



DIPARTIMENTO DI INFORMATICA
E SISTEMISTICA ANTONIO RUBERTI



SAPIENZA
UNIVERSITÀ DI ROMA

Groupware Mail Messages Analysis for Mining Collaborative Processes

Claudio Di Ciccio
Massimo Mecella
Monica Scannapieco
Diego Zardetto
Tiziana Catarci

Technical Report n. 1, 2011

Groupware Mail Messages Analysis for Mining Collaborative Processes

Claudio Di Ciccio ^{#1}, Massimo Mecella ^{#2}, Monica Scannapieco ^{*3},
Diego Zardetto ^{*4}, Tiziana Catarci ^{#5}

*#SAPIENZA – Università di Roma,
Dipartimento di Informatica e Sistemistica ANTONIO RUBERTI
Via Ariosto 25, Roma, Italy*

^{1,2,5}{cdc,mecella,catarci}@dis.uniroma1.it

**Istituto Nazionale di Statistica
Via Balbo 16, Roma, Italy*

^{3,4}{scannapi,zardetto}@istat.it

Abstract

Artful business processes are informal processes typically carried out by those people whose work is mental rather than physical (managers, professors, researchers, etc.), the so called “knowledge workers”. In this paper we outline how to automatically build, on top of a collection of e-mail messages, a set of workflow models that represent the artful processes which lay behind the knowledge workers activities.

Index Terms

Object Matching, Process Mining, Text Mining, Information Retrieval, E-mail Analysis, Knowledge Workers, Artful Business Process

January 4, 2011

I. INTRODUCTION

Nowadays, the most of the research related to workflows has considered the management of formal business processes. There has been some discussion of informal processes, often under names such as “artful business processes” (see [1]). Informal processes are typically carried out by those people whose work is mental rather than physical (managers, professors, researchers, etc.), the so called “knowledge workers” ([2]). With their skills, experience and knowledge, they are used to perform difficult tasks, which require complex, rapid decisions among multiple possible strategies, in order to fulfill specific goals. In contrast to business processes that are formal and standardized, often informal processes are not even written down, let alone defined formally, and can vary from person to person even when those involved are pursuing the same objective. Knowledge workers create informal processes “on the fly” to cope with many of the situations that arise in their daily work. While informal processes are frequently repeated, since they are not written down, they are not exactly reproducible, even by their originators, nor can they be easily shared. Their outcome releases and their information exchanges are very often done by means of e-mail conversations, which are a fast, reliable, permanent way of keeping track of the activities that they fulfill.

The objective of the research outlined in this work is to automatically build, on top of a collection e-mails, a set of workflow models that represent the artful processes which lay behind the knowledge workers activities.

A company can take many advantages out of this work. First of all, the unspecified agile processes that are autonomously used become formalized: since such models are not defined *a priori* by experts but rather inferred from real-life scenarios that actually took place, they are guaranteed to respect the true executions and not reflect the expected ones (often Business Process Management tools are used to show the discrepancy between the theoretical workflows and the concrete ones). Secondly, such models can be shared and compared, so that the best practices might be put in evidence from the community of knowledge workers, to the whole business benefit. Moreover, without any further computational cost, an analysis over such processes can be done, so that bottlenecks and delays in actual executions can be found out. We remark here that all of these utilities come out with almost no effort for workers, due to the automated nature of the proposed system.

The approach we propose, in order to retrieve a collection of process models out of an initial set of e-mail messages, involves many research fields at a time, each concerning a sequential phase of the overall processing. For first, we exploit *text categorization* techniques to filter the set of e-mails of interest out of the whole provided collection. Then, we make use of *object matching* algorithms to obtain clusters of related e-mail conversations, from the previous extracted subset. Every cluster is subsequently treated by *text mining information extraction* procedures, in order to find out which tasks e-mail messages are about. Finally, *process mining* is used to abstract process models representing the workflows, which the sets of subsumed tasks were considered traces of. Such sequential multiple-disciplines-based approach is reflected by the multi-layer architecture of the implementation tool, named MAILOFMINE.

The remainder of this paper is organized as follows: Section II will describe the state of the art of the research fields that concern what outlined before. The proposed approach is detailed in the successive Section III. Finally Section IV concludes the paper by outlining future work.

II. BACKGROUND AND STATE OF THE ART

A. Process Mining

Process Mining, also referred to as *Workflow Mining* (see [3]), is the set of techniques that allow the extraction of structured process descriptions, stemming from a set of recorded real executions. Such executions are intended to be stored in so called *event logs*, i.e., textual representations of a temporarily ordered linear sequence of tasks. There, each recorded *event* reports the execution of a *task* (i.e., a well-defined step in the workflow) in a *case* (i.e., a workflow instance). Beware that events are always recorded sequentially, even though tasks could be executed in parallel: it is up to the algorithm to infer the actual structure of the workflow that they are traces of, identifying the causal dependencies between tasks (*conditions*).

The idea to apply process mining in the context of workflow management systems was introduced in [4]. There, processes were modelled as directed graphs where vertices represented individual activities and edges stood for dependencies between them. Cook and Wolf, at the same time, investigated similar issues in the context of software engineering processes. In [5] they described three methods for process discovery: one using neural networks, another with a purely algorithmic technique, the last adopting a Markovian approach. The authors consider the latter two

the most promising approaches. The purely algorithmic approach builds a finite state machine where states are fused if their futures (in terms of possible behavior in the next k steps) are identical. The Markovian approach uses a mixture of algorithmic and statistical methods and is able to deal with noise. Note that the results presented in [5] are limited to sequential behavior.

Most of the nowadays mainstream process mining tools model processes as Workflow Nets (WFNs – see [6]), explicitly designed to represent the control-flow dimension of a workflow. From [4] onwards, many techniques have been proposed, in order to address specific issues: pure algorithmic (e.g., α algorithm, drawn in [7] and its evolution [8], α^{++}), heuristic (e.g., [9]), genetic (e.g., [10]). Indeed, heuristic and genetic algorithms have been introduced to cope with noise, that the pure algorithmic techniques were not able to manage.

A very smart extension to the previous research work has been recently achieved by the two-steps algorithm proposed in [11]. Differently from previous works, which typically provide a single process mining step, it splits the computation into two phases: the first builds a Transition System (TS) that represents the process behavior and the tasks causal dependencies; the second makes use of the state-based “theory of regions” ([12], [13], [14]) to construct a WFN which is bisimilar to the TS. The first phase is made “tunable”, so that it can be either more strictly adhering or more permissive to the analyzed log traces behavior, i.e., the expert can decide a balance between “overfitting” and “underfitting”. Recall, indeed, that event logs are not the whole universe of possible traces that may run: hence, on one hand, the extracted process model should be valid for future unpredictable cases; on the other hand, it should be checked whether such a process model actually adheres to the behavior that most of the gathered traces reflected in the past (we say “most” here to emphasize that a little percentage of the whole log may represent erroneous deviations from the natural flow of tasks). The second phase has no parameter to set, since its only aim is synthesizing the TS into an equivalent WFN. Thus, it is fixed, while the former step can be realized exploiting one among many of the previously proposed “one-step” algorithms (for instance, [9] seems to integrate well).

[2] describes the “ACTIVE” EU collaborative project, coordinated by British Telecom, currently ongoing (due date is February 2011). Such project addresses the need for greater knowledge worker productivity by providing more effective and efficient tools. Among the main objectives, it aims at helping users to share and reuse informal processes, even by learning those processes from the user’s behavior.

B. Object Matching, Record Linkage and Text Mining

1) *Object Matching and Record Linkage*: Object Matching (OM) is the problem of identifying pairs of data objects coming from different sources and representing the same real world object. The aim of OM techniques is typically related to the cleaning of large data sets from erroneous duplicates. Such duplicates usually derive from misspellings, abbreviations, lack of standard formats, or any combination of these factors. Analogous techniques can be used to achieve a *reference reconciliation* (see [15]), namely the identification of related references in complex information spaces where data corresponding to the same reference can be structurally heterogeneous (e.g., e-mail contacts, documents, spreadsheets). If data objects are records, the problem is known in literature as Record Linkage (RL) [16]. Integration of different data sources and improvement of the quality of single sources are only some of the real application scenarios that need to solve the OM problem. The RL version of the OM problem has received considerable attention by the scientific community. Most of the works (e.g. [17], [18]) focus on solving the problem within a relational DBMS: often, in the real world, entities have two or more representations in databases. Duplicate records do not share a common key and/or they contain errors that make duplicate matching a difficult task.

Some techniques for discovering strong entity associations in semi-structured data, such as document meta-data, are also known (e.g., see [19]). Here, the attention is moved towards a novel technique, proposed in [20]. Such a methodology differs from the others in the field, indeed, since (i) it is able to treat not only records but every kind of objects, provided that it is possible to define a distance for them, (ii) it is focused on effectiveness rather than on the ability to manage huge amount of data, (iii) its nature is completely automated. Regarding the third point, there were new unsupervised techniques (such as [17]) already proposed, but none of them, to the best of our knowledge, were fully automated. Indeed, though not requiring exactly a clerically prepared training set, such techniques still depend critically on some external inputs (e.g., human intervention is needed to set crucial parameters for the algorithms in [17]). It achieves such a result by applying a two-phases algorithm. It makes use of statistical models that allow to represent a probability distribution as a convex combination of two distinct probability distributions:

the one stemming from the sub-population of Matches (\mathcal{M}) and the other from that of Unmatches (\mathcal{U}). Thus, the two consecutive tasks are: (i) estimating mixture parameters by fitting the model to the observed distance measures between pairs; (ii) then, obtaining a probabilistic clustering of the pairs into Matches and Unmatches, by exploiting the fitted model. In the clustering step the fitted mixture model is used to search an optimal classification rule such that each pair can be assigned, based on its observed distance value, either to the \mathcal{M} or to the \mathcal{U} class. Such this constrained optimization problem is solved by means of a purposefully designed evolutionary algorithm [21].

2) *Text Mining*: Text Mining, or Knowledge Discovery from Text (KDT) deals with the machine supported analysis of text: indeed, it refers generally to the process of extracting interesting information and knowledge from unstructured text. It is a field in the intersection of related areas such as information retrieval, machine learning, statistics, computational linguistics and, especially, data mining.

Natural language text contains much information that is not directly suitable for automatic analysis by a computer. However, computers can be used to sift through large amounts of text and extract useful information from single words, phrases or passages. Therefore, text mining can be interpreted in the sense of an information extraction activity, i.e., as a restricted form of full natural language understanding, where we know in advance what kind of semantic information we are looking for.

Text mining covers many other topics that are out of the scope of this paper: for a comprehensive survey on it, please refer to [22] or, for a more extended explanation, [23]. A particular discipline of interest that belongs to it is *Text Categorization* (TC, also known as *Text Classification* or *Topic Spotting*), namely, the activity of assigning *categories* (symbolic labels), from a given set, to natural language texts, on the basis of *endogenous* knowledge only (i.e., knowledge is extracted from the documents only and not from other possible external sources). The categories in the given set can be two (*Binary TC*, i.e., a document can belong to a category or its complement) or more. TC is applied in many contexts, such as document filtering and automated metadata generation. For a comprehensive survey on Machine Learning in Automated Text Categorization, the reader can refer to [24].

As a successful example of application, we want to report here a case that deeply relates with the research proposal of this paper: [25] proposes a method employing text mining techniques to analyze e-mails collected at a customer center. The method uses two kinds of domain-dependent knowledge: one is a key concept dictionary manually provided by human experts and the other is a concept relation dictionary automatically acquired by a fuzzy inductive learning algorithm. Based upon the work exposed in [26], the depicted method takes as input the subject and the body of an e-mail, decides a text class for the e-mail, extracts key concepts from e-mails and finally presents their statistical information as well.

C. Two related famous industrial tools for e-mail content processing

Here are recalled just two famous examples of e-mail content processing tools that are related to our research proposal (see Section III).

The first is the *Priority Inbox*¹ for *GMail*, i.e., an add-on that automatically infers which e-mail messages are to be considered “important” and which are the so called “bologna” (i.e., less important mail) or “Bacn” (e.g., news alerts, social network messages and that kind of messages that the user is subscribed to or interested in but can stay in the inbox for a long time, unread). On the basis of such a partitioning within the e-mail messages set, the user interface shows the two groups separately, so that she can see the important messages at a glimpse, on top of a vertical tab that graphically distinguishes them from the others.

The second is the *Xobni*² plug-in for *Microsoft Outlook*. One of the feature it provides is the trace of e-mail contents with the ability to create connections among people that the user has conversed with and to analyze their behavior in terms of time of response, frequency of communications and so on.

III. MAILOFMINE

A. A definition of the context

As described in [2], knowledge workers make use of their practice and competence for leading complex tasks to a successful ending. Either if they work alone or in a collaborative context, each of them is an expert in her own

¹for further readings, the reader may visit the related Gmail Help page <http://mail.google.com/mail/help/priority-inbox.html> or the Google software engineer Doug Aberdeen’s post on the official Google blog, at <http://gmailblog.blogspot.com/2010/08/email-overload-try-priority-inbox.html>

²<http://www.xobni.com/>

fields. Thus, when they encounter situations which they previously coped with, they act on the basis of previous experience that lead their decisions. It is clear then that they figure out in their mind some informal process schemes that follow and refine time by time. They typically do not want to formally depict in a process model such “artful” ones (see Section II-A and [1]), since they are short and continuously subject to changes, whereas tools for Business Process Management are designed for large-scale workflows involving companies or many actors at a time.

“Classical” Process Mining aims at discovering actual workflows from event logs (see Section II-A) in order to compare them with the expected ones (and thus refine the ideal, or correct the real ones). Here an advantage for the target users may be discovering agile process models, in order to help them organizing the following conformant cases, and tidy up, though with a simple behavior modeling, their personal workflows. This way, they can share them or find out what to do to perform best, without having spent either a single minute to draw by hand a formalization of what they do in their everyday work. Of course, once the workflows are extracted, an additional outcome such as identifying bottlenecks in executions should be valuable to the users.

In collaborative contexts, these targeted users share their information and knowledge work outcomes with other knowledge workers. Typically, they achieve such exchange by means of several e-mail messages. For instance, the reader may figure out a European Project manager, who collaborates with other professionals working in different countries. Outcomes such as deliverables and reports, or information like dates for meetings, timetables and so on are all delivered through several e-mail-based conversations. They are provided as content bodies, inside attachments, or both.

Hence, such conversations are actual traces of running processes that knowledge workers adhere to and lead. Each message is a *semi-structured document*, due to the presence of free text in content bodies that is enriched by header fields and a (usually) related subject specifying the matter. It can also be considered, under a different point of view, as a rich log event (see Section II-A), since it records, at least, (i) the actors involved (as sender and recipients); (ii) the absolute timings (since every e-mail message header has a “date” field); (iii) some direct causal relations (e.g., think of “Re:” or “Fwd:” prefixes in e-mail subjects); (iv) the activity matters (typically, in subjects); (v) the possible activity outcome (e.g., in attachment).

Not only the first two characteristics come out for free by e-mail headers, but also the third, actually, since when no “Re:” or “Fwd:” is in the e-mail subject, a unique identification string of the original message that the reply refers to is specified inside “References” and “In-Reply-To” header fields. Still regarding this point, we highlight that there are only “some” causal relations mapped this way: indeed, it may happen that many e-mail messages, though related to the same argument or activity, might have subjects that do not match at all, and are not written in response of any of the previously exchanged messages. But the system should understand that subjects such as “Deliverable v1” and “Deliverable: first version” might refer to the same topic. Although, a typical attitude of users that comes to our help, in clustering related mail, is that they often include quoted text in responses (in such cases, the matching activity can be partially based upon a string equality check). These cases are more difficult to cope with, anyway, as the extraction of the forth and fifth characteristics becomes harder: they all require a text analysis of subjects, message bodies and attachment names.

B. The approach

Of course, not all the messages will be taken into account: a previous filtering on the whole set to extract the ones of interest only is needed. As knowledge workers are not expected to have time to formalize their processes, we can not either suppose to rely on predefined logical rules that describe whether a document belongs to the domain of interest or not. In general, it is easier to manually classify a set of documents than to build and tune a set of rules, since it is easier to characterize a concept extensionally (i.e., to select instances of it) than intensionally (i.e., to describe the concept in words, or to describe a procedure for recognizing its instances). Thus, the application of Machine Learning techniques for an automated binary Text Categorization (see [24] and Section II-B.2) perfectly fits our case: we can consider a category c_I of e-mail messages of interest and a complementary $c_{\bar{I}}$ of e-mail messages to be considered out of scope. The knowledge worker will be asked to distinguish, in an *initial corpus* of documents, which are *positive examples* (i.e., that belong to c_I) and *negative examples* (i.e., that do not belong to c_I – thus, they are classified under the $c_{\bar{I}}$ category). This set is then used to train the *learner* and test its effectiveness (like in the typical *train-and-test* approach).

Once the *initial corpus* is provided, information retrieval techniques for *indexing* messages contents (i.e., reducing

e-mail texts to a weighted vector of meaningful terms³) are applied, in order to allow the following inductive construction of *text classifiers*, whose output is a mapping between e-mail messages and either c_I (what we keep for the further processing to take place on) or c_T .

The next step is the division of such a unique set of e-mail messages into many *clusters* of related e-mail messages. As aforementioned, e-mail messages can be considered as *semi-structured data*. Thus, the scientific results reported in Section II-B.1 can be effectively exploited to this purpose. Indeed, [15] already mentions the suitability of its proposed technique to be adopted for Personal Information Management (PIM) usage.

Of course, it is a common attitude to include in single messages multiple themes or activities: for instance, the same e-mail can contain both a suggestion for the whole project developers community about the next review meeting deliverables, and a recommendation for the organizer about the hotel location. Thus, produced clusters will not be parts of a *tessellation*, since the boundaries of such aggregation sets can not be outright: a single message can pertain multiple matters and thus belong to many clusters at a time.

Despite these difficulties, we can take advantage of a very important assumption in this context. [20] (see Section II-B.1) required that the noise introduced by errors was not wild but limited enough to distinguish among Matches and Unmatches (\mathcal{M} and \mathcal{U}). Adopting a signal theory metaphor, it means that the noise is required to act as a small perturbation and not like a destructive interference over \mathcal{M} and \mathcal{U} signals. According to a probabilistic point of view, the probability distribution of the distance between data objects is significantly different inside the \mathcal{M} and \mathcal{U} classes. Turning back to our example, it does perfectly make sense, since the source of the e-mail messages is certainly not a stochastic process that generates subjects by concatenating words taken uniformly at random. E-mail subjects, for instance, are made by people that want the recipient to understand at a first glimpse what they are talking about: this is why typically we use prefixes among square braces such as “[SM4All] D4.3 next release date” or “[oonar] Upcoming show delayed”. Thus, we can state such hypothesis with a high degree of certainty, due to the origin of the data objects generation process itself.

Once e-mail clusters are obtained, a further discovery step is required on each of them: extracting what we may call *emerging process*. We start by bringing out the activities described in such e-mail messages sets (Information Extraction, IE) on the basis of a key concept dictionary, following the methodology described in [25] and [26] (Section II-B.2).

At that point, we exploit the previous extraction of causal dependencies among e-mail messages (which are sequential by nature) and treat the e-mail clusters as traces of the process (or event logs, to adopt the classical Process Mining term – see Section II-A). On top of them, we can run a two-steps algorithm like the one described in [11], exploiting the possibility to customize its first phase parameters, i.e., the one where the behavioral Transition System is extracted. As described in Section II-A, such parameters make the Transition System that represents the process behavior more (overfitting) or less (underfitting) strictly adhering to the gathered traces – conversely, less or more permissive with respect to future variations in the next executions. Indeed, emerging (artful) processes are agile and subject to changes, thus it is preferable to draw models that are more underfitting than overfitting.

The noise that affects this variable traces may be addressed by the combination of it with the technique in [9], that, as aforementioned, seems to correctly integrate with it.

Then, the final output, namely the process models that represent the previous traces implicitly recorded by means of e-mail based conversations, is returned.

C. The layers-based architecture of MAILOFMINE

As depicted in Figure 1, MAILOFMINE is intended to be divided into layers that act in series like in a waterfall model. Each layer is built to achieve a specific task (e-mail reading, e-mail filtering, e-mail clustering, tasks extraction, workflow mining). The initial input is a file storing a whole repository of e-mail messages. The final output is a set of process models inferred from the given input. In the middle, every layer is drawn to receive as input the result that comes from the upper one, starting from the raw mail, and in turn take the output to the lower one, down to the final process models.

The very first task is accomplished by a plug-in based system, that must be able to extract a common format for e-mail messages, stemmed from heterogeneous sources: for example, Microsoft Outlook, Mozilla Thunderbird, Evolution Mail files use different storage formats, but the system must be able to manage them all.

³in this context, weights are going to be differently scaled according to the relevance of words in the e-mail message structure – e.g., terms in subjects are generally more important than in content bodies

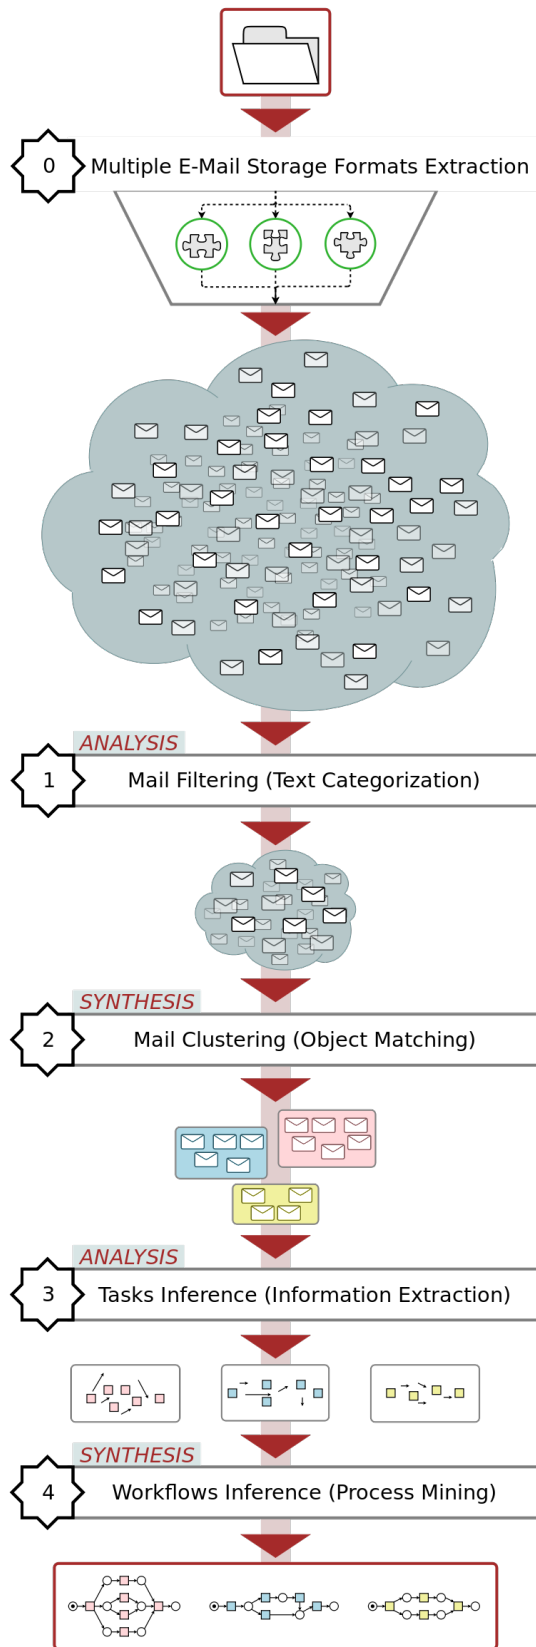


Fig. 1. The MAILOFMINE architecture

In Figure 1, such task corresponds to the layer 0 (*Multiple E-Mail Storage Formats Extraction*).

From that point on, the proposed approach follows this pattern, founded on a double analysis/synthesis passage:

- 1) (*Mail Filtering*) *analysis* on the set of heterogeneous e-mail messages to filter irrelevant messages out (Text Categorization, Information Retrieval – see Section II-B.2);
- 2) (*Mail Clustering*) *synthesis* on the set of relevant messages, to cluster related ones into homogeneous⁴ sets (Object Matching – see Section II-B.1);
- 3) (*Tasks Inference*) *analysis* on each cluster, to extract the tasks out (Information Extraction – see Section II-B.2);
- 4) (*Workflows Inference*) *synthesis* on tasks, to build a process model that conforms with the subsumed trace (Process Mining – see Section II-A).

IV. CONCLUSIONS

Throughout this paper, we have proposed a novel approach to infer artful business process models, that is to say agile workflows formal representations, from knowledge workers e-mail storage files. It is the basis for the design and implementation of the tool MAILOFMINE. This approach can take advantage of many previous works that succeeded in the heterogeneous research fields which are involved in this context (Text Mining, Object Matching, Process Mining). Another point of strength of the proposal is that it requires almost no effort for workers, because of its automated nature. Despite this, its advantages are valuable to companies that decide to make use of MAILOFMINE.

In future work, we want to address some interesting open issues. First of all, a challenge about cooperative activities: they may involve many knowledge workers at a time, and it can happen that a task, say *DoIt*, that Mr A. Bloggs demands to Mrs B. Doe, is in turn redirected to Mr C. Smith. Thus, Mr C. Smith fulfills a series of tasks, in collaboration with Mrs B. Doe, which Mr A. Bloggs is unaware of. It could be interesting, then, to indagate on how to relate these activities that are traced by separated e-mail sets (one belonging to Mr A. Bloggs, the other to Mrs B. Doe), so to unify them under a single workflow model.

Another point to cope with is the question of privacy: e-mail messages may contain sensible private information that the single knowledge worker, or the company, might not want to be processed. At this stage of the research, though, we can suppose that treated data are public at least to the company that the knowledge worker works for, since we consider just the company mailbox and not her personal one, and the inferred information would not be presented to other people than who the company wants to involve.

REFERENCES

- [1] C. Hill, R. Yates, C. Jones, and S. L. Kogan, “Beyond predictable workflows: Enhancing productivity in artful business processes,” *IBM Systems Journal*, vol. 45, no. 4, pp. 663–682, 2006.
- [2] P. Warren, N. Kings, I. Thurlow, J. Davies, T. Buerger, E. Simperl, C. Ruiz, J. M. Gomez-Perez, V. Ermolayev, R. Ghani, M. Tilly, T. Bösser, and A. Imtiaz, “Improving knowledge worker productivity - the active integrated approach,” *BT Technology Journal*, vol. 26, no. 2, pp. 165–176, 2009.
- [3] W. M. P. van der Aalst, “The application of petri nets to workflow management,” *Journal of Circuits, Systems, and Computers*, vol. 8, no. 1, pp. 21–66, 1998.
- [4] R. Agrawal, D. Gunopulos, and F. Leymann, “Mining process models from workflow logs,” in *Advances in Database Technology – EDBT’98*.
- [5] J. E. Cook and A. L. Wolf, “Discovering models of software processes from event-based data,” *ACM Trans. Softw. Eng. Methodol.*, vol. 7, no. 3, pp. 215–249, 1998.
- [6] W. M. P. van der Aalst, “Verification of workflow nets,” in *ICATPN*, ser. Lecture Notes in Computer Science, P. Azéma and G. Balbo, Eds., vol. 1248. Springer, 1997, pp. 407–426.
- [7] W. M. P. van der Aalst, T. Weijters, and L. Maruster, “Workflow mining: Discovering process models from event logs,” *IEEE Trans. Knowl. Data Eng.*, vol. 16, no. 9, pp. 1128–1142, 2004.
- [8] L. Wen, W. M. P. van der Aalst, J. Wang, and J. Sun, “Mining process models with non-free-choice constructs,” *Data Min. Knowl. Discov.*, vol. 15, no. 2, pp. 145–180, 2007.
- [9] A. Weijters and W. van der Aalst, “Rediscovering workflow models from event-based data using little thumb,” *Integrated Computer-Aided Engineering*, vol. 10, p. 2003, 2001.
- [10] A. K. Medeiros, A. J. Weijters, and W. M. Aalst, “Genetic process mining: an experimental evaluation,” *Data Min. Knowl. Discov.*, vol. 14, no. 2, pp. 245–304, 2007.
- [11] W. van der Aalst, V. Rubin, H. Verbeek, B. van Dongen, E. Kindler, and C. Günther, “Process mining: a two-step approach to balance between underfitting and overfitting,” *Software and Systems Modeling*, vol. 9, pp. 87–111, 2010.
- [12] J. Cortadella, M. Kishinevsky, L. Lavagno, and A. Yakovlev, “Synthesizing petri nets from state-based models,” in *Computer-Aided Design, 1995. ICCAD-95. Digest of Technical Papers, 1995 IEEE/ACM International Conference on*, nov. 1995, pp. 164–171.

⁴Here “homogeneous” has to be intended with respect to the activities to serve for the purpose of task inference.

- [13] —, “Deriving petri nets from finite transition systems,” *Computers, IEEE Transactions on*, vol. 47, no. 8, pp. 859–882, aug. 1998.
- [14] J. Desel and W. Reisig, “The synthesis problem of petri nets,” *Acta Informatica*, vol. 33, pp. 297–315, 1996.
- [15] X. Dong, A. Y. Halevy, and J. Madhavan, “Reference reconciliation in complex information spaces,” in *SIGMOD Conference*, ACM, 2005, pp. 85–96.
- [16] A. K. Elmagarmid, P. G. Ipeirotis, and V. S. Verykios, “Duplicate record detection: A survey,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 19, pp. 1–16, 2007.
- [17] S. Chaudhuri, V. Ganti, and R. Motwani, “Robust identification of fuzzy duplicates,” in *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on*, apr. 2005, pp. 865–876.
- [18] S. Guha, N. Koudas, A. Marathe, and D. Srivastava, “Merging the results of approximate match operations,” in *VLDB '04: Proceedings of the Thirtieth international conference on Very large data bases*, pp. 636–647.
- [19] N. Sarkas, A. Angel, N. Koudas, and D. Srivastava, “Efficient identification of coupled entities in document collections,” in *ICDE 2010*, pp. 769–772.
- [20] D. Zardetto, M. Scannapieco, and T. Catarci, “Effective automated object matching,” in *ICDE 2010*, pp. 757–768.
- [21] Z. Michalewicz, *Genetic algorithms + data structures = evolution programs (2nd, extended ed.)*. New York, NY, USA: Springer-Verlag New York, Inc., 1994.
- [22] A. Hotho, A. Nürnberger, and G. Paaß, “A brief survey of text mining,” *LDV Forum - GLDV Journal for Computational Linguistics and Language Technology*, vol. 20, no. 1, pp. 19–62, May 2005.
- [23] M. W. Berry and M. Castellanos, *Survey of Text Mining II: Clustering, Classification, and Retrieval*, M. W. Berry and M. Castellanos, Eds. Springer, September 2007.
- [24] F. Sebastiani, “Machine learning in automated text categorization,” *ACM Comput. Surv.*, vol. 34, no. 1, pp. 1–47, 2002.
- [25] S. Sakurai and A. Suyama, “An e-mail analysis method based on text mining techniques,” *Appl. Soft Comput.*, vol. 6, no. 1, pp. 62–71, 2005.
- [26] S. Sakurai, Y. Ichimura, A. Suyama, and R. Orihara, “Acquisition of a knowledge dictionary for a text mining system using an inductive learning method,” in *Proceedings of IJCAI 2001 Workshop on Text Learning: Beyond Supervision*, 2001, pp. 45–52.