# Beyond life streams: activities and intentions for managing personal digital memories

Jérôme Picault, Myriam Ribière and Christophe Senot

Bell Labs, Alcatel-Lucent,
Route de Villejust, 91620 Nozay, France
{jerome.picault, myriam.ribiere, christophe.senot}@alcatel-lucent.com

**Abstract.** In this paper, we expose a set of initial ideas related to an innovative way of structuring and organizing personal information. Indeed, users have to deal with a huge amount of information either coming from social connections, collected on the Web or generated by them. This phenomenon leads to new research challenges. In particular, how to structure, organize, and classify this personal information in order to *better manage the user's digital memory*? In this position paper, we present the concepts of activities and intentions as means for the user to structure efficiently all his past information, but also help him in the future, for example by suggesting relevant events, anticipating his information needs or providing opportunities to satisfy latent desires.

**Keywords:** personal information management, digital memory, timeline, activities, intentions, information container, anticipation of information needs

## 1 Introduction

Nowadays, due to the increasing development of communication technologies, social media, massive content production or diversification of knowledge sources, users tend to be overwhelmed with a huge volume of personal information such as emails, photos, e-books, blogs, social feeds, or various documents. These data are either created by them (e.g. through lifestream aggregators such as FriendFeed[1], Lifestrea.ms[2], etc.) or by others (e.g. through social services such as Twitter, Facebook). All this information are from near or far sighted centered on the user life - social exchanges, information gathered on the web, etc. and constitute what we call the *user's digital memory*.

However, today this information is only captured, stored, but not very-well organized from users' point of view and thus is not used as much as it could be. This phenomenon induces the following research challenges. First, how to keep track of important events? Which semantic structure would allow users to find the right information when needed and *organize their digital memory* properly? A second

[1] http://friendfeed.com/
[2] http://lifestrea.ms

challenge deals with the anticipation of information needs: we believe that a user-centric semantic organization of the digital memory may help the user in his current or future information needs.

Thus, we present some initial ideas towards a new way of indexing and structuring users' digital memories. Section 2 gives an overview of existing models and solutions for managing personal information. Section 3 introduces the notion of *activity* as a key concept to structure personal memory. Section 4 gives some clues on how to go beyond this first layer, by enriching this semantic structure with an additional meta-layer of information organization, based on the notion of *intention*. Section 5 illustrates how this intention-based personal information management model can be instantiated for improving content filtering and opportunistic recommendations.

## 2 Related art

The problem of organizing and structuring personal information is not new. This field has already been studied in the domain of personal information management, and several paradigms of document organization have been identified. Temporal paradigms organize documents according to a time line. This is the way how life streams[3] [4] are usually presented to the user. Life logs projects such as Microsoft MyLifeBits [6] aim at storing in a database a massive set of every activity and relationship a person engages in (books, music, photos, video, office documents, email, phone calls, meeting, web pages, etc.) and structure them according to two axes: time and life (personal vs. professional). However, according to Gemmel, "the collection is so large that the user cannot remember much of the contents, and will never *use* them." Some solutions use a spatial representation, such as in Data Mountain [3], a logical paradigm, based on keyword or content assignment, such as in Haystack [8], or a combination of dimensions such as TimeScape [10]. Search engines such as Google Desktop[4] are an alternative to structured information, but in the case of digital memory, they do not rely on an index with the right granularity from the user's point of view. Other approaches propose manual ways of structuring information. For example, Pearltrees[5] proposes to users a way to keep content they find everyday on the web and to let them structure their information through trees.

Finally, some research has been carried out in the perspective of anticipating information needs. Thus PackHunter [5] is a collaborative tool to share with a group of users web trails, which allow jumping to pages visited by others, etc.

However existing work are limited to an organization through a structure (e.g. timeline, hierarchical) with limited semantics which does not correspond effectively to the way users behave. So, there is a need to better structure this digital memory to make it useful and usable to the user. In this paper, we propose a solution using episodic memory [12] with two different layers: *activities* of the user and his *intentions*. We detail these concepts in the following sections.

---

[3] Cf. http://www.readwriteweb.com/archives/35_lifestreamin_apps.php for examples

[4] http://desktop.google.com

[5] http://www.pearltrees.com

## 3 Activity-based personal information management

In the human memory process, two main steps are fundamental: the *acquisition* (retention) and *recall*. Tulving in [12] showed that episodic memory, which receives and stores information about temporally-dated episodes and spatio-temporal relations among them, is a faithful record of a person's experience. Recalling a piece of information is easier when the user can remind himself in time and space. Besides, according to a recent study [1], users tend to think about and classify their personal information in terms of activities more than they do in terms of information type or just time. The positioning of information in a three dimension space (time, place and people) is already envisioned as a de facto standard to structure life logs [2]. Activities are adding to the event notion a semantic context, which defines another essential dimension for representing the user's daily life. Therefore, they may constitute a good paradigm to manage digital memory.

Thus, we can think of organizing user activities in a temporal way through a timeline of activities. This organization shows how activities can also address different research areas in the domain of multimedia content consumption according to their position in the timeline.

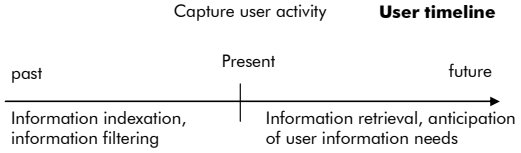


**Figure 1. Usage of the activity concept in the user timeline**

As presented in Fig. 1, the "present" part of the timeline consists in capturing the current user activity. Capturing user activity is a research area in itself, where different related work [13] could be used. The "past" side of the timeline enables to index content and people and keep track of user memory. Past activities are reference marks (i.e. episodes) for people to find information and content, and a support for social information sharing even after their end.

More formally, we define an *activity* as a personal activity (digital or not) or as a user's perception of a given social activity or event. Based on this definition examples of activity can be: reading a book and making notes and comments, or meeting someone in a conference and exchanging information, collecting multimedia content related to a user activity. An activity is composed of the following main properties:

- *A set of content* that the user has generated, consumed or bookmarked in the context of the activity. A consumed content can be any type of multimedia content or web bookmarks. A user generated content can be an important piece of information written about the user activity (document, comments and annotations, notes) or any interaction captured during the activity (phone call, IM, email, chat, or interactions through social media applications).
- *A semantic context* is inferred from the set of content. It is a key enabler for the awareness of the activity community, and for further information classification.
- *A social context* of the activity is the list of people that are sharing this activity (implicitly people around the user), or people following this activity (explicitly

defined by the user or gathered from interaction traces related to the activity semantic context).

- *A spatio-temporal context* of the activity. Time and place are the two dimensions that can be used to identify typical user contexts such as "at home", "at work", "on the move" or simply to position the activity in space and time for a better user recall.
- *A status*. An activity can have three distinct statuses: ended, ongoing and in mind. The *ended* status means that the activity belongs to the past and that it can be used as a piece of memory. An *ongoing* activity constitutes a recipient for new incoming information. An *in mind* activity is not yet started; this is used to describe latent activities that may be recommended in the future to the user.

The role of the activity is twofold: (1) a working space environment where all pieces of information (documents, emails, bookmarks, etc.) and pertinent contacts are gathered within a same structure, becoming a relevant index (on people and content) for structuring the user digital memory, and (2) a representation of the social environment of an activity, helping people to share information in a controlled way and to get information from their social networks around this activity.

## 4 Intention-based personal information management

The management of personal information through the notion of *activity* provides already a first organization layer. However, it does not consider interdependencies between activities. So, we propose to extend this semantic structure with the concept of *information container* as a semantic entity that encapsulates a set of coherent activities that are correlated according to the different activity dimensions. Ultimately, the observation of correlated activities may denote user's *intentions* in time and space, that describe what the user wishes to achieve at a high and pragmatic level [9].

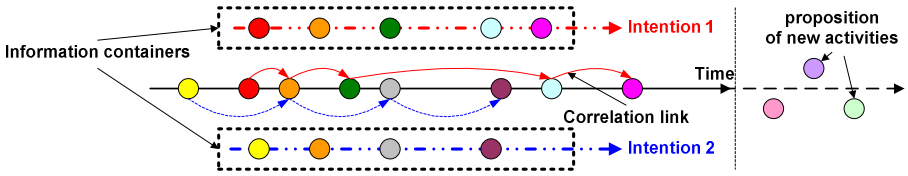


**Figure 2. Notions of information container and intentions**

The "past activities" of the user (Fig. 1) are structured through an additional layer, an *information container* (Fig. 2). The latter is composed of a *set of activities* and one or several properties describing the nature of the correlations between activities:

- A *content link* reflects the shared semantic context between all the activities;
- A *social link* contains the common contacts or social context (family, colleagues, etc.) between the activities;
- A *logical link* indicates how an activity relates to others. Possible links are *causality* (an activity is the follow-up of another one), *temporality* (an activity is the repetition of another one), etc.

Based on the analysis of these semantic links an *intentional link* can be inferred between the activities present in a given information container. An *intention* can be seen as the high level "glue" between several activities and describes the set of activities as a whole unit as in [11]. Contrary to previous works such as [14], we do not express an intention by a formal plan; nevertheless at a high level, it may be described thanks to an action verb, a complement and an intensity reflecting its certainty or feasibility.

In addition to its structuring role of past activities, the information container can be seen as an *active* recipient, in charge of helping the user towards the "future" side of the timeline (Fig. 1). Indeed, intentions act as a guideline that leads the user involvement through various activities. Thus, the knowledge of existing intentions can be used to recommend information associated to activities belonging to the container or which are completely new for the user. Additional exploitations of intentions can be envisaged through some forms of collaborative mechanisms for different purposes, for example: 1) to enrich / suggest activities to a user based on the detection of a common activity pattern with other users – this may help the user to find faster what he needs; and 2) to build a dynamic social network around people having a common intention, in order e.g. to help them to realize it jointly [7].

Moreover, an information container is not static, it may grow by acting as a kind of agent that enriches the information it contains with coherent new elements coming from specified *information streams* (email, IM, RSS feeds, notifications etc.).

The iPIM ontology (Fig. 3) describes more formally the concepts described above.

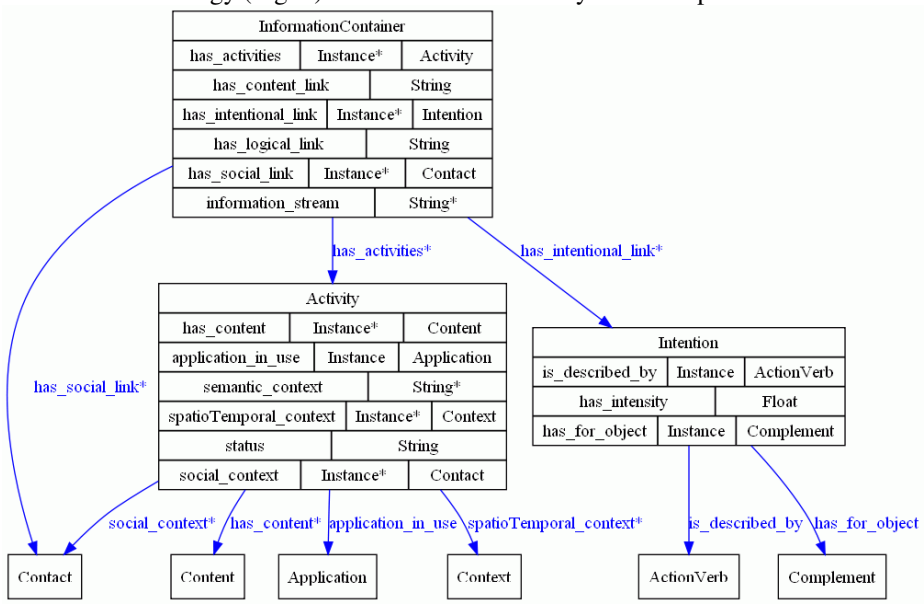


**Figure 3. Overview of the iPIM ontology**

This vision raises many research questions:

– *Construction of information containers*: how to correlate activities to build those information containers? When a new activity appears, to which information

containers should it belong to? Is it just a clustering problem? How are we able to modify the information containers if we detect an anomaly?

– *Identification of intentions*: detection of a precise user intention may be difficult. A possible solution is to use a learning model, where the user at the beginning explicitly describes the intention associated to an information container. After a while, the model could suggest the user relevant action verbs and extract knowledge from social and/or content links as complements. Another possibility would be to use a collaborative model which compares information containers of one user to the ones of other users to suggest possible intention labels.
– *Monitoring of intentions*: how to infer the progress with respect to an intention or an information container?
– *Usage and acceptance* – how to capture or confirm user activities (what is the part of automation and manual declaration) and present information containers to users?

## 5 Exploitation of iPIM to improve recommendations

In this section, we express through a scenario how the semantic structure described above can be used, in particular as a way to go beyond classical recommendation systems. Fig. 4 summarizes the different user's activities that occur during the scenario. This scenario shows how a system can monitor in real-time different user's activities, such as watching a documentary, browsing the web, meeting friends, etc. and the nature of the resulting intentions over the time.

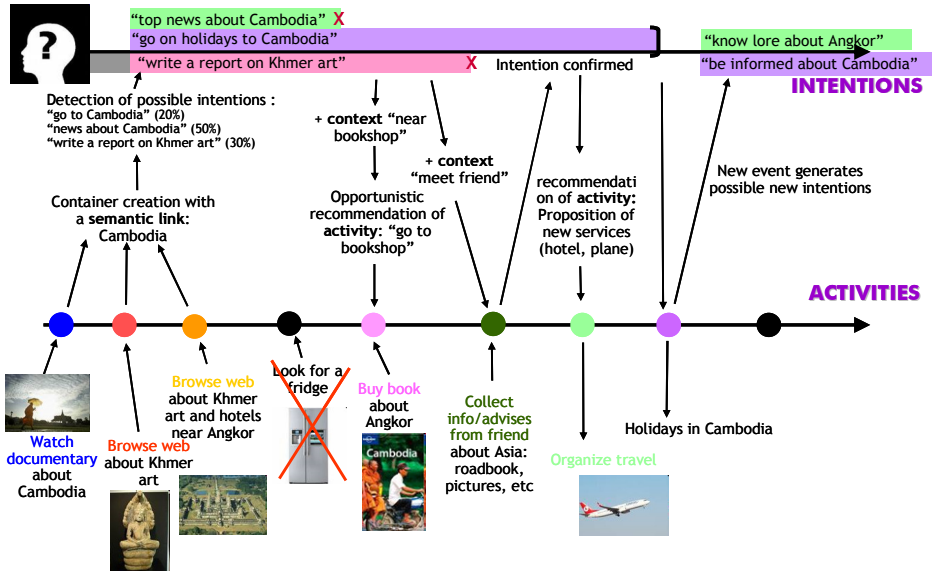


**Figure 4. Illustrative scenario**

The scenario can be decomposed through three main axes:

– *Activity indexing*: from the user timeline several *activities* are detected and then indexed by the system based on their *contexts* (e.g. for the activity "watch a

documentary" the *semantic context* is a documentary reference and its *status* is equal to ended).

- *Building of information containers*: in the scenario the construction of the *information container* is quite easy as most of the activities share at least the same *content link* related to Cambodia (except the "search of a new fridge"). By correlating more precisely the existing activities with past activities from other users (based on a collaborative approach) a *logical link* can also be inferred from the same information container (e.g. travel booking).
- *Intention detection*: within the Cambodia information container several user's *intentions* may be inferred based on the underlying information container links. For each intention the system tries to formalize its meaning (e.g. verb + complement form). In addition to the previous treatment a certainty degree is computed reflecting the current intention relevance according to several parameters (context, activities, etc). While new activities appear, the potential intentions are refined or simply removed from their information container. Thus, in Fig. 4, at the beginning three intentions were inferred, and at the end only one seems to be relevant: "go on holidays to Cambodia". Nevertheless first inferences are already useful for proposing relevant content or services – especially in an opportunistic way, where the user may not have thought about himself (e.g. meet a friend). Another interesting property of an information container is that even if an intention is ended (e.g. the holidays are now finished) it is still open to new activities; thus new intentions can emerged (e.g. know more about Angkor).

## 6 Conclusions and perspectives

In this paper we presented initial steps towards a new paradigm for structuring and organizing personal information. We believe that the concept of intention provides a relevant conceptual framework to anticipate user information needs, and opens the way to new service opportunities for context-aware multimedia content access and delivery. However we still need to understand if semantic and social contexts are appropriate indicators of relationships between activities to deduce user intentions. This can be learnt through a diary study, and further with experimentations on real captured activities. This new way of managing personal information may have a real social impact, e.g. by providing opportunistic interaction with people driven by intentions. To go a step further in the social exploitation, we envisage the use of collaborative algorithms for better inferring intentions through the co-relation of activities.

Besides, intentions could generate spontaneous social networks, i.e. communities of people sharing the same kind of intentions, which will ease social interactions, and help them collectively find the right path to fulfil it (joint realisation of an intention). A further perspective of this work could be the creation of communities of knowledge, based on people promoting their information container, and sharing with the community the solution they found. We could capitalize on this community of knowledge to identify similar patterns of activities to fulfil typical intentions, and propose appropriate compositions of services that can be seen as an intention-based service mash-up.

## References


1. Bergman, O., Beyth-Marom, R., & Nachmias R.: The user-subjective approach to personal information management systems. Journal of the American Society for Information Science and Technology 54 (9): 872-78. (2003)
2. Byrne, D., Lee, H., Jones, G. and Smeaton, A.F.: Guidelines for the presentation and visualisation of lifelog content. In Irish Human Computer Interaction Conference, (2008).
3. Cockburn, A., & McKenzie, B.: 3D or not 3D? Evaluating the Effect of the Third Dimension in a Document Management System. Conference on Human Factors in Computing Systems, Seattle, Washington, USA. (2001)
4. Freeman, E. & Fertig, S.: Lifestreams: Organizing your electronic life. In R.Burke (Ed.), AI Applications in Knowledge Navigation and Retrieval. AAAI Press. (1995)
5. Furmanski, C., Payton, D. & Daily, M.: Quantitative Evaluation Methodology for Dynamic, Web-based Collaboration Tools. Proceedings of the 37th Hawaii International Conference on System Sciences. (2004)
6. Gemmell, J., Bell, G. & Lueder, R.: MyLifeBits: a personal database for everything, Communications of the ACM, vol. 49, Issue 1, pp. 88.95. (2006)
7. Gold, N. & Harbour, D.: Cognitive Primitives of Collective Intentions: Linguistic Evidence of our Mental Ontology. Queen Mary, University of London. (2008)
8. Karger, D. R., & Quan, D.: Haystack: A User Interface for Creating, Browsing, and Organizing Arbitrary Semistructured Information. Conference on Human Factors in Computing Systems, Vienna, Austria. (2004)
9. Kemke, C.: About the Ontology of Actions, Technical Report MCCS-01-328, Computing Research Laboratory, New Mexico State University. (2001)
10. Rekimoto, J.: TimeScape: A time-machine for the desktop environment. Conference on Human Factors in Computing Systems, Pittsburgh, Pennsylvania, USA. (1999)
11. Searle, J. R.: The Intentionality of Intention and Action. Cognitive Science vol. 4. (1980)
12. Tulving, E.: Elements of Episodic Memory. Oxford: Clarendon Press. (1983)
13 Voida, S.: Activity Representations and Tagging in Support of Resource Organization and Collaboration. PhD thesis, Georgia Institute of Technology. (2008)
14. Zamparelli, R.: Intentions are plans plus wishes (and more). AAAI Symposium. (1993)