conference on Digital Government Research (DGO 2024)*, June 11–14, 2024, Taipei, Taiwan. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3657054.3657116>

1 INTRODUCTION

City council members create policies over multiple council meetings, building upon the existing policies. Previous policies may have been created years ago, by different council members. For instance, after deciding to construct a concert hall it takes years to construct and find out that the sound leaks between its music rooms. In this case, new council members first need to understand the original construction plan. They could use search engines to find relevant fragments of information (documents and meetings), but council members also need to understand how these fragments fit together and see the bigger picture [16]. Although doing this through searching individual documents would eventually lead to a complete picture, council members need to make decisions under external constraints such as time pressure [16]. Providing an overview of council information is invaluable, as such constraints limit how much information will people gather before making a decision [1] and could therefore lead to sub-optimal decisions.

Civil servants recognise the need for an overview of information, and therefore manually create timelines of complex policy proposals. The authors of timelines have no guidelines on how to create such histories. In section 4 we show how these represent the best effort of an individual, but are typically created with subjective inclusion criteria and are incomplete. Modern technologies enable a digital transformation of how we plan, record and archive the decision-making process of city council. This results in more transparent decisions and clearer accountability. In this paper we investigate how to generate timelines that complement, and eventually may replace, the manual timelines. Additionally, we investigate how to design an overview perspective (interface) based on these timelines. Our research questions are:

MRQ How can we create timelines of policy information?

- RQ1** Micro-level: What items should be included in the timeline?
- RQ2** Macro-level: How should the timeline be structured?
- RQ3** How can we algorithmically identify documents that should be included?
- RQ4** How should generated timelines replace manual timelines?

Drawing from informal interviews with the authors of manual timelines we develop an approach to generate timelines. We use two strategies based on extracting two types of links between documents. The first method identifies during which meetings the same (near) duplicate documents are discussed (see Figure 1). We interpret this as evidence that both documents are relevant in the same context, as determined by the staff that prepares the council meetings. The second method finds textual references in documents to other documents and meetings (see Figure 2). We interpret this as evidence that the referenced item is relevant for the current document, as determined by its author. In a user study we evaluate whether these two methods yield timelines with relevant information, and explore what qualities make for a good overview of policy information.

The main contribution of the paper is in characterising the need to generate timelines for council members, and in proposing a solution. This starts in section 2 by introducing related literature and introduces the council information dataset in section 3. Based on informal interviews, section 4 describes how experts manually construct manual timelines. Then two approaches to generate timelines are introduced (and combined) in section 5. In section 6 we present a user study where users are tasked to understand the development of a policy proposal based on a provided timeline. This study serves a dual purpose: to evaluate the relevance of documents selected by our method for a policy proposal, and to establish guidelines for timeline creation. We found that documents should be organised around their provenance (i.e. history), and that they are relevant to a proposal when both have origins in the same council decision. The generated timelines achieved a high precision (0.97) because the references between documents reflect their provenance. This precision score reflects that we do not need to account for weak links between documents, as references between documents are only created when they are directly related, according to staff with expertise.

The main limitation of our approach is that for ambiguous references we find multiple candidate documents, rather than just the intended target. This could be resolved in future work, or with a human in the loop. We conclude in section 7 that other decision-making processes that use shared meeting planners may also benefit from provenance-based timelines. An overview of a decision-making process makes the decisions more transparent, explainable and useful.

2 RELATED WORK

Transparent decision-processes tend lead to better results [8] and clarify who should be accountable for decisions. Making contextualised decision information easily accessible also lowers barriers to involving citizens [5, 12]. These factors enable an open government

[2] and the quality of information available, which fosters trust [9, 11].

2.0.1 Timelines. To generate timelines we can adapt techniques from the recent survey by Norambuena et al. [14]. They show works at three different levels of resolution: sentence level, document level and cluster level. Techniques that model events at document level are of interest to our setting, as council decisions are recorded in official documents. An influential approach to map out the narrative threads in a corpus is the metro maps approach [17, 18], which constructs maps of interconnected narrative threads. Links between documents are based on the similarity of documents, determined by identifying important words in the corpus. Other approaches are based on extracting entities and/or events and then temporally ordering them (e.g. [10, 13]). The structure of council information enables two further approaches: following co-references and textual references to older documents.

2.0.2 Co-reference approaches (approach 1). When multiple council meetings discuss the same (near) duplicate document, we can view this as one meeting referencing/building on an older meeting. Such citation patterns have been analysed in the scientific literature (see e.g. [6]). Timelines of policy proposals reason back in time (i.e. how did this political decision follow from the previous ones?), similar to the bibliographic coupling approach for citation analysis.

A more comprehensive way to record references is to track the provenance of documents. A provenance model was proposed that records the process of how information artefacts are created, and from which information it was derived [7]. Figure 3 gives an overview of the important concepts, with two examples: the city council acts as an agent that performs the activity of making decisions, and generates their decisions which are recorded in entities such as motions and council proposals. These decisions lead to new activities performed by civil servants (agents), and outcomes are reported in documents (entities) sent to the council.

2.0.3 Identifying in-text references (approach 2). Several approaches have used textual document references to identify how a narrative developed [3, 21, 22], often relying on URLs. Textual references can be less specific references, where multiple documents qualify (e.g. there are multiple documents on that date). This can be resolved by finding the candidate articles, and selecting the intended target based on further context (e.g. topic similarity) [3]. When establishing links between documents using both co-citation and in-text references, a graph is formed where documents may be linked despite being multiple steps removed. We can model the strength of a link between documents by assigning a weight to the edges between documents is, which allows us to account for how strongly documents are linked through network analysis.

3 DATASET

The dataset is the public council information of the city of Utrecht in the Netherlands. This dataset reflects the two active responsibilities of council members: to shape new policies, and to oversee whether the municipality properly executes those policies. These responsibilities are recorded and carried out during weekly council meetings. The structure of council information is shown in Figure




Figure 1: Plotting how document re-use can be used to generate timelines of council proposals. This example shows how the council might vote on building a concert hall, and how the light grey and dark grey documents can be used to link these meetings.

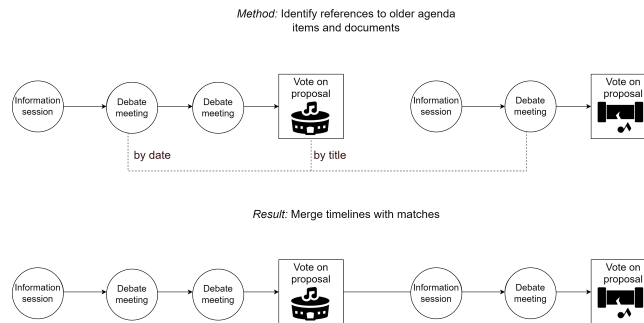


Figure 2: Plotting how temporal expressions can be used to find out which timelines should be merged. This example shows how the council might first vote on the construction of a concert hall, and later on a sound leak between the studios.

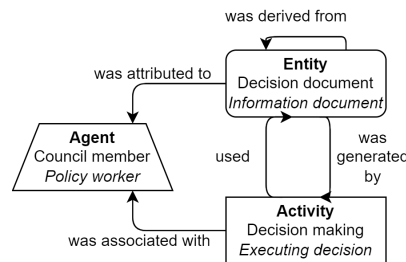


Figure 3: An overview of the key concepts in provenance, adapted from Groth et al [7]. In normal text, a council member example. In italics, a policy worker example.

4. During each meeting multiple items are discussed, typically concerning a variety of topics. Meeting items have documents attached which provide pertinent background information, as collected by council clerks.

Council documents are pdf files that come in various genres. The main ones being council letters, motions, memos, decision histories and council proposals. Council letters generally inform the council on upcoming matters. Motions are discussion points prepared by council members. Memos update the council on small matters. The manual decision histories were constructed by clerks for complex political topics, when it was necessary to give council members a better overview of the temporal context. The outcome of the policy making process is a policy proposal. This type of document contains the policy that was decided, and is voted on by the council.

The dataset includes all data from 2013 (the origin of this system) until 2022. There are 1648 meetings held between 2013 and 2023, containing 15,314 agenda items and discussing 29,229 documents.

In addition, the municipality can send the council documents that are not tied to an agenda item, but instead uploaded as ‘entries’. Each entry is a council document that may have attachments. Whereas the council has staff that select the documents attached to meetings, the entries originate from civil servants writing to the council. The dataset contains 19,156 entries that contain 22,908 documents. Entries and meeting documents can overlap, when entries are re-uploaded to discuss during a meeting item.




Figure 4: Documents are organised around meetings and around entries (documents from civil servants). These sets overlap when entries are discussed during a meeting.

4 MANUAL TIMELINES

The authors of manual timelines (and the authors of this paper) are not aware of existing guidelines on constructing these timelines, and therefore create them at their own discretion. Consequently there are differences in, for example, which document genres the author includes (only council proposals, or perhaps also council letters or motions). We conducted informal interviews with six civil (three female, three male) servants who submitted a manual timeline to the council to see how and why they constructed their manual timeline. Two participants were project managers who had delegated the creation of the timeline to a colleague (30 min interviews). The remaining four were authors of the timelines (50 min interviews) themselves. Authors worked in administrative roles, supporting policy creation. All interviewees were mediators or seniors.

While every council proposal includes a brief overview of key decisions, for complex topics the authors of manual timelines can decide to provide a more comprehensive overview. This begins with the author’s domain knowledge of key items, and is extended through search and asking colleagues for help. Although the authors have expertise in supporting policy makers, they are not necessarily domain experts. The reported inclusion criteria for items were subjective, as authors 1) determine what is important enough to include, 2) possess limited knowledge on what information exists and 3) may have ulterior goals when creating a timeline. For instance, one interviewee created an overly comprehensive document to emphasise how long the council had been undecided about a topic.

A recurring theme was that authors reference important documents when writing to the council. There are two types of references between documents: 1) co-citation of documents within the meeting planner, created by clerks attaching important documents to meetings, and 2) textual references in documents, used when authors describe why they are writing the council.

In summary, manual timelines are created for diverse purposes, in multiple formats and based on subjective inclusion criteria. Consequently, these do not reflect an objective or complete perspective. An algorithmic approach presents a more scalable alternative, capable of mitigating subjective factors.

5 IMPLEMENTATION

Our approach to generate timelines is based on the two types of links between documents. Both types are combined for a more comprehensive timeline.

The overall approach is to interpret each agenda-item and entry as a separate timeline, and progressively combine timelines. First we merge timelines where the same document is cited, by grouping timelines with (near) duplicate documents (approach 1). Then we identify the textual references between documents, and merge timelines that refer to each other (approach 2). Finally, we sort the timeline based on the meeting date (for agenda item) or upload date (for entries).

The source code we created for this project is available at github.com/UtrechtUniversity/expertsearch, and the dataset can be accessed through zoek.openraadsinformatie.nl (last accessed 11-9-2021).

5.1 Re-use of (near) duplicates

Co-citation of documents is found by identifying agenda items and entries containing (near) duplicate documents. Duplicate detection is performed by finding documents with the same filename and/or the same displayname in the metadata fields. For this set of documents it was ensured that all documents have a similar filesize, defined as not deviating by more than 5% from the average filesize in the set. This resulted in 43507 unique documents. Before grouping, filenames and displaynames were normalised by removing file extensions and any trailing white spaces. The requirement for similar file sizes prevented cases where a generic filename (e.g. ‘Proposal.pdf’) refers to completely different documents.

The algorithm models each agenda item as a small timeline, and for each (near) duplicate document a new timeline is created that merges all timelines that co-cite it. If the original timeline contained multiple duplicates, then the same timeline will be merged into multiple new timelines. Therefore an iterative process is started to identify and merge timelines that contain new duplicates. This yielded 3,006 timelines that are especially strong at showing the history of individual policy proposals, including the weekly meetings and documents.

5.2 Textual references

Figure 6 shows the processing pipeline for textual references. There are three types of textual references, which we define and extract as follows:

References by ID are very specific, and refer to a unique document. ID extraction is done by detecting hyperlinks in the text using `pdfminer` in python [19]. The contents of the URL indicate what type of meeting or document is referenced, as well as its unique ID.

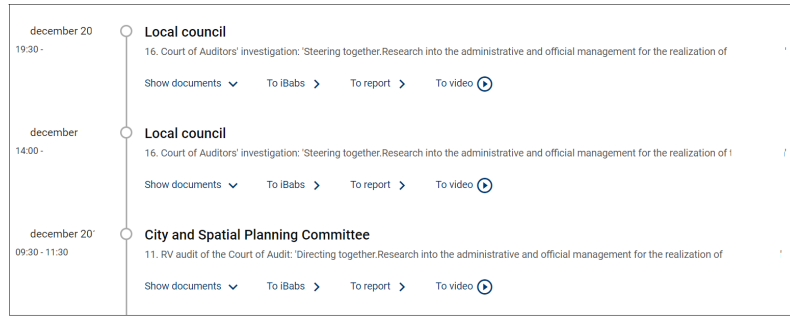


Figure 5: The resulting timeline as shown in the 'U-reka' search engine for council members (ureka.utrecht.nl/app/). All Dutch text was translated through Google translate.

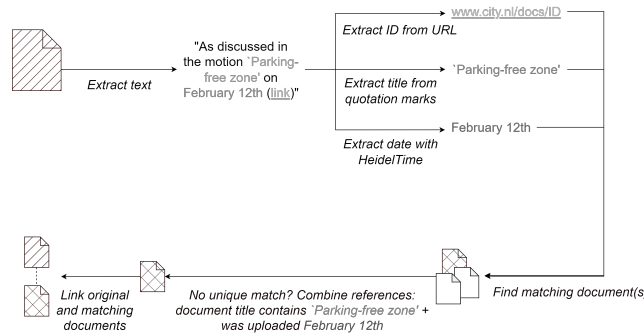


Figure 6: The processing pipeline for textual references to documents and meetings.

References by title are less specific, as documents may share titles or have non-descriptive and generic titles. Title extraction is performed using regular expressions to identify text strings enclosed with quotation marks (e.g. ' ' '). As document titles were usually shorted, we tested if each string is a substring of an existing document title.

References by date are the least specific, as sets of documents are typically submitted as a batch on the same date. Additionally, it is ambiguous whether a date refers to a document or a meeting. Dates were extracted using HeidelTime [20]. Relative expressions such as 'yesterday' were normalized to the document's upload date. Because documents are uploaded in batches, the exact target is typically ambiguous. Additionally, it can be unclear whether a date refers to a document or a meeting.

Ambiguous references (e.g. multiple documents are uploaded on the same date, or share the same title) are disambiguated to the target document using other references within the same sentence. We ignore references that we cannot disambiguate to a single target to maintain a high precision (rate of true positives) in our timelines. This increases user trust, and it prevents the co-citation approach from including irrelevant documents.

The resulting timelines typically show the progression between council decisions, on a time-scale of months to years. The textual references allowed us to merge the 3,006 timelines found through co-citation into 2,751 timelines. There is an overlap in the links found by both types of references, but the co-citation approach

is particularly adept at finding short-term connections whereas the textual references are better at finding long-term connections. On average the timelines consist of 4.55 agenda items and entries, spanning a period of 12.5 months. Two outliers spanned 64 months and 32 months. Two others spanned less than 1 month.

5.3 Technical evaluation: identifying references

We evaluate what proportion of in-text references we successfully extracted from documents, and how many of these we disambiguated successfully. We annotated all textual references in ten randomly chosen manually constructed manual timelines (as other document genres contained fewer references). These annotations were compared to the extracted references.

Table 1 shows that we successfully extracted 91% of references from the documents, but that we were only able to disambiguate 39% of references to an exact document. References by URL were identified and trivially disambiguated, but a part of these include dead links (to pages no longer in the dataset) or link to pages outside of the dataset. References by date and title often refer to a subset of documents, including the target document. In the following user study we investigate whether the references that we found resulted in relevant documents.

Table 1: How many unique references were annotated, how many we were able to extract, and how many we were able to disambiguate to a single target document.

	By ID	By date	By title
Total	100% (45)	100% (46)	100% (46)
Extracted	100% (45)	93% (43)	80% (37)
Disambiguated	82% (37)	22% (10)	9% (4)

6 QUALITATIVE EVALUATION AND EXPLORATION

We performed an exploratory study with a dual purpose: to investigate the qualities of a useful overview of policy information, and to assess whether the documents found by the proposed approach are relevant.

During interactive sessions users were provided with a timeline and tasked to understand what led to a policy proposal. They also assessed the relevance of each document in the timeline. Each timeline consisted of a chronologically ordered list of entries and agenda items.

Participants were presented different types of timelines to investigate the qualities of both manual and generated timelines. For four proposals, the timelines were created by combining the manual and generated timeline. One further proposal only consisted of generated items and the last one of only manually selected items. To investigate inter-rater agreement, four participants were assigned identical tasks to previous participants. These timelines were between 6-12 items long (8.5 on average), with one outlier containing only 4 items.

6.0.1 Participants and tasks. As council members were unavailable, we invited policy workers. These are domain experts who work with council information (for more detail, see Schoegje et al. [15]). Ten policy workers (six female, four male) were invited to participate. Four participants had over five years of experience, three had between five and one years, and three had less than a year. Nine of these sessions were conducted in person, and one over Microsoft Teams.

Participants were invited for an interactive session (30-45 minutes) where they were tasked to use a timeline to understand what led to a policy proposal. Six policy proposals were selected, each accompanied with a timeline. The proposals were randomly selected, although one proposal was replaced with another, because its timeline was too large to discuss during a single session.

6.0.2 Procedure. After introductions and securing informed consent, participants were directed to read the policy proposal prepared on the screen. Participants first explained to the interviewer what the proposal was about, and then chronologically started reading the documents in the timeline. Per document, the participant was first asked to clarify its contents, and what happened in between this document and its predecessor. They were asked whether the document was relevant, and whether it was useful for understanding the policy proposal (both on a three-point scale). They were asked to give reasons for these assessments.

Within this structure, the interviewer would allow room to discuss themes and questions that arose. In the early interviews, these

themes were primarily about individual documents and what made them relevant or useful. In later interviews, discussions shifted about the timeline and policy-making process as a whole, including themes about completeness and conciseness.

6.0.3 Analysis. All interview responses were analyzed with a thematic analysis, by grouping responses based on recurring themes. In multiple iterations the themes were refined to be more descriptive and better reflect the participant responses.

6.1 Results

Three main themes emerged: inclusion criteria; completeness and conciseness; and overview perspective. These themes and the main points are briefly summarized in Table 2.

Inclusion criteria [RQ1]: In this setting it was not important to find whether a document was relevant to the topic or useful to the task. Instead, a document should be included in the timeline if both the document and the associated policy proposal can trace their provenance (as introduced in section 2.2) from the same council decision. This provenance is not explicitly tracked, but it is of such importance that authors of council documents are consistent (and trained) in referencing to previous decisions and documents. Documents typically state why the council has to read them. The importance of council decisions in particular were reflected in every interview, as shown e.g. when P3 stated “[this document] is important because the council apparently has thoughts about this”. P5 termed the document types which contain council decisions as “milestone documents”, and references to these milestones when persuading other civil servants to take their requests seriously (e.g. in emails).

Participant P4 illustrated the importance of provenance over relevance, stating that “Sometimes the council asks questions about a different topic”. In these cases “There can be a whole [internal] discussion on [which civil servant] should address that question, but we don’t want to record that”. Such semi-related documents need to be included “in the same package”, to keep track of “who is responsible for it”. P9 was reluctant to use relevance as an organising principle as “[something] can be relevant, but it’s a side issue”, and “there comes a point where everything is relevant to everything”.

Participants rated nearly all documents found by our approach as relevant, and no difference between any document’s relevance or usefulness. The only documents that were not relevant were 1) a single entry that was incorrectly included because the duplicate detection malfunctioned (and the co-citation approach included a non-duplicate document), and 2) a few agenda items during administrative meetings. These administrative meetings should not be included in the timeline, because these same items were discussed during a separate meeting later that week. After excluding this type

Table 2: Summary of the main themes and findings in the qualitative study

Theme	Main findings
Inclusion criteria	Structure timelines around tracing document provenance Correct provenance leads to relevant and useful documents The proposed method finds documents from the correct provenance
Completeness and conciseness	Include a view on only decisions Include a view that includes the steps in between decisions
Overview perspective	We can present a layered view, that expands from decisions to all provenance Show how timelines intersect and interrelate Overview perspective aids in understanding big picture Overview perspective aids in finding holes/curiosities in big picture Linked data in overview avoids challenges in selecting keywords for search
Comparing types of timelines	No qualitative differences found

of meeting from the generated timelines, 97% of timeline items were found to be relevant. This precision score reflects that we do not need to account for weak links between documents, as references between documents are only created when they are directly related, according to staff with expertise. Participants always agreed in their assessments of items, with the exception of P5. P5 only considered council decisions as relevant, and other documents (those of an informative nature) as semi-relevant. All participants gave input around this theme.

Completeness and conciseness [RQ2]: Participants noted both the importance of conciseness (P6: *“Less is better”*) and completeness (P4: *“Maybe I’m too careful, but I want the complete picture”*). P5 described how these relate: *“There is information at two ranks. Decisions are at the first rank ... informing [documents are] at the second rank”*. They suggested a layered overview, which initially shows a timeline of the decision layer, which can be expanded to also show a provenance layer by toggling a button. The decision layer displays the main items, and expanding the provenance layer displays the sub-steps. This is illustrated in Figure 7 The decision layer encompasses the decisions made by the council, as recorded in motions, policy proposals, formal questions to alderpeople and formal promises by alderpeople. These result in duties and activities from municipal staff, which in turn lead to documents whose provenance is displayed in the provenance layer.

P5 suggested that the decision layer should support people who *“primarily need to know what the council decided, and the current state of affairs”*. As such, the decision layer should include which decisions have been resolved, and the latest information on those unresolved. P5 noted that such information also serves as a form of accountability, where the staff shows *“we haven’t been idle”*. Note that ‘latest information’ refers to the latest official document sent to the council, as more recent working documents *“are usually not ready yet to show [to the council]”* (P10). Conversely, the provenance layer is more comprehensive, providing background information and showing how individual decisions were made. P10 noted the history of individual decisions can be important *“when something seems awkward”* about them. All participants had responses about this theme.

Overview perspective: Although some decisions follow a linear process, P8 highlighted that *“sometimes multiple [time]lines*

converge, and I can’t see if [these documents] are in the same line”. Therefore, overviews should display the different lines, and how they interrelate (akin to the metro maps timelines [18]). P6 shared their experience tracing document references without an overview perspective, noting that *“[they] couldn’t see the forest for the trees”*. P10 cited an example where this approach took them 2-3 hours to understand what lead to a policy proposal. P7 appreciated that a timeline draws attention to time gaps between documents, which they considered important clues that *“something might be missing”* and could prompt further investigation. P4 commended on the robustness of the timeline against changes in the keywords in documents, highlighting an example when *“the building changed names”* which would be less obvious when searching by keywords. Seven participants had responses about this theme.

Comparing types of timelines: No obvious differences emerged between participants who used different types of timeline (manual, generated, combined).

6.2 Discussion

We reflect upon the practical implications, and then the theoretical implications of the study. We find two themes with regards to practical implications:

How to structure council decisions: The importance of provenance explains why authors of council documents are consistent in referencing past documents in the text of their documents. This practice ensures that we can generate timelines based on references to documents (RQ3). Although previous work typically generated timelines based on document similarity, our approach can leverage the provenance information that is provided explicitly in the council’s work processes. The importance of provenance in this setting suggests that similarity based approaches are less appropriate for decision-making processes. Provenance ties into the main responsibilities of council members: shaping policies that generate the municipality’s activities, and overseeing the execution of those activities. Hence a timeline should be presented as an overview with two layers: a decision layer that conveys these two things (decisions and current state of activities), and a provenance layer which includes more comprehensive information.

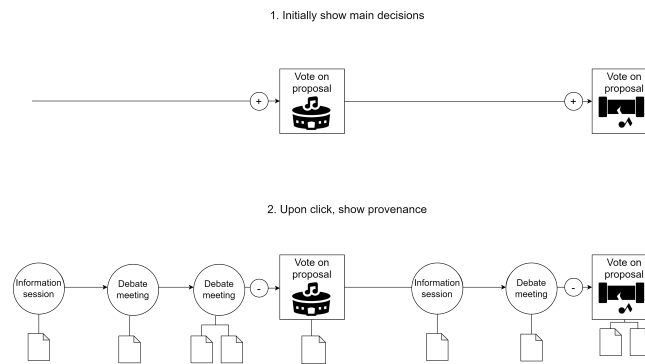


Figure 7: Example of how an interface can show information in two layers. The main decisions are initially shown, and the user can click the small circle with a '+' symbol to expand the view. The expanded view also shows the sub-steps within that decision-making process. This second layer of information shows the provenance of decisions and documents.

Investigating decision-making: To quantify the value of an overview perspective, future work could investigate how much information is available at the moment of making a decision (with or without the overview), and how the availability of information affected the decisions that were made. The establishment of reliable overviews of council decisions also facilitates further research on the nature of those decisions. For instance on identifying critical decision moments [4], and whether these moments are more likely to arise when the topic is discussed in information sources like the news.

With regards to the theoretical implications, we reflect:

Generalisation: The proposed method depends on the process of how decisions are made, as well as how they are recorded and archived. The decision-making process is similar for many governmental organisations at the municipal and (sub)national level in Country, although smaller organisations may record these in less detail as they have less administrative staff available. The approach likely generalises to many (country) governmental organisations, especially as most municipalities in the country use the same two meeting planner systems and hence already structure their data similarly.

The method itself is domain-agnostic, and could be adapted for similar decision-making processes at other organisations. Potential future work could investigate whether this approach could be adapted broadly, specifically for use in organisations that have integrated software for their email, calendar and content management (e.g. organisations using SharePoint). Although these processes will be less structured and recorded less accurately, a timeline might still be a useful way to organise and revisit information.

Limitations: Although our approach achieved a high precision in retrieving relevant documents (97%), it is only an explorative step towards generating overviews of policy information. The main limitation is that, although 91% of references are extracted from the text, only 39% of total references find exact matches. Future work should improve the detection of true positives, both in the reference disambiguation (e.g. using the domains given in document metadata), and in the duplicate detection. One approach to disambiguate

references and identify missing documents is to involve a domain expert in the loop (RQ4).

As the domain experts reference the most vital documents known to their expertise, we assume that our timelines cover the most important moments towards a decision. Future work could include documents that did not directly lead to a given council decision, but might still include contextual information. Improving the recall of these timelines could be done by first generating high-precision timelines, and using these documents as the basis for content-based recommendations. A domain expert could prepare a timeline for council members by first disambiguating references, and then expanding the timelines based on recommendations.

A limitation of the experimental setup is that each user only used one timeline. Although no obvious differences between the manual and generated timelines emerged, properly investigating these differences would require a comparative study.

Finally we wish to highlight the limited number of participants in the study ($n = 10$). Although this is a useful sample size for an iterative design process, later stages in this research direction will require studies with a larger participant pool to quantify how such an overview aids in task performance compared to searching individual documents.

7 CONCLUSION

Understanding the decision-making process of a city council necessitates an understanding how council documents are interrelated. In this paper we considered a digital transformation of how we plan, record and archive the decision-making process of city council can result in more transparent decisions and clearer accountability. Specifically, we (re)constructed the timelines of the policy-making process from the existing council information. As an informal user study indicated that authors of policy documents consistently reference important documents, we proposed an approach to generate timelines based on two types of references between documents: document co-citation during meetings and textual references within council documents. We generated timelines of individual policy proposals by identifying meetings that discussed the same documents. Timelines of how policy proposals extend one another were identified by examining textual references between documents.

A user study with policy workers investigated both 1) guidelines for designing an overview interface and 2) whether the generated timelines included relevant documents. Experts considered documents relevant if the document and council decision both result from the same council decision (RQ1). Such provenance is recorded through references between documents. Policy workers need to see timelines from an overview perspective that balances conciseness with completeness by providing a decision layer and a comprehensive provenance layer (RQ2).

Creating an overview based on references between documents nearly exclusively yields relevant documents, with a precision of 97% (RQ3). The main limitation is that our approach identified the exact target document for only 39% of textual references. A further 52% of the references is ambiguous, finding a subset of documents. Although future work can enhance this aspect by extracting more context around references, we recommend involving a domain expert in the loop to select the exact matches and identify if further missing documents (RQ4).

A practical application of this work is to include overviews of council information in the municipality's search engine, such that users can click individual search results to view that document in the context of a larger decision history. This presents a step towards better supporting council decision making, by making pertinent information more accessible. The method we proposed generates timelines of a local government's decision-making process with high precision. As other municipalities in the Netherlands use a similar decision-making process and similar software to plan and archive their meetings, we expect that our functionality can be adapted to those organisations with minimal adjustments.

ACKNOWLEDGMENTS

We are grateful to the contributions of Jos Nicolai and Anne de Wildt to early discussions of this work, and the code written by Jos Nicolai to parse council information.

REFERENCES

- [1] Jennifer M. Berryman. 2006. What defines 'enough' information? How policy workers make judgements and decisions during information seeking: preliminary results from an exploratory study. *Inf. Res.* 11, 4 (2006), 14 pages. <http://www.informationr.net/ir/11-4/paper266.html>
- [2] Fabian Birghan, Robert Hettenhausen, Christine Meschede, and Tobias Siebenlist. 2019. Informing Citizens via Council Information Systems. In *20th Annual International Conference on Digital Government Research, DGO 2019, June 18-20, 2019*, Yu-Che Chen, Fadi Salem, and Anneke Zuiderwijk (Eds.). ACM, Dubai, United Arab Emirates, 280–286. <https://doi.org/10.1145/3325112.3325220>
- [3] Thomas Bögel and Michael Gertz. 2015. Time will Tell: Temporal Linking of News Stories. In *Proceedings of the 15th ACM/IEEE-CE Joint Conference on Digital Libraries, June 21-25, 2015*, Paul Logasa Bogen II, Suzie Allard, Holly Mercer, Micah Beck, Sally Jo Cunningham, Dion Hoe-Lian Goh, and Geneva Henry (Eds.). ACM, Knoxville, TN, USA, 195–204. <https://doi.org/10.1145/2756406.2756919>
- [4] Michael D Cohen, James G March, and Johan P Olsen. 1972. A garbage can model of organizational choice. *Administrative science quarterly* 17 (1972), 1–25.
- [5] Noella Edelmann and Valerie Albrecht. 2023. Designing public participation in the digital age: Lessons learned from using the policy cycle in an Austrian case study. In *Proceedings of the 24th Annual International Conference on Digital Government Research, DGO 2023, July 11-14, 2023*, David Duenas-Cid, Nadzeya Sabatini, Loni Hagen, and Hsin-chung Liao (Eds.). ACM, Gdańsk, Poland, 300–308. <https://doi.org/10.1145/3598469.3598502>
- [6] Stefano Ferilli, Domenico Redavid, and Davide Di Pierro. 2023. Holistic graph-based document representation and management for open science. *Int. J. Digit. Libr.* 24, 4 (2023), 205–227. <https://doi.org/10.1007/S00799-022-00328-Z>
- [7] Paul Groth and Luc Moreau. 2013. PROV-overview. An overview of the PROV family of documents. Retrieved February 8, 2022 from <https://www.w3.org/TR/prov-overview/>
- [8] Osama Ibrahim and Aron Larsson. 2019. Intelligibility and Transparency in Model-based Collaborative Governance. In *20th Annual International Conference on Digital Government Research, DGO 2019, June 18-20, 2019*, Yu-Che Chen, Fadi Salem, and Anneke Zuiderwijk (Eds.). ACM, Dubai, United Arab Emirates, 214–226. <https://doi.org/10.1145/3325112.3325247>
- [9] Anas Kanaan, Ahmad AL-Hawamleh, Anas Abulfaraj, H Al-Kaseasbeh, and Almuhammad Alorfi. 2023. The effect of quality, security and privacy factors on trust and intention to use e-government services. *International Journal of Data and Network Science* 7, 1 (2023), 185–198.
- [10] Egoitz Laparra, Itziar Aldabe, and German Rigau. 2015. From TimeLines to StoryLines: A preliminary proposal for evaluating narratives. In *Proceedings of the First Workshop on Computing News Storylines*. Association for Computational Linguistics, Beijing, China, 50–55.
- [11] Yan Li and Huping Shang. 2023. How does e-government use affect citizens' trust in government? Empirical evidence from China. *Information & Management* 60, 7 (2023), 103844.
- [12] Amal Marzouki, Sehl Mellouli, and Sylvie Daniel. 2018. Spatial, temporal and semantic contextualization of citizen participation. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age, DGO 2018, May 30 - June 01, 2018*, Marijn Janssen, Soon Ae Chun, and Vishanth Weerakkody (Eds.). ACM, Delft, The Netherlands, 63:1–63:8. <https://doi.org/10.1145/3209281.3209385>
- [13] Anne-Lyse Minard, Manuela Speranza, Eneko Agirre, Itziar Aldabe, Marieke van Erp, Bernardo Magnini, German Rigau, and Ruben Urizar. 2015. SemEval-2015 Task 4: TimeLine: Cross-Document Event Ordering. In *Proceedings of the 9th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2015, June 4-5, 2015*, Daniel M. Cer, David Jurgens, Preslav Nakov, and Torsten Zesch (Eds.). The Association for Computer Linguistics, Denver, Colorado, USA, 778–786. <https://doi.org/10.18653/v1/s15-2132>
- [14] Brian Keith Norambuena, Tanushree Mitra, and Chris North. 2023. A Survey on Event-based News Narrative Extraction. *CoRR* abs/2302.08351 (2023), 39 pages. <https://doi.org/10.48550/arXiv.2302.08351>
- [15] Thomas Schoegje, Arjen P. de Vries, Lynda Hardman, and Toine Pieters. 2023. Improving the Effectiveness and Efficiency of Web-Based Search Tasks for Policy Workers. *Inf.* 14, 7 (2023), 371. <https://doi.org/10.3390/info14070371>
- [16] Thomas Schoegje, Arjen P. de Vries, and Toine Pieters. 2022. Adapting a Faceted Search Task Model for the Development of a Domain-Specific Council Information Search Engine. In *Electronic Government - 21st IFIP WG 8.5 International Conference, EGOV 2022, September 6-8, 2022, Proceedings (Lecture Notes in Computer Science, Vol. 13391)*, Marijn Janssen, Csaba Csáki, Ida Lindgren, Euripides N. Loukis, Ulf Melin, Gabriela Viale Pereira, Manuel Pedro Rodríguez Bolívar, and Eftihimos Tambouris (Eds.). Springer, Linköping, Sweden, 402–418. https://doi.org/10.1007/978-3-031-15086-9_26
- [17] Dafna Shahaf, Carlos Guestrin, and Eric Horvitz. 2012. Metro maps of science. In *The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, August 12-16, 2012*, Qiang Yang, Deepak Agarwal, and Jian Pei (Eds.). ACM, Beijing, China, 1122–1130. <https://doi.org/10.1145/2339530.2339706>
- [18] Dafna Shahaf, Carlos Guestrin, and Eric Horvitz. 2013. "Metro maps of information" by Dafna Shahaf, Carlos Guestrin and Eric Horvitz, with Chingman Au Yeung as coordinator. *SIGWEB Newsl.* 2013, Spring (2013), 4:1–4:9. <https://doi.org/10.1145/2451836.2451840>
- [19] Yusuke Shinyama. 2022. Elasticsearch 7.16.2 [computer software]. Retrieved March 10, 2022 from <https://pypi.org/project/pdfminer/>
- [20] Jannik Strötgen and Michael Gertz. 2013. Multilingual and Cross-domain Temporal Tagging. *Language Resources and Evaluation* 47, 2 (2013), 269–298. <https://doi.org/10.1007/s10579-012-9179-y>
- [21] Mikhail Tikhomirov and Boris V. Dobrov. 2017. News Timeline Generation: Accounting for Structural Aspects and Temporal Nature of News Stream. In *Data Analytics and Management in Data Intensive Domains - XIX International Conference, DAMDID/RCDL 2017, October 10-13, 2017, Revised Selected Papers (Communications in Computer and Information Science, Vol. 822)*, Leonid A. Kalinichenko, Yannis Manolopoulos, Oleg Malkov, Nikolay A. Skvortsov, Sergey A. Stupnikov, and Vladimir Sukhomlin (Eds.). Springer, Moscow, Russia, 267–280. https://doi.org/10.1007/978-3-319-96553-6_19
- [22] Giang Binh Tran, Tuan A Tran, Nam-Khanh Tran, Mohammad Alrifai, and Nattiya Kanhabua. 2013. Leveraging learning to rank in an optimization framework for timeline summarization. In *SIGIR 2013 Workshop on Time-aware Information Access (TAIA)*. ACM, Coast, Queensland, Australia, 4 pages.

Received 26 January 2024; revised 25 April 2024; accepted xxx