Active Artifacts as Bridges between Context and Community Knowledge Sources.

Federico Cabitza
Università degli Studi di Milano-Bicocca
20126, Viale Sarca 336
Milan, Italy
cabitza@disco.unimib.it

Carla Simone
Università degli Studi di Milano-Bicocca
20126, Viale Sarca 336
Milan, Italy
simone@disco.unimib.it

ABSTRACT

The aim of the paper is twofold: i) understanding how to provide additional information that is reflective of current organizational context in knowledge production and use; ii) proposing an architectural solution that can be applied to this need. To this aim, we introduce the concept of Active Knowledge Artifact (KA), i.e., an electronically augmented (i.e., active) artifact that puts together the archival functions of artifacts belonging to organizational ISs with context- and content-aware functionalities to promote collaboration awareness and support knowledge management. Through a case study in the hospital domain, we illustrate an approach where documents are augmented with information intended to support context interpretation and evoke the knowledge that actors need to coordinate their actions in that context. The autonomous provision of Awareness Promoting Information (API) and Knowledge Evoking Information (KEI) by means of modular and reactive mechanisms embedded in each KA is what makes KAs active computationally.

Categories and Subject Descriptors: H.5.3 Group and Organization Interfaces: Computer-Supported Cooperative Work.

General Terms: Design, Human Factors.

Keywords: Active Artifacts, Knowledge Artifacts, Collaboration Awareness, Knowledge Evocation.

1. THE SCOPE OF INVESTIGATION

The persistent popularity of the term 'community', which has been used to refer to even deeply different situations, requires starting from a clear connotation of the kind of community we are considering. To this aim, we are not proposing any sort of taxonomy; rather, we aim to define the background of our study to put its findings in the appropriate context. In this paper, we consider communities of professionals that work in relatively small groups, belong to a greater institution delimiting their autonomous behavior and defining their objectives, and cooperate through a well defined and predefined set of artifacts; these contain the information produced to achieve both the organizational objectives and the needed coordination. For these communities, the relation between institutional goals, their constraints and the local practices needed to reach the necessary effectiveness is a unavoidable part of the domain, both from the organizational and technological point of view. Our study shows that this relation is made tangible by the usage of a web of artifacts; these, at the same time, reify the institutional business rules within an organizational Information System and support the everyday work and its contingent needs. In other words, these artifacts are the point where Information Systems (IS) and Knowledge Management (KM) have to reach a factual compromise.

In the last years, we have studied a number of Information Systems deployed in clinical settings, in order to consider the pragmatic point of view of their actual users: doctors, nurses and hospital employees. The study helped shed light on the double nature of the artifacts compounding an IS: these artifacts can be characterized not only in terms of the conceptual categories of the IS research [17], but also from a KM perspective. In fact, users of an organizational IS see its artifacts also as Knowledge Artifacts (KA), that is, as artifacts that reify or evoke some organizational and/or domain-dependent knowledge [15], either in its design, in its use or in regards to both. Therefore, our point is that designers must conceive and treat these artifacts as such. To this aim, the conceptualization of knowledge creation proposed by Nonaka and Takeuchi [30] is a good reference point since it clarifies the role of tacit and explicit knowledge in operational terms and the role of technological solutions in the management of knowledge sources. In fact, while Information System technologies aim to organize and manage information (i.e., explicit knowledge) with the ultimate goal of reaching the maximum level of data and information quality [7], KM technologies aim to transform information into forms of usable and useful knowledge. So far, various techniques have collected a lot of research and design efforts: information wrapping and extraction, data mining, process mining, business intelligence tools (and the like) are typical techniques and tools that support the phases of knowledge externalization and combination [30]. The same is not true for tacit knowledge: often, different communication means (from more traditional e-mail to blogs, wikis, etc.) are proposed as the only means to deal with this more problematic kind of knowledge. All these solutions support socialization, but they are usually proposed without taking the maturity level of the target community into due consideration [21]. In fact, a mature community increasingly develops means and tools to share and use the knowledge that characterizes it: these means and tools are peculiar to this community and can be effectively maintained by its members only.

C&T'09, June 25–27, 2009, University Park, Pennsylvania, USA.
ACM 978-1-60558-601-4/09/06.

In this light, our analysis focused on the mainly underspecified and tacit knowledge that is practically involved in either the proper or the conventional use [16] of functionalities like filling in, querying, retrieving and transforming the data of a traditional IS. Knowledge on proper (and "intended") use is "native" of the IS, in that it is tied to the requirements on data/process consistency and quality that is externalized at design and development time. Conversely, knowledge on conventional (and "actual") use is native of a specific community of users and characterizes the number of practices and ad-hoc means that users (re-)invent and put to work to better perform their collaborative activities [8, 4]: any support of this dissipative process [30] requires the understanding of the nature of these practices, and proper means to avoid proposing useless solutions. Our research effort aims to contribute to a better understanding of how tacit knowledge is conveyed through explicit knowledge, and how technology can amplify, or at least support, this action. The paper is organized as follows: in the next section the requirements emerged from the studied settings are illustrated by some significant vignettes: they regard the need to preserve and support current practices in terms of context interpretation and knowledge "evocation". Then, a specific interpretation of the notion of KA is used to characterize the functionalities that respond to these requirements. We also propose an architectural pattern as an extension of standard MVC patterns that could be used in different domains that share the same requirements on knowledge evocation. The achievements and further developments of this approach are illustrated in the concluding section.

2. CLUES TOWARDS KNOWLEDGE

In the last years, we accomplished an empirical study in two wards of an important hospital in Northern Italy [12, 16, 15]: a Neonatal Intensive care Unit (NICU) and an Internal Medicine (IM) ward. These wards were significantly different not only with respect to the medical specialty, but also as regards their communities of professionals and the nature of the work being done therein. We focused on how clinicians use the heterogeneous web of artifacts that compound their patient-centered clinical records [6]. We analyzed how these records are used a) to create an accurate and complete picture of the illness trajectories of their patients, b) to be supported in making apt clinical decisions and coordinating each other, even in frantic and often interrupted situations. Our observations together with interviews of the involved stakeholders constituted the basis for a phase of participatory requirement elicitation and conceptual application design. In fact, the study was deeply influenced by the current concern among practitioners about the impacts of the foreseen introduction of a computerized Patient Record fully integrated with the Hospital Information System.

During this phase, it emerged that even small cues, which clinicians were used to attach to their paper-based artifacts (e.g., lines, underscores, marks and post-it notes), were able to convey a lot of additional information to their colleagues. These additional pieces of information were informal and heterogeneous in nature, distributed across the whole web of artifacts, and they regarded either relevant aspects of the clinical context and the medical knowledge involved in the decisions reported in the records or, more generally, what of relevant was going on in the hospital ward.

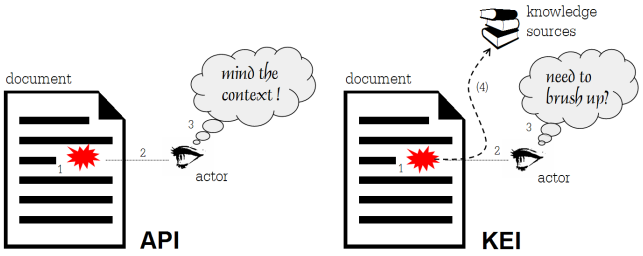

Figure 1: The conceptual difference between API and KEI.

A systematic analysis of the role of these additional annotations in the context of the cooperative work activities allowed us to identify 13 types of *Awareness Promoting Information* (API) [16]. These API types regard either a) coordinative aspects related to record keeping and related cooperative work (namely, Articulation API, Responsibility API, Appropriateness API and Schedule API); b) cognitive aspects involved in the interpretation of recorded data and work context (namely, Inquiry API, Change, Provisionality, Revision, Criticality and Inconsistency API); or c) intertwined dimensions of record keeping, work and community life that are partly cognitive and partly bound to cooperative conventions and mutual expectations of the community of professionals (i.e., Deviation API, Quality and Safety API).

Initially, we proposed the concept of API to address two intertwined research questions: on the one hand, to address the question "what should actors be made aware of, while they skim, consult and inscribe their documental artifacts?"; on the other hand, to translate general requirements on context awareness into specific computational requirements for the design of an innovative electronic clinical record able to raise and promote collaboration awareness [29]. Yet, while we were undertaking our study on API provision, we saw that even small additional indications conveyed through the analyzed artifacts had an important role in suggesting, promoting and even consolidating specific lines of action. We observed that these small indications could exert this pragmatic role only in virtue of a *specific knowledge*, which the visual and textual clues were able to *evoke* in the minds of practitioners. We therefore saw that some of the API types detected in the medical domain could be also considered KEIs, i.e., *Knowledge Evoking Information*. The difference between API and KEI is subtle but significant with respect to ICT design. In fact, although KEIs and APIs might share the same graphical/textual indication, they differ in both *affordance* [31] (i.e., what functionality related to the indication the system "affords") and associated (conventional) meaning (see Figure 1): therefore, an API is any indication or data that is "additional" with respect to regular (and official) content and that can make (or keep) users aware of some relevant condition or event; on the other hand, a KEI is any additional data that either allows users get a direct access to useful[1] knowledge resources or helps them remind of a more or less externalized notion.

[1]Here, usefulness can vary according to the textual context referred by either the structure or content of documents, as well as to the work context (where known) and local peculiarities (e.g., user profiles, access rights).

Observations of paper-based forms and practices at the NICU provide examples of these two cases. At the NICU, the head doctor fostered the habit to attach the sketch of a flowchart to the medical records of newborns whose illness was diagnosed as pneumonia. The flowchart represented a Pneumonia Pathway, i.e., a process schema that suggests how to manage a pneumonia regular case according to various conditions that can unfold during treatment. He asked his collaborators to add an explicit reference to the pathway in their usual work of documenting clinical decisions and prescriptions. In doing so, the medical record of pneumonia cases ended soon by being enriched with paper-based KEI, which would refer readers (at any time in future) to the procedural knowledge represented by the pathway flowchart attached at the front cover.

On the other hand, while a specific committee was working on the redesign of the hospital templates for the official medical record, neonatologists were asked to use the standard History Form template used in all the other hospital departments to report history of current complaint and past medical history. Yet, in order to duely frame the clinical history of a child, pediatricians needed to report more things than what was encompassed in the general-purpose History Form (which were initially designed for adult patients). So pediatricians decided to use the wide and unspecified text area at the bottom of the form to report their own notes. To this aim, they used to annotate an acronym beside the text area that would work as medical mnemonics and help them ask the child (or her parents) the additional questions not reported in the form: BIFIDA, i.e., Birth details and problems, Immunisations, Feeding Infection (exposure to), Development (normality of), Allergies. In so doing, pediatricians ended by enriching their history form with a KEI that would refer themselves to a conventional (and not more specified than that) way to remember due questions. These two examples can obviously suggest ICT designers to add simple expedients to the graphical interface of electronic forms to evoke either explicit (cf. the pneumonia pathway) or implicit (cf. the BIFIDA acronym) forms of knowledge that would be meaningful and agreed in a community of practitioners.

3. TAKING API AND KEI SERIOUSLY

In this section, we will survey the main types of API and sketch some simple vignette of knowledge use from our observational studies in order to make clear how API and KEI can exert their function in digitized media. In what follows, we will refer to either API or KEI according to the context, or more generally to AP/KEI meaning the general construct irrespectively of the content.

3.1 AP/KEI for Context Interpretation

Clinical work is a paradigmatic example of a massive usage of reading and writing inscriptions on a shared support, especially in the case where this support is based on paper [8]. For this reason, clinical practices are based on a rich set of annotations or conventional implications that all contribute towards providing a context to the content of the inscriptions itself [9]. Clinicians add annotations to their formal and due inscriptions to enrich the interpretative context of their colleagues or even as memorandum for themselves. The simplest case we observed regarded particular inscriptions that clinicians used to add beside either a passage or field of the record in order to indicate the *opportunity* of further or more in-depth inquiry; these inscriptions were usually small arrows with a label or a short note like "see ...". The closest example of this affordance that practitioners recognized in electronic documents was that of entries and sentences associated to hyperlinks: the different look of the hyperlinked resource can be seen as an example of what we called **Inquiry** API; the corresponding Inquiry KEI could be any indication (even a small hint) at the documental resource that is "pointed" or referred by the entry. The provision of this AP/KEI regards two main requirements: support double checking of the consistency of mutually correlated data (whereas the system is not capable of detecting such inconsistencies automatically); and support knowledge retrieval and retention, e.g., by binding fields, terms or entries with a corresponding legend, medical glossary or clinical guideline or by relating the current case to a previous clinical case recognized by clinicians as similar and comparable. Inquiry AP/KEI is annotated directly by clinicians when they record an item that is related to another one, which is possibly stored in a different document, to make this relationship more explicit. In doing so, clinicians can imply a linkage between entries of the patient record that can be as precise or ambiguous as needed [13]. Irrespectively of who or what creates Inquiry AP/KEI, this is conceived to make practitioners aware of additional data that could help them interpret a specific inscription, in order to leverage what has been called *redundancy by supplementary data* [12]; this positive redundancy is created in documents to facilitate practitioners' access to educational materials [38], improve their knowledge retention [28] and the learning of the specific notions and conventions that characterize a community of practice [46, 16].

The identification of the need of Inquiry API led practitioners to envisage more sophisticated functionalities aimed at promoting their awareness: namely, Revision, Inconsistency and Criticality AP/KEI. In regard to the former two, practitioners emphasized the fact that record's content, which out of its context could be seen as erroneous or jeopardizing patients' safety, does not always necessarily require an amendment or correction. In fact, they asserted the difference between an API that must be aimed at prompting the verification of its data sources (**Inconsistency API**) and an API that conveys a stronger suggestion, to check and correct record's data (**Revision API**). In fact, we observed that actors may have a very good reason to deal with a partial inconsistent state of the world, or even to supersede the conventions by which a situation is fallaciously considered inconsistent, e.g., as when a pregnant woman must undertake a C.T. for life-threatening complications or as when a vegetarian patient suffering from serious anemia is resistant to dietary supplements. The practitioners we interviewed expressed the need for being prompted of the presence of possible mistakes and inconsistencies in records and the need to be reminded of what resources state consistency constraints, accuracy requirements and improvement techniques (respectively, Revision and Inconsistency AP/KEI) because they told us that any mechanisms of automatic spell-checking or even smarter correction in official record-keeping would be unsuitable and even potentially harmful: they told us that they would prefer to have to cope with unsolicited warnings and alerts, which the computerized system could raise according to specific data constraints (e.g., numbers used in fields used to record names, phone numbers recorded with-

out the area code), rather than having to rely on automatic checking scripts. The reason for this is twofold: On the one hand, these spell-checkers could change clinical reports without even notifying it to the practitioners (they cited the automatic correction of MS Word as an example), or worse yet, give users the false confidence that a (syntactically) accurate report does not contain (semantic) mistakes. On the other hand, practitioners agreed that being actively reminded of the presence of a *possibly* inconsistent or just wrong information could stimulate a reflection about the possible reasons that led to the mistake and provide useful insight on how safety barriers are bypassed even unintentionally. In particular, in regards to Inconsistency AP/KEI, we discussed both the trivial cases that can be related to possible mistakes according to data constraints and types (e.g., body temperatures can not be higher than 44 degrees; dates for next examinations can not be scheduled in the past), as well as more local and articulated cases that can be related to the conventions according to which clinicians consider data as mutually inconsistent. In these latter cases, inconsistency is usually verified with respect to either other data recorded previously in the record, or to knowledge sources deemed as reliable: e.g., drug administrations can be considered inconsistent with respect to some particular disease, or to the allergies reported by the patient.

In the same vein, **Criticality** API can be related to situations that the community of practitioners deems as critical, also on a conventional and local basis. For instance, in newborn intensive care, practitioners wanted to be alerted of all the critical situations in which a colleague of theirs would record an *APGAR score* lower than 4 after five minutes since delivery. Similar conventions regarding situations that are deemed critical and worthy of attention within a specific community of clinicians could be based on specific thresholds of body temperature with respect to specific diseases, or on the basis of age, weight and skin surface in newborns. Criticality AP/KEI, and the corresponding alerts and pointers, could be raised on the basis of either local conventions or more global and widespread guidelines, constituting a reliable source of knowledge for the community. In particular, clinicians saw in Criticality AP/KEI a useful aid to call for a closer supervision of the critical newborns by the nurses on duty, to remind them of highly significant conditions in the scheduled handing-over conferences and to give them direct access to the relevant excerpts from the documentation regarding the critical conditions at hand (e.g., the table of the APGAR scores). In regard to more complex conditions of criticality, a doctor expressed a more complex convention: for operated inpatients low blood pressure is normal *unless and until* also signs of an anaemia show up, when instead low pressure could be an indication of internal hemorrhage. Similar conventions can be applied to all those cases where data become significant only *after* their insertion into the record, i.e., outside the context in which they were originally recorded. In those cases, clinicians expressed the requirement that an alert should be raised as soon as vital signs become serious under some other condition.

3.2 AP/KEI for Activity Articulation

The above mentioned case depicts how ward-specific conventions can externalize both a community need (i.e., to be alerted only in particular situations) and a knowledge notion (i.e., a critical condition occurs when pressure drops below a certain threshold) that a local community has either defined by itself or made its own from the literature or wider communities in the same domain. As we reported in [11], within a specific hospital ward, clinicians articulate care activities in virtue of a complex mix of either tacit conventions or explicit relations between tasks, which are expressed in terms of the so called Clinical Pathways. These process maps can be seen as the instantiation of the local combination and harmonization of heterogeneous and general guidelines to provide a visual representation of the course of action that clinicians agreed to comply with, when managing a specific disease. Clinicians maintain and refer to this kind of knowledge for a selection of diseases that the community considers as worthy of the related effort of modeling: typically, these diseases are either frequently recurring situations that have to be managed in a uniform way or highly critical diseases for which any mismatch in coordination can imply serious consequences in safety and effectiveness [15]. We described an example of convention that leverages clinical pathways, and therefore an externalized model of work flow, in [10], where we illustrated the case of the nurses of the emergency department that are pre-alerted in order to schedule with proper advance the transfer of a patient to the Stroke Unit as soon as the triage doctors have requested a chest-ray exam for that patient within the pathway of suspect stroke. We called the visual and textual indications that can warn nurses in cases like that as either **Schedule** AP/KEI or **Articulation** AP/KEI. Schedule API is used to alert nurses of what activities have been performed or are to be performed (e.g., an exam order, a test result); schedule KEI can either refer to assignment records and agreed process models (declarative knowledge), as well as hint at how to perform tasks on time (past experiences, best practices) or to articulate their execution efficiently (procedural knowledge [36]). On the other hand, Articulation API is used to indicate what activities follow the scheduled activities in an agreed and common process map; the corresponding KEI aims to either evoke coordinative patterns (declarative knowledge) or hint at how prevent, manage and solve coordinative bottlenecks and conflicts. An interesting convention that we collected in the field study, and then traced back to a Schedule AP/KEI requirement, states the following: if a certain passage has been recorded by a nurse long after the scheduled end of her work-shift, this could mean two different yet correlated things: on the one hand, that the recorded item could refer to the handling of a serious emergency (whose management had priority over record-keeping); on the other hand, that the recorded item should be considered with some caution. Also in this case, we see how a convention lies across different AP/KEI, namely Schedule, Criticality and Revision, and can inform them all, notwithstanding that most of the possible implications must be necessarily left tacit and underspecified. In fact, if a system could make the situation depicted above just more evident, the practitioners of the following shift would be invited to apply the appropriate interpretation and align their actions accordingly. In this case, the involved knowledge regards multiple aspects, like the overall illness trajectory of the patient, her current conditions, who wrote the inscriptions and in which context. On the other hand, in regards to Articulation AP/KEI, we collected the strong interest especially by nurses. They proposed that this AP/KEI could be conveyed in all those cases where actors are involved in activities that can block oth-

ers's activities. In this case, these actors could be supported in understanding the needs of the colleagues involved in blocked activities. This would contribute to avoiding underutilization of resources and would limit the occurrence of situations where practitioners are kept idle and their time is wasted. We also observed the occurrence of these conflicts and schedule clashes in [12] and found these occasions quite frequent and time consuming, besides being a source for resentment and frustration that had an impact on the quality of the collaboration both between different professions (typically, nurses and doctors) and between different facilities (typically, a hospital ward and a diagnostic service).

We speak of **Responsibility** API and KEI, when artifacts provide either a direct or indirect indication of who should accomplish a task, in all those cases an explicit organizational hierarchy is known, as well as job descriptions and assignment records[2]. Clinicians advocated the provision of Responsibility AP/KEI for those situations when practitioners consult the history or log of updates of a certain section to have a quick glance of who did what. They can need this information in order to either ask for a clarification on a clinical case or to even assess the reliability of the entry [33, 18], especially in case of disagreement between colleagues. For this reason, Responsibility AP/KEI can be assimilated to something between social and task-oriented awareness [35].

Whereas Articulation API regards indications that make practitioners aware of what activities *can be* accomplished after some other activity (according to some process model), what we called **Appropriateness** API regards indications of what activities *should be* accomplished according to the context [11]. The corresponding KEI would be aimed at evoking how to perform tasks compliantly to either official or informal models of actions (e.g., Standard Operating Procedures and conventions, respectively). The context in which an activity can be considered appropriate or not pertains to either the current content of the record (i.e., data) or other contextual condition, such as absolute time (e.g., it's noon) or time relative to a past event (e.g., it's the second day of the chemotherapeutical cycle). Clinicians agreed with us that conveying this AP/KEI effectively would be extremely difficult: first of all it would require a substantial modeling effort that can be justified only in the case of specific critical situations; moreover, Appropriateness AP/KEI would require a tight integration between hospital information systems and electronic clinical records. In this light and similarly as the previous APIs and KEIs, Appropriateness AP/KEI was first of all perceived as a meaningful category for the phases of requirement elicitation and domain analysis towards the deployment of a supportive knowledge artifact.

3.3 AP/KEI for Compliance and Safety

Deviation AP/KEI is conveyed to make actors aware of what data either regard or represent a variation with respect to an *expected* outcome, trend or indication by the physician. Therefore, Deviation API regards changes with respect to some "expected" therapeutic interventions or physiological value range, which have not yet been associated to any Criticality API. In this case, "expectations" are based on specific diagnostic/therapeutic patterns defined in terms of local and evidence-based conventions, referred or just evoked by means of corresponding KEIs. We have observed clinicians generate this API on their own initiative: for instance, to justify a prescription that involves dosages that are significantly different from those normally computed by formulas taking into account clinical data (as body surface in chemotherapy, or age of newborns in neonatology). Clinicians create Deviation API also to point out to their colleagues that they deem a certain clinical trend (e.g., in the vital signs or blood levels) as unusual, unexpected or just different from the hoped reaction to the treatment.

When the records' content "deviates" with respect to precise quality requirements and targets, we speak of **Quality** API and KEI. This kind of AP/KEI is conveyed to make actors aware (and remind) of the current level of data quality of either specific sections of the Clinical Record or of the whole of it. The need for this AP/KEI emerged after observing the initiative taken by the pharmacy and radiology supervisors to sporadically return order forms to the ward in case of inadeguate information: they attached a rough indication of the intrinsic quality of the request to the forms to raise practitioners' awareness of the importance of filling forms completely and accurately, or else their departments could not process requests efficiently and safely. For instance, requests lacking the indication of the patient's gender, weight, as well as reporting inaccurate identification details and incomplete reasons why a certain exam has been requested, usually lead the referred unit to contact the referring unit or the hospital admission by phone in order to complete this information with consequent delays, rework and responsibility bouncing back. The basic requirement behind this kind of AP/KEI is to help clinicians become aware of the requirements pertaining to the activities that follow the merely clinical part of the hospital process, leaving them free to improve the quality of the document either on the spot or later in the work shift. As practitioners told us, data quality of clinical records is a topic of increasing importance. In fact, data produced by clinicians during their care activities play a central role in risk management and in a number of heterogeneous activities that use these data to reach their organizational objectives [43]: e.g., epidemiological and pharmacological research centers, hospital and regional administrations, national policy makers and international committees. It is known that, since the level of accuracy and completeness of medical records is approximately 95% [2], relatively small improvements in their data quality can yield significative improvements in process quality and safety.

Clinical safety is addressed by the last kind of AP/KEI we survey in this section, **Safety** AP/KEI, which is conveyed to make actors aware that either the activity they are performing or that they are about to perform can have a strong impact on the patient's safety and is correlated with a significative occurrence of adverse events (i.e., events implying a damage for the patient, e.g., adverse drug reactions [45]). Recent cases of adverse events that aroused a great deal of attention (e.g., the Jowett case [45]) led some of the clinicians we interviewed to create specific conventions by which certain situations were deemed safety-critical, like drug administrations performed outside their regular schedule; administrations at close time intervals of drugs that differ in their administration way or with strong and known interactions; or, more generally, any unexpected situation where administration time, drug type or anything else has changed in the end without planning. The need of this AP/KEI was

[2]In Italian: mansionario dei profili professionali

associated to a number of cases that represent just a small amount of all conditions by which a clinical intervention can go wrong. For instance, Safety AP/KEI can be associated i) to the prescription and administration of specific treatments (e.g., cytotoxic and chemotherapeutic drugs): in this case the corresponding KEI would refer to the drug detail sheet; ii) to the use of any abbreviation and acronym (which are frequently misunderstood and associated with errors [19]); the corresponding KEI would refer to a hospital or ward-wide glossary; iii) to the prescription of dangerously invasive procedures (e.g., angiography, rachiocentesis) (KEI would refer to detailed standard operating procedures to deal those critical cases); iv) to the prescription of drugs with names similar to those of dangerous drugs (called look-alike and sound-alike drugs [39]) or close to dangerous drugs in nation or local drug list.

4. THE ROLE OF THE COMMUNITY

In our empirical study, the phase of collecting relevant situations for the community's life and of defining the related requirements of awareness promotion and knowledge evocation have given rise to some considerations on the ongoing elicitation process and on the relation between its outcome and the more general issues regarding the hospital milieu. As anticipated, the two considered wards showed some interesting differences. At the NICU, the high level of participation in our research and the relative easiness by which the condition-KEI-knowledge associations were externalized made it evident that this community of practitioners was highly engaged in creating and maintaining a shared, although not always explicit, pragmatic knowledge. When questioned about this point, the NICU practitioners acknowledged the fundamental contribution of the head doctor and his approach to the department management, both from the organizational and human point of view. Despite the tight nature of NICU work, he scheduled several opportunities in which nurses and doctors were invited to reflect and brief each other on what was going on in the unit, both during rounds and in dedicated meetings at shift handovers. The elicitation process undergoing during the study was therefore perceived as an opportunity to externalize what was already "there". In addition to that, they welcomed the opportunity to get actively involved in the twofold process of envisioning possible technological solutions that could respond the requirements and of thinking how these solutions could be integrated in their practices.

Conversely, at the Internal Medicine ward, the situation was not that positive since people were more skeptical towards any effort additional to their daily schedule and especially towards any initiative that could be associated to the possible deployment of an electronic patient record. The contribution of a couple of competent key users to the identification of interesting cases was positive but their attitude towards leveraging their work conventions and practices in order to conceive a technological support was more circumspect. This difference confirms the fundamental role of the head doctor: in fact, at the NICU he played both the role of community member – let's say "au pair" with the other members – and the role of "strategist" [46], i.e., who protects and promotes the creation of local knowledge within his community in relation to the other communities and the overall organization. In either cases, however, practitioners recognized that the practices they identified were generated as a reaction to the mismatch between their ever-changing local needs of knowledge support and the too rigid structures – of both data and business process – that were progressively incorporated in their Hospital Information System and computer-based applications. They appreciated these technologies for most of their basic archival and retrieval capabilities but found them progressively circumscribing their interventions and limiting their ambits of discretion. This was true also in the case of the NICU, where however some of the artifacts that had been incorporated in the hospital IS had been designed with the active participation of some NICU's representatives [12]: in this case, clinicians recognized the need of a much more flexible layer that could address their peculiar information and knowledge-oriented needs. They advocated a layer that, on the one hand, could be fully integrated with the hospital IS ("same data are to be reported only once"), and, on the other hand, could be locally managed and maintained to support their cooperative practices and decision making. These considerations led us to conceiving a technological support that aims to fulfill the two above requirements as described in the next section.

5. THE ROLE OF AP/KEI IN DESIGN

Irrespectively of the specific content of each type of AP/KEI, all of them share some basic features that emerged during the interactions with practitioners. First, each AP/KEI specifies a set of information requirements about context interpretation and knowledge retrieval. The specification language we have defined and used [16, 14], LWOAD, allows designers to express requirements in an associative and modular way, i.e., in terms of rule-based mechanisms endowed with data patterns. In their more general format, each requirement is then expressed in terms of an *association* between a relevant situation and specific functionalities that a computer system must exhibit to support actors in that situation. In a mechanism, situations are expressed as sets of declarative statements (facts) that are all true in that situation and that are bound together on a conventional basis in the antecedent of the mechanism. In each mechanism, the set of contextual conditions is associated with GUI-related functionalities – declared as primitives in the consequent part of the mechanism – in order to display users indications that could either a) make them aware of those contextual conditions that are relevant for the community's life (provision of API, as illustrated in [16]); or b) recall them what resources they have the opportunity to consider to interpret the situation correctly, with respect to the community's conventions, and manage the situation accordingly (provision of KEI).

Second, it is noteworthy that each triple association context-KEI-knowledge and context-API-notification is reified by a single mechanism. In so doing, each specific situation is described as a standing alone "module" that is "exhaustively", but yet "minimally", described without reference to any other association or chain of reasoning. What exhaustivity means for the practitioners involved is fully related to pragmatic and local justifications that go beyond any consistency or completeness criteria that someone outside the community could reasonably propose or even impose. We took "their minimality" as the leading criterion, basing on the belief that well established common practices suffice community's members to fill in the missing information and recognize the target situation. The same holds for the cues

that the technological system should employ to evoke the appropriate knowledge and make them aware of either related resources to consult or even more tacit knowledge to retain.

Third, the role of conventions. Practitioners expressed the need to be aware of some contextual condition only because they could respond to this condition relying on some specific convention, policy, rule or corpus of notions: in fact, this kind of knowledge that they had previously (and in some way) internalized and shared could make that condition meaningful to hem. As anticipated, the association between context and knowledge can not be based on a formal and fully specified model of the application domain [42, 27]. In fact, practitioners considered this effort not only open-ended but simply useless since what they need is a mix of tacit and explicit knowledge whose nature cannot help being in any case underspecified, incomplete and sometimes even inconsistent. While the two former attributes respond to the criterion of minimality, the third one carries a value in itself; in fact, inconsistency, when it is not related to a trivial mistake, is the evidence of either divergent opinions, lack of reconciliation between different perspectives or critical situations that deserve specific attention.

These three features let the practitioners consider each association as a sort of single "reactive" mechanism. In fact, each association can be seen as condition-action construct that, in its antecedent (if-part), defines relevant conditions requiring attention and competence and, in its consequent (then-part), identify *what* AP/KEI to convey to users (but not *how* to convey it). These reactive mechanisms contain the pieces of behavior that make the hospital web of artifacts not only a set of Input/Output interfaces towards the underlying IS; rather, our point is that, these behaviors have the potential to transform each artifact of this web into an Active Knowledge Artifact, aimed at supporting both the articulation and knowledge work that are required to cooperate within (and across) communities of practitioners.

6. ACTIVE KNOWLEDGE ARTIFACTS

The vignettes discussed in Section 3 confirm that the documental artifacts that practitioners inscribe and use in their everyday work are actually KAs since they either refer to or incorporate different kinds of tacit and explicit knowledge. Our point is that these KAs exert their function by means of KEIs, i.e., of annotations that, in virtue of their conventional value, have the power to evoke knowledge by making it – or the sources that represent it explicitly (e.g., guidelines, diagrams, procedures, memorandums) in a given community –'present-at-hand' [22] for the interpretation of the current context and future action.

In [4], we proposed a notion of KAs that emphasizes how they mediate the process of *knowledge externalization* [30] and how they are *collectively defined* as the result of a progressive stratification of experiences, local practices of use and lessons learned to solve problems. This definition of Knowledge Artifact enhances alternative formulations [26, 40] that take the notion of KA as a mere intuitive and decontextualized juxtaposition of the concepts of knowledge and artifact and that therefore disregard its role of community collective memory [15].

Actually, cooperative settings provide a wealth of significative examples of the *social* and *participatory* nature of the core knowledge therein involved, as well as of the *dynamic* and *cumulative* nature of the knowledge artifacts that reify it: almost any manual, internal report, bulletin and circular that has been collaboratively edited and that can be re-edited, amended and referred to by their "consumers" can be considered a KA, as long as it "incorporates" some core competencies and "best practices" in which members of a community *recognize themselves*.

However, the scenario of the study that we report here is different, although being quite common in organizational settings: in fact, the artifacts are given to practitioners – typically in terms of forms that constitute the input sources of the organizational IS – and with little or no possibility of further customization and tailoring to local needs. In this case, the community members have to develop alternative ways to define the memory characterizing their community: they have to take into account the existing information, its predefined structure, and to build an additional layer "on top of it", which is fully under their control and is manageable in an effective and efficient way. The study confirmed that this stratification is easier when artifacts are in a paper-based format thanks to the well known flexibility of this medium [41] and emphasized the current concern among practitioners about how to preserve this possibility in the foreseen introduction of a computerized Patient Record.

The interaction we had with the practitioners leads us to propose to support KAs, and their community-gluing capabilities, by endowing them with computational capabilities, i.e., we propose to build KAs as specific types of active artifacts [20]. Active artifacts are data structures capable of assuming an active (i.e., either reactive or proactive) role in mediating information exchange and coordination among cooperative actors. To this aim, active artifacts incorporate aspects of the coordinative protocols and conventions that refer to the portion of context that the artifact can perceive (e.g., time, task status), as well as to the content of the artifact that users progressively inscribe on it. The most notable research on active artifacts in documental and cooperative settings has been the Placeless Document Project [21] developed at Parc. In our research, we extend this idea by considering artifacts that exhibit behaviors that i) facilitate the coordination of actions and the sharing of the memory among the community of their users, ii) promote awareness of the collaborative interactions occurring in the field of work [29]; and iii) support (event tacit) knowledge retention and exploitation.

The former two functionalities have been discussed in other works [20, 4, 16]: in this paper we focus on the ability to support knowledge work by embedding "user-generated" and modular active code in documental artifacts; this code makes the reactive mechanisms mentioned in Section 5 computable and hence augments the artifacts' content according to the context to make users aware of any significant condition.

This kind of augmentation differentiates our research from other proposals where content is either annotated on the basis of formal models (e.g., [44]) or filtered adaptively according to the context and the channel of its provision (e.g. [34]), and where various kinds of reasoning (e.g., case based, ontology based, logic based) are used to achieve either context-adaptiveness or context-awareness (e.g. [37]).

7. AN ARCHITECTURE FOR COMMUNITIES

In this section, we describe the notion of Active Knowledge Artifact in the context of an architectural pattern extending a common design pattern, the Model-View-Controller (MVC). We also outline how KAs can be used to make an MVC architecture "content-aware" and oriented to community's conventions and rules. As mentioned in Section 6, Active KAs are digital artifacts (parts of an electronic information system) whose aim is to remind users of the knowledge these should retain in order to interpret the artifacts' content (and hence, the overall context data refer to). Active KAs result from associating the interpretable specification of mechanisms, which have been designed to evoke knowledge, to data templates, which have been designed for the View layer (see Figure 2).

In our architecture, Active KAs are processed by a layer that we conceive as an augmentation of the typical Controller of legacy information systems and that in Figure 2 we called "community-oriented layer". This layer is conceptually on-top-of the "business-oriented layer" that constitutes the regular Controller-Model stack of an Information System (e.g. the Hospital Information System coupled with a traditional Electronic Patient Record application). While these give access to data, store them and permit their modification, the community-oriented layer is designed to enrich these same data with API and KEI according to a set of mechanisms, acting as sort of "rewriting rules". In fact, these mechanisms are active rule-based constructs that generate either API or KEI on the basis of conditions expressed over the data contained in the "passive" portion of the artifacts (bottom side of Figure 2 and left side of Figure 3) and possibly over other data coming from either the View layer (e.g., user-driven events) or other third-party context managers (not depicted in Figure 2 for simplicity's sake). To this aim, a rule interpreter (see middle section of Figure 2) matches the mechanisms' if-part with the data that the Controller fetches from the underlying Model to execute their then-part (see Section 5). These then-parts contain instructions to build specific metadata to be associated with application data; these metadata are RDF statements that associate data fields, values and KEI types (see also the KEI Metadata File in Figure 3). These metadata are then processed by a Markup Tagger (see the top side of Figure 2 and right side of Figure 3). This component automatically annotates the Data pages with markup tags that are associated to specific style classes so that specific KEI types can be rendered in terms of specific affordances, icons and text formatting as defined in corresponding style sheets (styles in Figure 2). Finally, a Layout Engine (e.g., an Internet browser), in the View layer, takes the annotated pages as input, interprets them and display the final document to the user.

To implement the prototype that we proposed to the users to get a preliminary validation of the architecture, we expressed the AP/KEI mechanisms by means of the LWOAD language [16, 14] and used the corresponding interpreter (LWOAD Interpreter in Figure 3) to execute the mechanisms and generate the appropriate metadata. With reference to Figure 3, the LWOAD interpreter takes the KA as input and executes the mechanisms contained in the KA's header by matching their antecedents with the data retrieved from the underlying DBs. As output of the mechanisms' consequents, the LWOAD interpreter produces a KEI Metadata file containing KEI tuples (the same holds for API tuples as well): these are XML triplets <KEI-type>,<ID Source>,<KEI content> that are associated univocally to the KA's data structures (i.e., templates) through the ID tag. The KA's body and the KEI Metadata are then inputs of the Markup Tagger. As said above, this produces a Data Page in HTML where raw data from the KA's body are enriched with metadata associated to Cascading Style Sheets and Javascript functions according to the type of KEI.

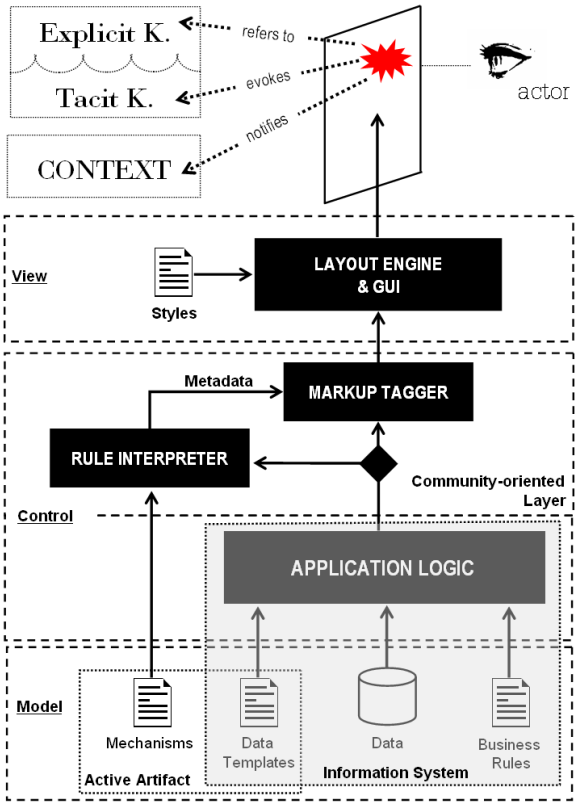

Figure 2: The architectural pattern of the KA-based solution.

In this architecture, users can play a twofold role: according to the End-User Development approach we undertook in [14], users can be involved either in defining the mechanisms that characterize each Active KA; in defining the configuration of the Markup Tagger, which specifies the mapping between KEI types and style sheets; or in defining the style sheets themselves. In doing so, users can take an active role in defining how specific types of KEI are associated to standard HTML tagging, i.e., to specific ways to render them. On the other hand, Active KAs are proactive in displaying additional clues regarding data but do not prompt users to take any action: they just notify users that a local convention, which associates a certain data in a certain template to a specific knowledge resource, has become significant; users are then left free to exploit this indication to either recall a whole bunch of (otherwise unspecified) notions or to consult the suggested knowledge sources for further inquiry.

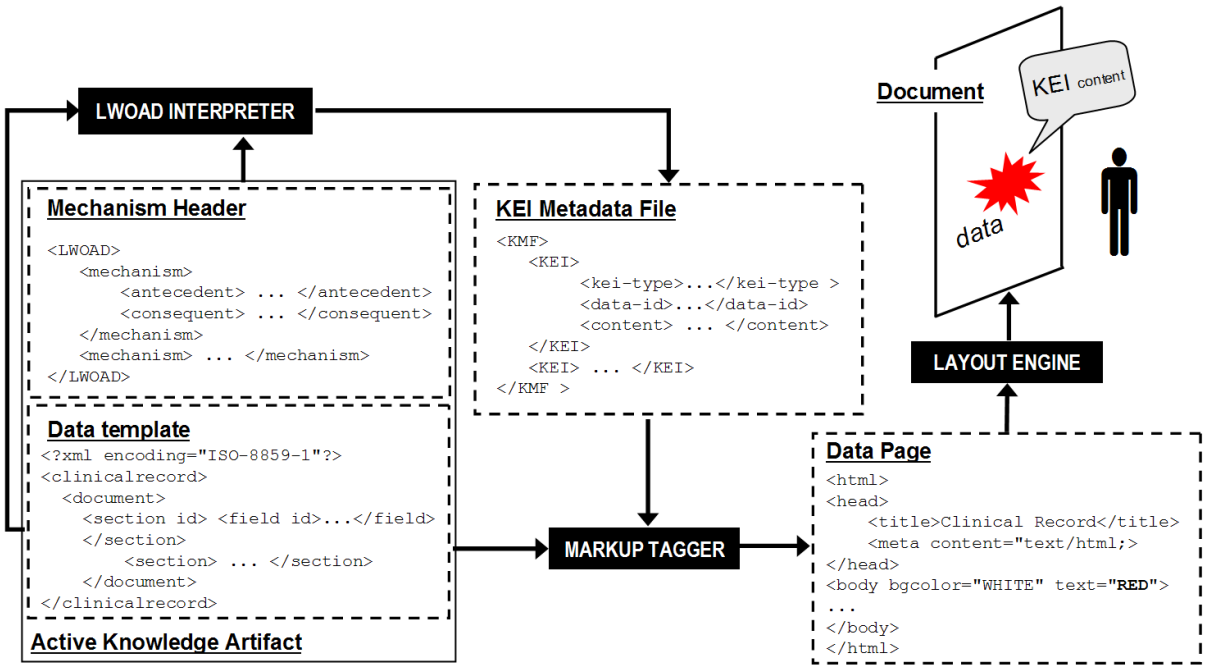

Figure 3: The main components of the Active Knowledge Artifact Management System.

8. CONCLUSIONS AND FUTURE WORK

The paper discusses the case of communities whose members are professionals working in an organization that provides them with the predefined information structures of an Information System. In this situation, there is an evident conflict between the need of flexibility required in everyday work and the rigidity of the predefined structures of the organizational IS. An empirical investigation in the domain of hospital work highlighted some practices that doctors and nurses put to work to overcome this conflict: these practices are mainly based on particular annotations that have the goal to either evoke conventional tacit knowledge or to point relevant sources of information. The paper presented an approach that leverages these practices in order to reduce the above mentioned conflict when organizational artifacts are digitized, in our case, the Electronic Clinical Record. Our point is to augment the documental artifacts of an organizational IS in order to make them play the role of Active Knowledge Artifacts, i.e., artifacts that are able to "react" to contextual situations and convey additional information (namely AP/KEI) able to evokes the often tacit and underspecified knowledge necessary to effectively act and coordinate action through the artifacts. The most important aspect of our empirical study is that APIs and KEIs are effective in recalling significant contextual conditions and knowledge resources, respectively, only if they are tightly grounded on the community's conventions. In [16, 14], we illustrated how designers can refer to conventions in order to build LWOAD mechanisms without worrying of having to cope with world models that are underspecified, incomplete or even inconsistent. In this paper, instead, we propose to leverage conventions also to have the indications conveyed as output of the LWOAD mechanisms meaningful in the local context of a community of practitioners. In other words, what designers can easily externalize from the community's conventions can be used to design tailored patterns able to describe significant conditions of the context; conversely, what of conventions is more difficult to express in declarative and explicit ways can be evoked by means of graphical cues, indications and messages, which become pragmatically useful in virtue of these tacit conventions.

This approach has a deep impact on the kind of technological solution that can be proposed: therefore, the originality of the solution that we propose does not lie in the individual components of the architecture, which we can easily recognize in other applications (namely, a rule-based interpreter, an HTML renderer, a DBMS); rather the originality of our approach lies in the goal towards which these individual components are combined together, on the basis of a detailed empirical study. This corroborates the view of other user-centered investigations claiming that, when designers have to realize a technology supporting Knowledge Management in real settings, the main problem is not technological, but rather it resides in the phase of requirement collection to identify usable and useful functionalities. Technological problems concern how to express the computational mechanisms that provide these functionalities and how to develop the overall architecture where IS and KM functionalities can interoperate to support the members of a community effectively. In order to address these points, the paper proposed an operative concept (Active KA), design-oriented requirements derived from the empirical investigation and an architectural pattern that can be applied to different application domains. In this effort, our main points are that i) KAs can be endowed with the capability of evoking knowledge in users by conveying AP/KEI, i.e. graphical and textual cues; ii) users must be then left free to recognize the relevance of what is evoked and exploit it if needed; iii) AP/KEI can be triggered by simple reactive mechanisms that do not presuppose complete representations of the domain knowl-

edge, rather they are derived from the investigated practices. Our study has not yet focused on a systematic, incremental and participatory analysis on how to validate the graphical rendering of the proposed AP/KEIs; therefore, future work will be devoted to this aspect and to the possibility of incrementally leaving mechanism definition and AP/KEI visualization under user control.

Acknowledgements

The work presented in this paper has been partially supported by the Italian fund F.A.R. 2008. The authors would like to thank the management and the Internal Medicine and Neonatal Intensive Care Unit personnel of the "Alessandro Manzoni" Hospital in Lecco (Italy) for their kind collaboration. In particular, we would like to acknowledge the invaluable help and courtesy of Dr Roberto Bellù and Mrs Daniela Colombo and Rossana Pezzotta.

9. REFERENCES

[1] G. Alvarez and E. Coiera. Interruptive communication patterns in the intensive care unit ward round. *International Journal of Medical Informatics*, (74):791Ů796, 2005.

[2] D. Arts, N. D. Keizer, and G. Scheffer. Defining and improving data quality in medical registries. *Journal of American Medical Informatics Association*, 9(6):600–611, 2002.

[3] J. S. Ash, M. Berg, and E. Coiera. Some unintended consequences of information technology in health care: The nature of patient care information system-related errors. *Journal of the American Medical Informatics Association*, 11:104–112, 2004.

[4] S. Bandini, E. Colombo et al. The role of kmowledge artifact in innovation management. In *C&T2003*, 327–345, Amsterdam, The Netherlands, 2003.

[5] J. E. Bardram. Temporal coordination: On time and coordination of collaborative activities at a surgical department. *CSCW*, 9(2):157–187, 2000.

[6] J. E. Bardram and C. Bossen. A web of coordinative artifacts: collaborative work at a hospital ward. In *GROUP'05*, 168–176, New York, NY, USA, 2005.

[7] C. Batini and M. Scannapieco. *Data Quality: Concepts, Methods, and Techniques*. Springer, 2006.

[8] M. Berg. Accumulating and Coordinating: Occasions for Information Technologies in Medical Work. *CSCW*, 8(4):373–401, 1999.

[9] S. Bringay, C. Barry, and J. Charlet. Annotations: A functionality to support cooperation, coordination and awareness in the electronic medical record. In *COOP'06*, France, Provence, 2006.

[10] F. Cabitza, M. P. Locatelli, and C. Simone. Promoting process-based collaborative awareness to integrate care teams. In *ProHealth '08, Milan, Italy, September 2008*, 2008.

[11] F. Cabitza, M. Sarini, and C. Simone. Providing awareness through situated process maps: the hospital care case. In *GROUP'07*, 41–50, New York, NY, USA, November 2007.

[12] F. Cabitza, M. Sarini et al. "When Once Is Not Enough": The role of redundancy in a hospital ward setting. In *GROUP'05*, 158–167, Sanibel Island, Florida, USA, November 2005.

[13] F. Cabitza and C. Simone. Supporting practices of positive redundancy for seamless care. In *CBMS'08*, 470–475. IEEE Computer Society, June 2008.

[14] F. Cabitza and C. Simone. LWOAD: a specification language to enable the end-user develoment of coordinative functionalities. In *IS-EUD2009*, LNCS, 5435, 146–165. Springer 2009.

[15] F. Cabitza, C. Simone, and M. Sarini. Knowledge artifacts as bridges between theory and practice: the clinical pathway case. In Knowledge Management In Action, IFIP 270/2008:37-50, Springer, 2008

[16] F. Cabitza, C. Simone, and M. Sarini. Leveraging Coordinative Conventions to Promote Collaboration Awareness. The case of Clinical Records. *CSCW*, 18, 2009 *to appear*.

[17] P. Checkland and S. Holwell. *Information, Systems and Information Systems*. Wiley, 1997.

[18] K. Clarke, M. Hartswood et al. Trusting the record. *Methods of Information in Medicine*, 42:345–352, 2003.

[19] M. Cohen, editor. *Medication Errors*. American Pharmaceutical Association, Washington, USA, 1999.

[20] M. Divitini and C. Simone. Supporting different dimensions of adaptability in workflow modeling. *CSCW*, 9(3):365–397, 2000.

[21] P. Dourish, W. K. Edwards et al. Extending document management systems with user-specific active properties. *ACM Transactions on Information Systems*, 18(2):140–170, 2000.

[22] T. Winograd and F. Flores. Understanding Computers and Cognition: a new foundation for design Addison Wesley, 1986.

[23] P. Gorman, J. Ash et al.. Bundles in the Wild: Managing Information to Solve Problems and Maintain Situation Awareness. *Library Trends*, 49(2), 2000.

[24] G. Hardstone, M. Hartswood et al. Supporting informality: team working and integrated care records. In *CSCW '04*, 142–151, New York, NY, USA, 2004.

[25] G. Hardstone, M. Hartswood et al. Supporting informality: team working and integrated care records. In *CSCW '04*, 142–151, New York, NY, USA, 2004.

[26] C. W. Holsapple and K. D. Joshi. Organizational knowledge resources. *Decis. Support Syst.*, 31(1):39–54, 2001.

[27] H. D. Jorgensen. Interaction as a framework for flexible workflow modelling. In *GROUP '01*, 32–41, New York, NY, USA, 2001.

[28] A. Knight, S. Kravet, et al. The effect of computerized provider order entry on medical student clerkship experiences. *Journal of the American Medical Informatics Association*, 12(5):554–560, 2005.

[29] C. J. Lauwers and K. A. Lantz. Collaboration awareness in support of collaboration transparency: Requirements for the next generation of shared window systems. In *CHIŠ90*, 303–311, Seattle, Washington, USA, April 1990.

[30] I. Nonaka and H. Takeuchi. *The Knowledge Creating Company*. Oxford University Press, Oxford, UK, 1995.

[31] D. Normann. *The Psychology of Everyday Things*. Basic Books, New York, 1988.

[32] J. Parker and E. Coiera. Improving clinical communication: a view from psychology. *Journal of the American Medical Informatics Association*, 7(5):453–461, 2000.

[33] V. L. Patel, K. N. Cytryn et al. The collaborative health care team: The role of individual and group expertise. *Teaching and Learning in Medicine*, 12(3):117–132, 2000.

[34] B. Pernici, editor. *Mobile Information Systems: Infrastructure and Design for Adaptivity and Flexibility*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.

[35] W. Prinz. NESSIE: an awareness environment for cooperative settings. In *ECSCW'99*, 391–410, Norwell, MA, USA, 1999.

[36] L. Prusak. *Knowledge in Organizations*. Butterworth-Heinemann, 1997.

[37] A. Ranganathan and R. H. Campbell. An infrastructure for context-awareness based on first order logic. *Personal Ubiquitous Comput.*, 7(6):353–364, 2003.

[38] S. Rosenbloom, A. Geissbuhler, and W. Dupont. Effect of CPOE user interface design on user-initiated access to educational and patient information during clinical care. *Journal of the American Medical Informatics Association*, 12(4):458–473, 2005.

[39] L. Schulmeister. Look-alike, sound-alike oncology medications. *Clinical Journal of Oncology Nursing*, 10(1):35–41, February 2006.

[40] R. Seiner. Metadata as a knowledge management enabler. *TDAN.com & KIK Consulting Services The Data Administration Newsletter (TDAN.com).*, 15, January 2001.

[41] A. J. Sellen and R. H. R. Harper. *The Myth of the Paperless Office*. MIT Press, Cambridge MA, 2003.

[42] F. M. Shipman and C. C. Marshall. Formality considered harmful: Experiences, emerging themes, and directions on the use of formal representations in interactive systems. *CSCW*, 8:333–352, 1999.

[43] D. M. Strong, Y. W. Lee, and R. Y. Wang. Data quality in context. *Communications of ACM*, 40(5):103–110, 1997.

[44] V. S. Uren, P. Cimiano, et al. Semantic annotation for knowledge management: Requirements and a survey of the state of the art. *J. Web Sem.*, 4(1):14–28, 2006.

[45] C. Vincent. *Patient Safety*. Elsevier Limited, Oxford, UK, 2006.

[46] E. Wenger. *Communities of Practice: Learning, Meaning, and Identity*. Cambridge University Press, 1998.