Context Aware Customer Experience Management: A Development Framework Based on Ontologies and Computational Intelligence

Hafedh Mili, Imen Benzarti, Marie-Jean Meurs, Abdellatif Obaid, Javier Gonzalez-Huerta, Narjes Haj-Salem, Anis Boubaker

LATECE Laboratory, Université du Québec à Montréal
mili.hafedh@uqam.ca

Abstract. Customer experience management (CEM) denotes a set of practices, processes, and tools, that aim at personalizing customer's interactions with a company according to customer's needs and desires [29]. E-business specialists have long realized the potential of ubiquitous computing to develop context-aware CEM applications (CA-CEM), and have been imagining CA-CEM scenarios that exploit a rich combination of sensor data, customer profile data, and historical data about the customer's interactions with his environment. However, to realize this potential, e-commerce tool vendors need to figure out which software functionalities to incorporate into their products that their customers (e.g. retailers) could use/configure to build CA-CEM solutions. **We propose** to provide such functionalities in the form of **an application framework within which CA-CEM functionalities can be specified, designed, and implemented**. **Our framework relies on, 1) a cognitive modeling of the purchasing process**, identifying potential touchpoints between sellers and buyers, and relevant influence factors, **2) an ontology to represent relevant information about consumer categories, property *types*, products, and promotional material, 3) computational intelligence techniques to compute consumer- or category-specific property *values***, and **4) approximate reasoning algorithms to implement** some of **the CEM functionalities**. In this paper, we present the principles underlying our framework, and outline steps for using the framework for particular purchase scenarios. We conclude by discussing directions for future research.

Keywords: Customer Experience Management, Customer Behavior Modeling, Service Design, Customer Profile Ontology, Product Ontology, Recommendation Systems, Context Awareness, Computational Intelligence, Data Mining.

1 Introduction

Customer experience management (CEM) denotes a set of practices, processes, and tools, that aim at personalizing customer’s interactions with a company according to customer's needs and desires [29]. Ubiquitous computing is a computing paradigm where a computation of interest to a stakeholder is collaboratively performed by a variety of often specialized devices with limited capacities that interact "spontaneously" on behalf of the stakeholder [19, 27]. E-business specialists have long realized the potential of ubiquitous computing to develop context-aware CEM applications (CA-CEM). Many e-business visionaries fantasize about what our shopping experience could be like at our favorite food store, entertainment place, or clothing store if the "system" knew about our lifestyle, our age, our heartbeat, what we ate this week, the weather outside, the people we hang out with, in real life and social networks, the places we visit, on foot or by mouse, what we like, what we wrote in our blogs, etc. What would be possible if we shared the same identity across devices, media, and applications, and if all the "things" within our environment were connected [29], including our cell phone, the store we walked into, the steak slab we put in our shopping cart.

This raises a multitude of questions, both for the e-commerce software vendors who want to offer CA-CEM functionalities, and for customers of such applications (e.g. retailers), who wish to take advantage of such functionalities.

Examples of these questions include:

- *What CEM functionalities do we want to provide?* Different actions can be taken, depending on the phase of the consumer experience. For example, during pre-purchase, we can try to anticipate a customer need ("your car is due for an oil change"), or respond to a product search on the company website by promoting a particular product. During purchase, we can develop functionalities for cross-selling ("would you like to buy a tie with your shirt? ") or up-selling ("did you check [the better] brand X? "). After the purchase, we can conduct customer satisfaction surveys, or monitor customer posts on product forums, blogs, or social networks.
- *What information about the customer, the product, the promotional material, and the environment do we need in order to support various CEM functionalities?* For example, to anticipate a customer's need, what kind of information should we know about the customer? This depends on the need. For some basic needs (e.g. food), knowing little information is enough. For *lifestyle* purchases, a lot more is necessary. The same can be said about products. For some commodities, price may be all that consumers care about. However, health- or socially-conscious consumers may care about the *production process* of what they are buying, including environmental considerations, fairness, sustainability, or labor practices.
- *Once identified, how to capture customer's data supporting the CEM functionalities?* This involves a number of issues. Consumers are - understandably - loathe to supply personal/demographic information, and marketers have to be creative to get that information from actual or potential customers. Then, there is the concern of *subjective* information about the consumers, such as their beliefs, values, and sen-

timents about products, processes, and issues. Finally, there are *legal* and *ethical* issues related to the capture and usage of such data.

We propose to address these challenges within the context of *CA-CEM development framework* that enables us to:

1) specify the CA-CEM functionalities that we wish to support within the context of a purchasing process/experience;
2) translate these CA-CEM functional requirements into *software specifications*, in terms of required data structures and algorithms to support the CA-CEM functionalities;
3) translate those software specifications into actual code using predefined software artefacts (libraries, templates, generators, etc.).

Traditional so-called *application frameworks* embody an architecture/a design, with predefined *variation points* that can be instantiated for the application at hand. As such, they support only the third step of our framework. By comparison, the *development framework* that we propose would cover all the steps from business/user requirements to code.

To support the requirements phase, we need to understand/delimit the *problem space*, i.e. the *requirements space for CEM*. Toward this end, we rely on **a cognitive modeling of the purchasing process to identify** the various **decision points**, and the decision criteria/**influence factors** relevant to each decision point (see Section 3). This modeling helps us identify the touchpoints between sellers and customers that are needed to manage the customer experience. In other words, this cognitive model enables us to *script the interactions between sellers and customers in a way that allows us to* ***read*** *or* ***change*** *the mind of customers at pivotal moments of the purchasing experience* (see Section 4).

The cognitive/functional design of interactions between sellers and customers relies on:

1) the identification of the relevant information about: a) the customer (*customer profile*), b) the products, and c) the communication content between the customer and the seller (e.g. promotional material);
2) the selection or design of algorithms and tools to get that information from available sources;
3) the selection and design of algorithms to customize the interactions between the seller and the customer based on the available information.

The **relevant information is presented in the form of ontologies** in Section 5 **to be specialized or instantiated for specific purchasing scenarios** described in Section 7.3. We rely on **computational intelligence techniques to:** 1) **fill out some of the property values including subjective information from, and about the customer** (e.g. *preferences*, *sentiments*, or *beliefs*, see Section 6.1), and 2) **customize the interactions between seller and customer** through product recommendation, targeted marketing, and the like as described in Section 6.2.

This paper presents the principles supporting our approach. Section 2 introduces the process view of CEM, using a retail example. Section 3 describes the purchasing process, from an operational and a cognitive point of view. It identifies the different steps of the purchasing process, and the various factors that can influence customers in their decision making. Section 4 lays the foundations of our framework by going from the cognitive model of the purchasing process to, a) a generic CEM pattern, and b) sample instantiations of the pattern for step-specific of the purchasing process. These instantiations help design the ontologies needed to implement CEM functionalities, which are presented in Section 5. Section 6 shows how computational intelligence techniques can be used to fill out consumer profiles, and to customize the interactions with them. Section 7 shows how this framework can be instantiated to handle a particular purchasing scenario. We conclude in Section 8 with a discussion about future work.

2 A Process View of CEM

In this section, we use an example of a purchasing scenario that illustrates some of the possibilities and requirements of a CA-CEM framework. We start by presenting the scenario, and then present our process view of CEM.

2.1 A CA-CEM Scenario

Chris is a young urban professional in her 30's who walks into her favourite grocery store where she usually shops. As she drops items in her shopping cart, the food labels are automatically displayed on her (latest) iPhone. As she drops a box of crackers, she gets a warning because of its high-level of sodium - given her family history of blood pressure. She walks through the produce section, and gets notices about latest arrivals of fair trade certified products. She is pleased, of course, being an active member of Équiterre[1], but wonders if every shopper gets notified. Walking into the meat section, her attention is drawn to a special on lamb chops that her significant other enjoys immensely. She picks a rack and drops it into the shopping cart. She gets wine sug-

[1] www.equiterre.org, whose mission statement includes "Équiterre helps build a social movement by encouraging individuals, organizations and governments to make ecological and equitable choices, in a spirit of solidarity. We see the everyday choices we all make - food, transportation, housing, gardening, shopping - as an opportunity to change the world, one step at a time..."

gestions: two thumbs up for a Syrah from Northern Rhône, and one thumb up for a Shiraz[2]. The Shiraz wins with the ongoing special on Australian wines. Cheese with that? A French baguette? In the seafood section, Chris picks up a slab of Tuna steak. She gets a warning to the effect of having consumed big fish already four times already this week[3]. While getting toothpaste, she gets a notification about size 4 diapers on special. Considering that she has been buying size 3 diapers for the past six months, it is about time she switched to size 4!

Some might find this "idyllic" shopping experience frightening, and rightfully so. We do not presume that this scenario is desirable: we will simply explore the capabilities that make it possible. In particular, we explore the kind of data/knowledge needed about the products and the customer to make such a scenario possible. We will do this step by step:

1) "As she drops items in her shopping cart, the food labels are automatically displayed on her iPhone". This, of course, is easy enough, and can be done in many different ways, either by having short range RFIDs, which are probed by emitters in the cart, or simple bar codes scanned by devices upon being dropped in the cart. The technology for this already exists in many different forms.
2) "As she drops a box of crackers, she gets a warning because of its high-level of sodium - given her family history of blood pressure". This requires the system to have access to a medical profile of the customer: what conditions they already suffer from, or they are predisposed to.
3) "She walks through the produce section, and gets notices about latest arrivals of fair trade certified products". This notification depends on the system having access to the customer's social values. These values could have been entered directly by the customer when signing up for a particular service or social network. Alternatively, we can infer the customer's interests, based on their likes (Facebook profile), on the groups they are following (LinkedIn), on their membership to various advocacy groups that promote specific values (e.g. Équiterre), or on posts in various social medias[4].
4) "Walking into the meat section, her attention is drawn to a special on lamb chops that her significant other enjoys immensely". Knowing that someone likes a particular food product is easy enough. However, underlying this recommendation is a deeper understanding of human relationships, and what they entail: 1) people *live* with their spouses, and 2) as such, they eat togeth-

---

[2] Gracieuseté of an entry in the INTOWINE web site, http://www.intowine.com/best-wine-pair-lamb-chops, accessed 1/12/2014

[3] Big fish, because they are higher up in the food chain, contain more toxic substances such as mercury.

[4] Analyzed using big data and text mining techniques such as sentiment analysis for instance.

er. If Chris' office buddy, or sibling like lamb chops, the system should make no such recommendation.

5) "She gets wine suggestions: two thumbs up for a Syrah from northern Rhône, and one thumb up for a Shiraz. The Shiraz wins with the ongoing special on Australian wines. Cheese with that? A French baguette?". This is standard market basket analysis, combined with recommender functionality: it does not rely on any data specific to the customer.
6) "Chris picks up a slab of Tuna steak. She gets a warning to the effect of having consumed big fish already four times already this week". Such a functionality depends on the system embedding a number of health advisories, and a history of the customer's purchases.
7) "she gets a notification about size 4 diapers on special. Considering that she has been buying size 3 diapers for the past six months, it is about time she switched to size 4". This is an extreme/fine-grained case of *family life cycle marketing*. The concept of *family lifecycle marketing* recognizes that families go through a predictable set of stages (*family lifecycle*) during which their needs, means, and decision patterns, evolve [30]. Marketers take advantage of this lifecycle to better identify the target market segment, an adapt their marketing message accordingly. In our scenario, not only does our system recognize that Chris is in the so-called "full nest I" stage[5] - which can be inferred from her track record of repeated diaper purchases - but it is also predicting her *child's progression* through her/his lifecycle stages.

Hence, to make this idyllic (or frightening) purchasing scenario possible, the system needs:

- electronically accessible detailed product information, including nutritional information (1), production mode, toxicity (6), and value-based assessments and certifications (3),
- product associations (5),
- a detailed customer profile, including medical history (2), demographic data and lifecycle stage (7), relationships (4), tastes (4, spouse), beliefs and values (3),
- a history of purchases (6 and 7).

Note the history of purchases can be used to infer - or more accurately to guess - other customer profile data, such as tastes or beliefs and values. For the time being, we will not consider how the different pieces of information can be obtained. That concern will be addressed later.

[5] Identified as "young married couple with dependent children".

In the next section, we analyze these different pieces of information within the context of a process-oriented view of the interaction between an organization offering a product, and its customers.

2.2 A Process View of CEM

CEM aims at managing the interactions between a company and its customers. An actual or potential customer interacts with a company to fulfill a particular need, in the form of a product or service provided by the company. A company interacts with its customers, actual or potential, because that is its raison d'être: fulfilling the needs of its customers. A for-profit organization gets paid in return for the fulfillment of that needs at a price that is higher than its production cost. A public service organization (e.g. a government department) draws its value from the fulfillment of the needs of the citizenry.

A *customer* has a 'life of his/her own' pursuing his/her objectives of survival and happiness. In the process of pursuing those objectives, they have needs that can or need to be fulfilled by a company. Customer experience deals with the interactions between customers and companies around the act of need fulfillment: the product or service sale. Fig. 1 illustrates this interaction. The family lifecycle theory recognizes that the needs of people evolve during their life, as they enter different stages of the cycle. Those needs depend on many factors, including the processes underlying those stages (e.g. raising children), as well as the means that are typically available to customers at those various stages.

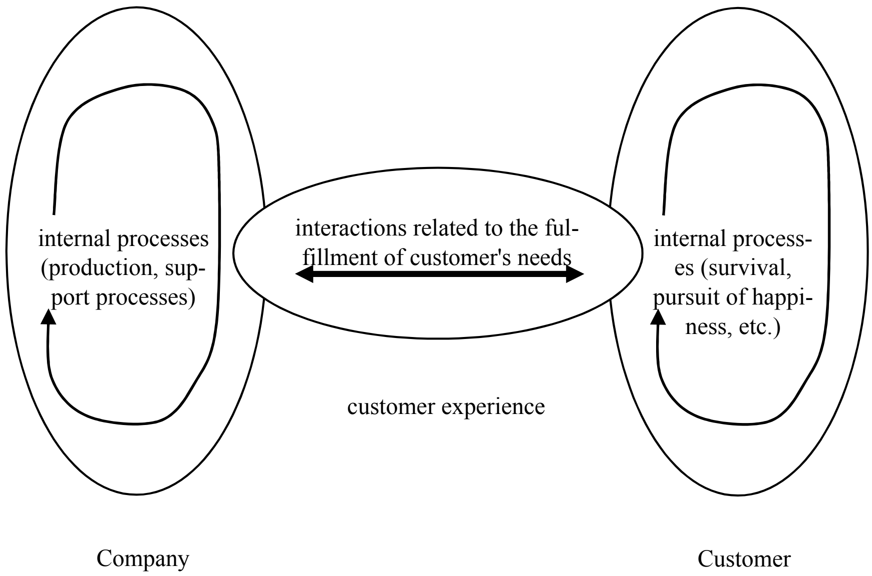

Fig. 1. Customer experience as the set of interactions between a company and a customer around the fulfillment of customer's needs

In the next section, we will look into the mechanisms of interactions covered by the customer experience by identifying the different stages of a purchasing process, and the various factors that influence the purchasing decision.

3 Understanding the Purchasing Process: a Cognitive Approach

Consumer behavior has been studied thoroughly by marketers and social psychologists trying to understand its mechanisms. They use the term "consumer", in a broad sense where the object of consumption can be a product/service (food, jeans, a cell phone service package, a car, a house), a behavior (exercising, dieting), or a belief (social values, political affiliation). At a fairly basic level, consumption is a conscious purposeful behavior, whose goal is to address a need or a desire. A number of psychological models have been proposed, including the *theory of reasoned action* [14], the *theory of planned behaviour* [1] the *MODE model* [13], the *theory of trying* [9], the *theory of self-regulation* [6], and subsequent variations thereof. The theory of reasoned action (TRA) perceives all actions as a purposeful behavior that starts by forming an *intention* to act, followed by the performance of the action itself [14]. TRA recognizes that *intentions* are influenced by two factors, namely, the *attitude*(s) towards the action, and the so-called *subjective norm*. The attitude towards the action is defined as the perceived likelihood of some outcome occurring. The subjective norm refers to the actor/consumer's perceived belief that (certain) people expect him/her to behave/act or not. To use a simple example, I intend to buy Nike shoes, because I think I will look cool (expected outcome for me), and I think my peers expect me to wear sneakers ("peer pressure"). Researchers and experiments have poked holes in this theory, which was amended to include other *salient* influences, such as actors' belief in their ability to complete the action (*perceived behavioral control*), leading to the theory of planned behaviour [1], and to account for *impulsive* action/consumption, leading to the MODE (*Motivation and Opportunity as DEterminants)* [13]. Other amendments took into account the *complexity* of the actions needed to consume, where the actions towards the consumption goal become themselves goals, leading to the *theory of trying* [9].

For the sake of clarity, we recall in this Section the synthesis presented by Bagozzi et al. [8], which integrates all of the influences that have been identified by researchers. The model is shown in Fig. 2. We briefly explain the various steps, and the influencing factors. The significance of 's model for CEM will be explored in Section 4.

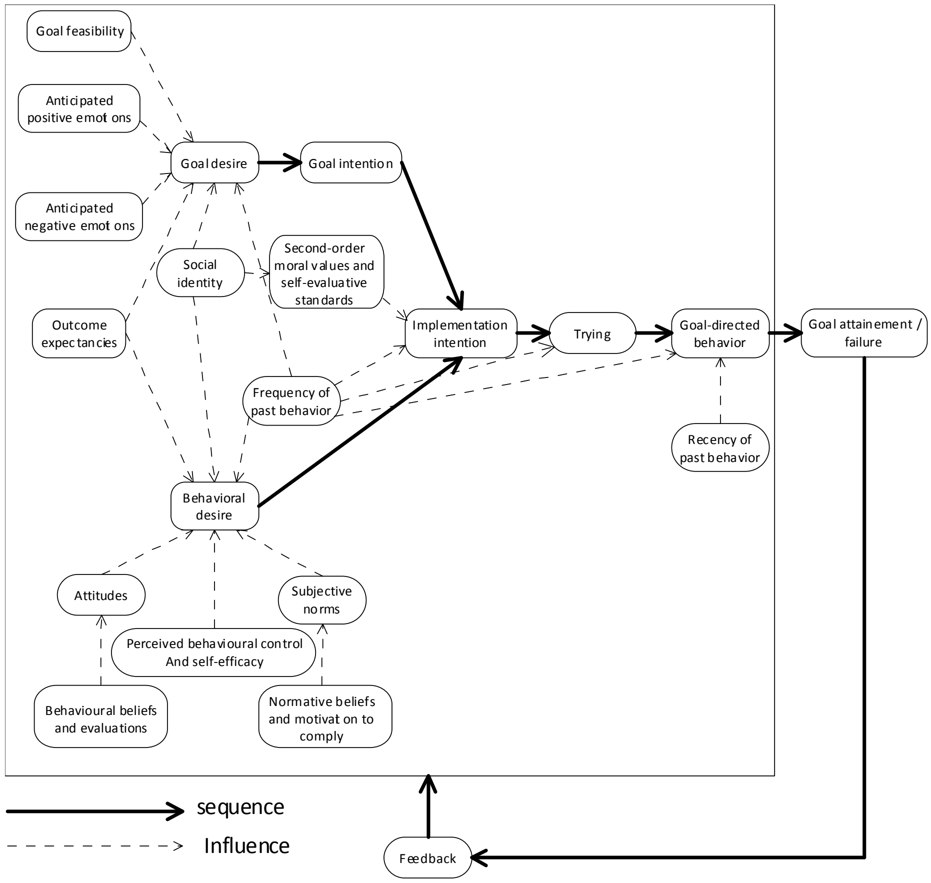

Fig. 2. A comprehensive model of consumer behavior adapted from [8].

First, we start with general observations/principles. First, there are different paths through the consumption process, depending on the type and complexity of consumption decisions. For example, habitual or low-risk consumption activities (e.g. grabbing a carton of milk) involve deliberation or planning. Second, social psychologists make the distinction between desirable ***goals*** and desirable ***behaviors***, with the former often preceding the latter. For example, my goal could be to lead a healthier life, which represents a desirable end state. This goal may entail alternate (set of) behavior(s), such as dieting and exercising, where the objective is a *behavior*[6]. For the sake of simplicity, we represent them as two independent process paths, ignoring the precedence relationship that can exist between goals and behaviors. Third, this consumption model makes a distinction between *desires*, *intentions*, *plans*, and actual *actions*:

---

[6] A good number of the studies that helped build these models concern *behaviors* like smoking, drinking, dieting, recycling, exercising, etc. Thus, the consumption process under study actually starts with the behavior desire.

- A consumer may have several desires (***goal desire*** or ***behavior desire*** in Fig. 2), but intends to pursue only one (***goal intention*** or ***behavior intention*** in Fig. 2). I would love to have a car to take my kids to school, and a motorcycle to take leisurely rides during the week-end (two goal desires), but I am going to stick with buying a car (goal intention). The same applies to behaviors - otherwise known as New Year's resolutions.
- Having decided to pursue a particular goal (goal intention) or behavior (behavior intention), the consumer needs to figure out *how* to achieve that goal, i.e. s/he needs a *plan of action*. For example, having decided to purchase a car, I need to figure out how to do it. This is as ***implementation intention*** in Fig. 2.
- Having devised a plan, the consumer then executes the plan by performing the actions of the plan. This is covered by the two steps, ***trying*** and ***goal-directed behavior*** in Fig. 2.

We will comment on some of the finer distinctions below. In the following, we describe each step of the process depicted in Fig. 2, along with their relevant influence factors:

1) ***Goal desire***: corresponds to identifying the various needs and desires of the customer. Several factors influence the setting of a goal:
    a) *Goal feasibility*: I may fantasize about owning an executive jet, but I will not consider it as a potential goal because it is unattainable.
    b) *Anticipated positive emotions*: this represents the perceived reward of attaining a particular goal, and it is a combination of positive emotions resulting from success (how good I will feel), and the expectation of success (how likely I am to succeed).
    c) *Anticipated negative emotions*: this represents the perceived penalty of failing to attain the goal, which is a combination of negative emotions resulting from failing (how bad I will feel if I fail), and the expectation of failure (how likely I am to fail).
    d) *Outcome expectancies*: the expected outcome of the pursuit of the goal (success vs. failure)
    e) *Social identity*: an individual's membership to a particular group, and the "emotional and evaluative [positive or negative] significance of this membership" [7].
    f) *Frequency of past behavior*: this was not included in 's framework[7], but is an important aspect of the *theory of trying*, which is one of the foundations of this framework. The basic idea is that customers may undergo the complex cognitive processes involved in goal setting once, or the first few times, but

[7] Fig 4.5, page 97, in [7]

after that, they make a cognitive shortcut: "I have thought this through many times, and found it is worth considering or pursuing this goal"

2) ***Goal intention***: Goal desire is concerned with the desirability of goals, but not with the decision to pursue them. This is step is where the decision is made.

3) ***Behavioral desire***: Researchers have identified six influences: 1) social identity, 2) outcome expectancies, 3) frequency of past behavior, 4) attitudes, 5) subjective norms, and 6) perceived behavioral control and self-efficacy. The first three were discussed above. We discuss the remaining others hereafter:

   a) *Attitudes*: were identified by the *theory of reasoned action* (TRA, [14]) as influencers of the pursuit, or not, of some actions. A *positive attitude* about an action leads the consumer to act, whereas a *negative attitude* inhibits him/her. The attitude towards an action is a combination of my belief about what would result from the action (what is the outcome), and my evaluation of that result. For example, if I buy jeans, I ***believe*** that I will look fashionable (the outcome), and I ***like*** being fashionable.

   b) *Subjective norms*: were defined by Ajzen & Fishbein as an individual's *perception* of the social pressure to perform, or not perform, an action [2]. That pressure is a combination of how confident I am that my spouse/colleague/peer group expects me to perform the action (a probability), and the perceived reward (or penalty) from complying or not with that expectation. For example, I am ***90% confident*** that my colleagues expect me to wear jeans, and if I do not comply, I ***believe*** that I will be eating alone.

   c) *Perceived behavioral control and self-efficacy*: Perceived behavioral control was introduced by the *theory of planned behavior* as a determinant in the decision to undertake or not, some actions/behaviors. It is defined as an individual's perception of how easy or difficult it is to perform a particular behavior [3]. It can be thought of as combination of my belief in me possessing a factor I think is needed to perform the behavior, and my belief in how important that factor is. For example, to stop smoking, I think that will power is everything, and I concede that I do not have strong will power. For the purposes of this paper, we will consider perceived behavioral control and self-efficacy as synonymous[8].

4) ***Implementation intention***: This is the planning phase. Once I have chosen a goal or a goal-directed behavior to perform, I plan for it. This is particularly relevant when the purpose of my consumption is an *end state*, i.e. when we are dealing with a *goal desire* as opposed to a *behavior desire*. In such a case, implementation intention consists of identifying the steps/individual actions I need to perform to reach my goal. These steps can themselves become goals, in their own

[8] There is some debate as to whether perceived behavioral control and self-efficacy are the same thing. Ajzen thinks so (e.g. http://people.umass.edu/aizen/faqtxt.html). Armitage & Conner do not [4].

right, if their performance is problematic [7]. The planning phase is influenced by the frequency of past behavior: if I solved the problem (of goal attainment) several times before, I can reuse the same solution (implementation plan). According to [7], a second influencing factor relates to *second-order moral values and self-evaluative standards*. Values are defined as the criteria or frame of reference by which people justify actions and judge actions and other people [24]. *Self-evaluative standards* represent how consumers see themselves, or what they want to become [7]. Our shopper from section 2.1 is a socially responsible consumer, and that is why the grocery store pitched the latest arrival of politically correct products. Fig. 2 shows the moral values and self-evaluative standards are influenced by the consumer's *social identity*.

5) ***Trying***: Within the context of the original *theory of trying*, the act of *trying* was thought of as "a singular subjective state summarizing the extent to which a person believes s/he has tried or will try to act" [9]. The definition was later extended to the actual execution of the plan established in the previous step, where each action in the plan can be problematic, and be forestalled [7]. Trying involves monitoring progress towards the objective, and making adjustments to the plan, as appropriate. uses the example of showing appreciation to a friend at their birthday (a *goal intention*) by buying her/him a gift (a *behavior desire*), which involves a plan of going to a shopping mall, selecting stores to visit, and browsing through merchandise to select a gift within the selected budget (*implementation intention*), to actually executing the plan (*trying*), which involves going to the mall, etc. [8].The consumer is not guaranteed to be able to do any of the steps required to make the purchase, including: 1) getting to the mall (transportation problem), 2) finding a set of stores of interest (not the right kind of merchandise), and 3) finding merchandise that fits the taste of the gift receiver, and the budget of the gift giver. The *trying* phase is influenced by the frequency of past behavior, discussed above, and the recency of past behavior. Within the context of the *theory of trying* [9], the recency of past behavior can influence current behavior in the same way that the frequency of past behavior does, because the behavior is still fresh in the consumer's mind

6) ***Goal-directed behavior:*** This refers to the final act in the consumption process, for example making a purchase.

7) ***Goal attainment/failure***: This is the step where consumers assess the extent to which they have reached their goal.

8) ***Feedback***: Based on the previous assessment, the consumers can adjust any of the choices or actions made in the consumption process, including the choice of goals to pursue.

This comprehensive model of consumer behavior, proposed by [8] with very minor adaptations of our own, is more or less the result of merging consumer behavior models pertaining to different kinds of consumptions for which some of the steps, or the influencing factors, may be irrelevant. Our framework for CA-CEM should provide guidance (or tools) for customizing this process for specific consumption or CEM scenarios. That is discussed in Section 7.1.

4 A Cognitive Approach to CA-CEM Design

4.1 A CEM Pattern

Section 2.2 showed a view of CEM as managing the interactions between two processes executed on behalf of two entities, the enterprise/seller and the customer, each with their own objectives. Those interactions are centered on the act of consuming. Section 3 presented the different steps of the consumption process, and identified the factors that influence each step, without considering interaction points between enterprise and consumer. In this Section, we put it all together.

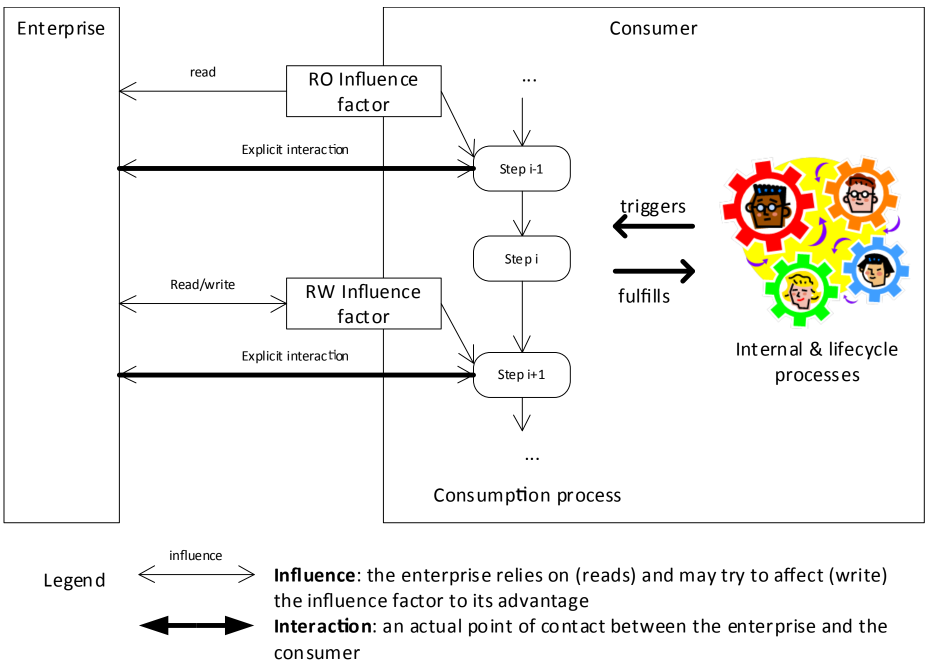

Fig. 3. Towards a systematic view of consumer experience management

Fig. 3 shows the idea. Consumers are active systems (living organisms) whose internal processes (to stay alive, pursue happiness, etc.) require a number of resources (food, clothing, transportation means) or conditions (fulfillment, happiness, etc.), triggering consumption processes to replenish the resources ("we are out of cereal") or to attain those conditions ("I need to go/dine out"). Those consumption processes involve a number of steps.

Decisions and choices taken in those steps are influenced by a number of factors, which have been identified by psycho-sociological studies of the consumer, and discussed in Section 3. Some of the steps of the consumption process involve interactions between enterprise and consumer. The interaction can be triggered by the consumer, for example by accessing the enterprise's web site, or triggered by the enterprise, by communicating with the customer, nominally or through advertising.

To the extent that some of the steps of the consumption process involve choices and decisions that are influenced by a number of psycho-sociological factors, the enterprise gains from knowing the values of those factors. Some of these influence factors are malleable, such as the *anticipated positive emotions*, and the enterprise may gain from modifying their values to its advantage, for example, by showing potential consumers how good life would be if they would purchased its products, or how easy it is to get approved for credit to make the purchase (e.g. perceived behavior control). Such factors are referred to as *read-write influence factors* in Fig. 3, as opposed to *read-only* (RO) influence factors, that enterprises cannot modify or act upon. Furthermore, while all the *read-write influence factors* are *potentially actionable* from a CEM point of view, some may not be worth pursuing.

4.2 Towards an Integrated View of Customer Experiences

Fig. 4 shows a first cut naive application of the pattern shown in Fig. 3 to all the steps of the consumption process, incorporating the full roster of influence factors identified by the literature on social psychology of consumer behavior [8]. However, such a view is reductive in the following ways:

1) A seller, acting through its e-commerce system, does not necessarily accompany a (potential) buyer through all the steps of the purchasing process. This is particularly true for the early steps of the process, such as the goal desire stage. This is not to say that sellers have no influence on goal desire but quite the contrary. For example, the tobacco industry, the beer industry, and the fashion industry, to name a few, all act upon the anticipated (positive) emotions, and to a lesser extent, social identity associated with the consumption of their products. However, this influence does not happen within the context of a one-to-one interaction between the customer and the seller's e-commerce system
2) As a corollary of the first observation, different channels will typically accompany different subsets of the purchasing process, often corresponding to different entry points into the purchasing process. For example, I can get into an e-commerce site through a web search for a product, where the returned page will display the product specs, and propose to add it to a shopping cart, thereby skipping "goal/behavior desire", "goal intention", "implementation intention", and part of "trying". However, a walk into a particular store will instantiate/trigger the full cycle, starting with a welcoming display featuring and promoting the desirability of a specific product.

3) A second corollary is related to the fact that a consumer may interact with several sellers for some steps of the process. This may be the case early in the process before the customer commits to a product and a seller. For example, you walk into a shopping mall and the location-based services on your phone send you advertisements from the different restaurants or shoe stores in the mall. Knowing that you are in competition with other contenders will impact the way you interact with the (potential) customer.
4) A lot happen in the trying step (executing a consumption plan) and the goal directed behavior step, that may include information search, comparison shopping, consulting product reviews, and so forth, i.e. lots of sub-steps that do involve interactions between seller and consumer. The pattern of Fig. 3 will need to be applied to those sub-steps.

With this is mind, in the next Section we study the various steps of the purchasing process, and see *whether* and *how* we can apply the pattern of Fig. 3.

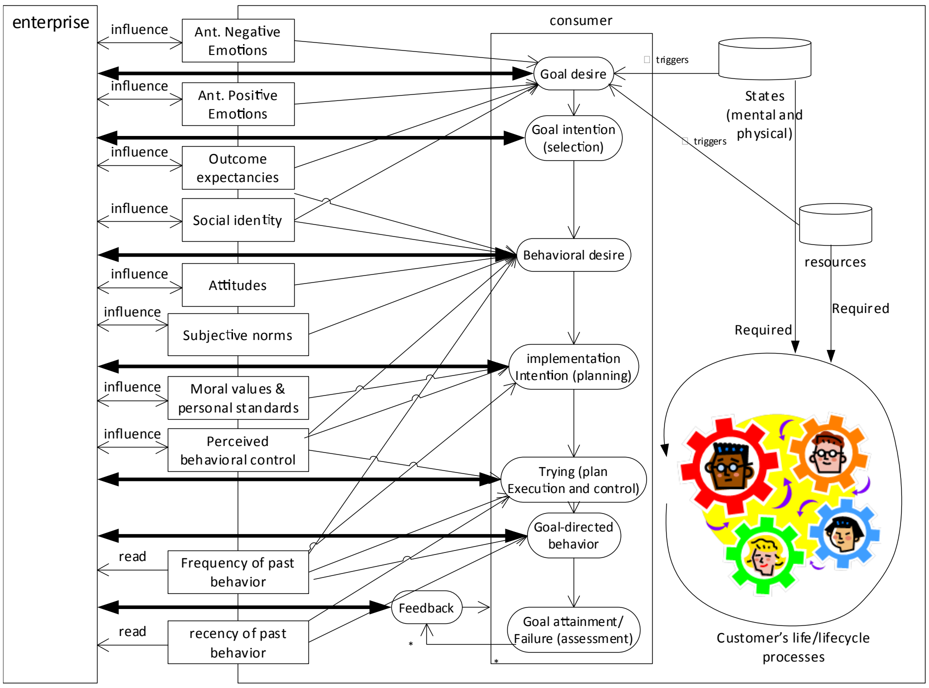

Fig. 4. A naive view of CEM

4.3 CEM as Service Design

The design of E-commerce applications in general, and CEM applications in particular, can be studied within the context of *service design*. We all have an intuitive understanding of what services are (as opposed to goods). They include things such as health care, education, transportation, hospitality, phone or telecommunications, etc. Wikipedia defines a *service* as an *intangible commodity*, where a *commodity* is a "marketable item produced to satisfy wants and needs"[9]. Service delivery typically involves a process that orchestrates people, components, structures, and automated systems - typically IT - to satisfy the needs/wants of the customer. *Service design* is concerned with the selection of the people, components, structures and automated systems needed to deliver the service and the design of the interactions between them in a way that is economically viable for the service provider and that maximizes customer satisfaction. Service design theory involves a number of disciplines, including marketing, operations management, organizational design, and technology (see e.g. [12, 15, 18]). The outputs of service design (good and bad) include the work organization at your local bank branch, your Department of Motors and Vehicles bureau, or your favorite department store.

There has been a growing interest in the research community to look at experience management within the context of service design (see e.g. [12, 26]), with the intent of incorporating *customer experience requirements* along with the other requirements that drive the service design. Teixeira et al. introduced *customer experience modeling* as a "model-based method [..] to represent and systematize customer experiences for service design." [26]. Customer experience modeling draws on Constantine's *Human Activity Modeling* (HAM) [11] on *customer experience requirements,* themselves from Mylopoulos' et al. *goal-oriented requirements analysis* [17], and on Patricio's et al. *multi-level service design* [20]. Fig. 5 shows the concepts of CEM and their relationships. Note that customer experience requirements apply to both *activities* and the *context* within which these activities take place, including other actors (e.g. a sales clerk), artifacts (e.g. a passive display), and technology enabled systems, including computing devices, sensors, and the like.

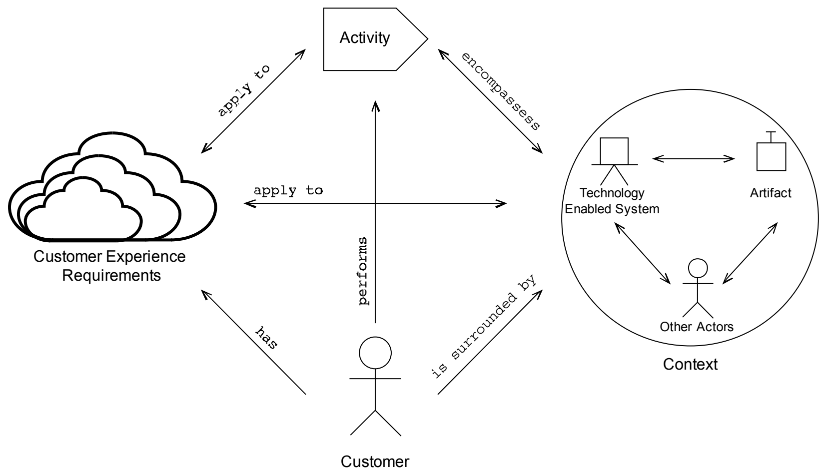

Fig. 5. Incorporating *customer experience requirements* in service design. From [26].

[9] See https://en.wikipedia.org/wiki/Commodity, accessed on 18/8/2015.

An important aspect of service design is *customer scripting* (see e.g. [12]), which is the specification/design of the interaction process between a customer and a service provider delivering the service. This involves the careful design of both, a) the individual interactions and interaction touchpoints, and b) the orchestration of such interactions. As an example of a), we explore different ways for the consumers to identify themselves, and select the easiest ID method. An example of b) is whether to request identification from the beginning of a service script (which would result into more appropriate/customized service interactions), or adopt a light process, and wait to request identification until the customer proceeds to checkout. For the most part, the *customer experience modeling* method is *descriptive*, providing modeling concepts and notations. However, it offers little guidance as how to design a service in a way that it conforms to the customer experience requirements. Cook et al. [12] identified some of the *human* and *subjective* issues surrounding customers' appreciation of their service encounters/experiences, independently of the intrinsic qualities of the product/service rendered, and of the customer specifics. For example, one of the script design guidelines, based on psychological studies, suggests that we design a customer script that puts unpleasant interactions/steps first, and ends with a high note [12].

Within the context of service design, the description of the purchasing process shown in Fig. 2 provides a *high-level solution space for consumer scripting*. As mentioned in Section 3, our e-commerce (or experience management) application will have to implement a sub-script (sub-process) of the process described in Fig. 2, depending on the type of purchase (carton of milk, versus car) and on the channel (brick or click). The pattern of Fig. 3 shows us what to look for to design the individual interactions/steps of the script. We see next how to turn the psychological influences described in Fig. 2 into specific actions/interactions to incorporate into a consumer script.

4.4 Turning Psychological Influences into Actionable Consumer Experiences

In this Section, we apply the pattern of Fig. 3 to the steps of the purchasing process depicted in Fig. 2. We start by outlining the principles (Section 4.4.1), and then we look at specific examples (Section 4.4.2).

4.4.1 Principles

We start by categorizing the purchasing process steps, and their influence. Then, we present the different ways we can exploit those influences.

Roughly speaking, the purchasing process involves two kinds of activities, 1) *internal* cognitive activities, taking place within the consumer's head, and 2) *external,* possibly *physical*, activities, some of which involving explicit interactions between the customer and the seller/service provider. Examples of cognitive activities include *goal/behavioral desire*, *goal intention* (selecting one worthwhile goal among many) and *goal implementation* (planning a way to achieve the goal). Examples of external activities include *trying* (plan execution) and *goal directed behavior*. An example of external activity that involves an interaction with the seller is doing an on-line product search on the seller's portal, asking an in-store sales clerk for advice, or checkout (online or in store).

We categorize the influence factors based on two dimensions:

1) *Subject/scope*: customer-specific factors (socio-demographic data, social identity, moral values and personal standards), product/service-specific factors (product specifications), and factors related to the <customer, product> relationship (anticipated positive emotions, anticipated negative emotions, history of past behavior),

2) *Type*: objective/factual factors (customer socio-demographic data, history of past behavior, product specifications) versus subjective/emotional/perception-based factors (attitudes, subjective norms, perceived behavioral control).

From a service script/interaction design point of view, for each activity of the purchasing process, we can do different things, depending on the kind of activity, and the kind of factors influencing that activity:

1) Purely cognitive activities. These activities happen inside the consumer's head (e.g., goal intention), and do not require interaction with the seller/service provider. Thus, we need to *provoke* an interaction, as part of our service script design. This interaction can take place:

    a) *prior to the activity* to modify the subjective/emotional/perception-based factors. For example, prior to goal desire/goal intention, I can strengthen the anticipated positive (negative) emotions associated with a goal, to make the customer desire it *and* select it over *other* competing goals.

    b) *after the activity*. In this case, we cannot influence the activity, but we can at least hope to get a *reading* on the choices made so that we can better prepare the response of the service provider to the subsequent activities. For example, as a retailer I may not be able to influence goal intention (selecting one goal to pursue, among many desired, e.g. which item of clothing you came into the store to purchase, among all the ones we offer). But if I can tell (or ask!) which item you came in for, I can direct you.

2) External activities that *normally* involve an interaction with the seller/service provider. In this case, the activity is part of the normal script, and the customization concerns the information content, and how it is delivered, both adapted to what we know about the customer. An example of such activities includes searching the seller's online catalogue. If we know who the customer is, we can refine the search query, and customize the presentation of the results.

3) External activities that do not normally involve an interaction with the service provider. Here too, we need to *provoke* an interaction with the customer, either before the activity takes place, to influence its outcome, or, failing that, after it has taken place, to get a reading on what the customer did, to plan the subsequent steps of the process/service script. An example of such an activity is searching for

product reviews (pre-purchase), or entering a product review (post-purchase), *in a third party site*.

The difficulty with the activities that do not (normally) involve an interaction with the seller/service provider (via their e-commerce software) is to *guess* where the customer stands along the purchasing process. A seller may adopt a default strategy of assuming that you are at the (pre) goal desire stage, and send you advertisements until/unless you do something specific. However, when we get into the intermediate stages (e.g. goal intention, goal implementation, and the various substeps/sub-stages of trying), things get tricky. Not knowing where the customer stands in the purchasing process can be challenging even when the customer initiates an interaction. When you walk into a dealership, the salesperson first asks you whether you are looking for a model in particular, or just looking, and if you answer the latter, they may go back to playing Solitaire. Similarly, when you do a product search on an online catalog, you may be looking for a particular model whose model number you forgot, or looking for any product that matches your specifications. Depending on the context, there are different strategies for completing a sale!

4.4.2 Examples

In this section, we show how to use the concepts and principles described above to design interactions for specific purchasing process steps. The purpose here is to illustrate the kind of analyses a customer experience designer needs to make. The final answer would rely/depend on marketing knowledge. This analysis will also help us start determining how the identified influence factors translate into data that an e-commerce application can manipulate. Section 4.4.1 will propose a first-cut ontology/metamodel to support the experience management functionalities illustrated throughout the paper. We start by a thorough analysis of the goal desire stage, and then present pointers related to other stages.

At a basic level, goal desire is triggered by the lack or deficit of physical resources ("we are out of milk") and physical or mental states("I am out of shape","I need some fun") (see Fig. 4). With the exception of biological needs, for which we need no prompting, companies/marketers can *create* a need[10] or strengthen a need through advertising. According to 's framework (see Fig. 2), the *strength* of such desire is influenced by anticipated emotions (positive and negative), outcome expectancies, and social identity. This raises two kinds of questions:

1) What is the appropriate advertising message *content* to strengthen the desirability of a goal? In particular, we need to think of the extent to which a (better) knowledge of the customer profile can help customize/select the appropriate message

2) How to deliver that message within the context of a CEM application, *considering the stage of the purchasing cycle we are at*?

---

[10] Marketers would disagree with the manipulative "create": they prefer the term "recognize".

At a basic level, we can strengthen the desirability of a product/service or behavior by projecting images that: a) strengthen the anticipated positive emotions ("how good it will feel to own this product or adopt this behavior"), b) strengthen the anticipated negative emotions ("how bad I will feel by *not* owning the product or not adopting the behavior"), and c) show how easy it is to succeed (in acquiring the product or adopting the behavior).

In what ways does our knowledge of the customer make such images more effective?

1) By depicting someone the customer can identify with. If we want to "sell" the latest Android smart phone to a teenager or professional, we present an advertisement that features a teenager, or professional, respectively. Idem with gender, ethnicity, geography, etc[11].

2) By sending images of more specific goals, appropriate for customers' category/social identity. While owning the latest and greatest in smart phones may be a broadly shared desire, the advertising message could focus on social media and multimedia capabilities when targeted towards teenagers, or focus on office productivity tools, when targeted towards professionals. Similarly, if a product is too expensive for my social category, the marketing message can stress the possibility or ease with which I can finance the purchase, playing on outcome expectancy.

3) By sending images of goals *we know the customer has not yet achieved*. My cell phone provider, or the dealer who sold me my last car, know the features of the current product I own, and know which feature(s) to highlight or stress in the marketing message.

Having determined/selected the advertising material, we need to determine ways to deliver it. This depends on the combination of two things: whether the customer is known to the seller, and the channel. For example, if the seller already knows the customer, then we have several possibilities:

- We can push the advertising material through an unsolicited e-mailing campaign (**on-line**).
- We can present the material to the customer when they log in (**on-line**).
- When the customer enters the store, and s/he is positively IDed, the message is displayed on an appropriately located monitor (**in-store**).

If we do not know the customer, then we broadcast the message variants, both through the company portal and in-stores.

In summary, for goal desire:

- We can act upon the desirability of goals through advertising.

[11] Typically, marketers produce different variants of the same marketing theme, aimed at different populations

- Knowing that anticipated positive emotions, anticipated negative emotions, and outcome expectancies influence the desirability of a goal helps determine the *orientation* or *content category* of the message carried by the advertising material.
- Knowing the customer (social identity, history of consumption) can help customize the message.
- The channel determines the advertising material delivery method.

Similar analyzes can be performed on the subsequent steps of the purchasing process. We provide some highlights for some of those steps:

- *Goal intention*: Here, the issue is in selecting a goal to pursue among many that are deemed desirable. This is another internal cognitive activity, that is practically indistinguishable from the goal desire. The priority of goals depends on means, lifecycle stage, values, social identity, etc. I can influence goal selection through advertising by matching advertising content to consumer profile. For instance, if I am a specialty outdoor equipment manufacturer, and I am talking to a young person who can afford ***either*** a new car or a (top of the line) mountain bike, I can show her/him an add that features healthy and handsome young people taking their fancy bikes to an extreme mountain trail in a cheap ride, or even hitchhiking there!
- *Goal implementation*: This is the planning stage. Planning makes sense in a multichannel shopping experience. For example, I have already made the decision to buy a new fridge, and I am planning my shopping experience. How can a retailer influence the planning stage to his advantage? Several general strategies includeoffering a simpler shopping experience (e.g., order from a catalog, or a better online ordering system), and making sure that any shopping plan includes an interaction with the retailer ("we offer the best price guaranteed", or "we will best the lowest price you will find by 5%"). Retailers can also take advantage of knowledge about customers like their location, by suggesting stores to visit to shop or pick up their merchandise, or their shopping style, by activating/proposing appropriate shopping scripts following an on-line search.
- *Trying*: This is the execution of the consumption plan (*goal implementation*). If the plan includes in-store browsing, then we can think of the CA-CEM scenario presented in Section 2.1, and various product recommendation strategies especially if we know who the customer is. Imagine walking into a department store and having your profile displayed on the PDAs of nearby salespeople who hence call you by name, and direct you to your style section.

Through these analyses, we do not pretend to provide perfect answers from a marketing point of view. However, we hope to provide the *design vocabulary* that service designers, including marketing specialists, can use to design CEM-enabled services.

5 Ontologies for CA-CEM

Section 4 showed how the combination of cognitive modeling of the purchasing process and service design theory can help us in:

1) *identifying the steps* (cognitive or physical) of the purchasing process *that*, *either naturally lend themselves to an interaction between seller and customer, or would benefit from provoking such an interaction.*

2) *selecting the modalities for such interactions*, depending on the stage of the purchasing process, and the purchasing channel.

3) *selecting the contents of such interactions,* depending on our knowledge of the customer and the product.

In this section, we focus on the last aspect, **the modeling of the knowledge we need to have about the customer and the products.**

We present the required knowledge models in the form of *ontologies*. We use the term *ontology* in two complementary ways:

1) In the computer science/information systems sense, as specification of a set of representational primitives modeling a domain of knowledge [Gruber, 2009]. If we think in terms of UML and MOF, the above definition corresponds to the *UML metamodel* (M0), i.e. the UML modeling language, a subset of UML itself. An ontology can also be understood in terms of description logics, where an *ontology language* can be seen as the realization of an underlying description logic.

2) In the sense of reflecting a *shared conceptualization of a domain*.

Indeed, within the context of our CA-CEM development framework, we want to specify the modeling ingredients that analysts can use to represent user's and product profiles required by purchase scenarios. We also use the term *ontology* in the sense of representing shared knowledge about domains, namely marketing knowledge in the context of CEM. Such marketing knowledge may include things such as known customer categories, their likes, tastes, and spending habits, known product categories, and their appeal, etc. However, to *specify* such knowledge, we need representational primitives - hence the complementarity of the two perspectives.

In this section, we focus on the *representational primitives* needed to represent the data/knowledge about customers, products, promotional material that may be needed, within the range of possible purchasing scenarios, and CA-CEM functionalities.

5.1 Consumer Data

Knowing the consumer is critical to a successful CEM. As shown in Sections 2 and 3, there is a lot of relevant data. The different kinds of data are discussed in the following.

As mentioned in Section 4.1, the influence factors that are relevant to the consumption process steps represent information that consumers use to select their choices and make their decisions. Thus, companies benefit from knowing what that information is, and should incorporate it in the consumer's *profile*. Let us take a couple of examples from the scenario presented in Section 2.1:

1) The customer's like for lamb chops[12]. This is an example of *anticipated positive emotion* (enjoyment) resulting from eating/getting lamb chops.
2) The customer's sustainable/equitable development values. This is part of the *moral values and personal standards* factor, which depends on the consumer's *social identity*[13] [5].
3) The same-week purchases of big fish. This is part of the *frequency/recency of past behavior* factors. In this case, it consists of customer's transaction histo

Part of this data is specific to a given customer, for example her/his transaction history or her/his liking of lamb chops, while other parts pertains to a category that the customer belongs to, for example fair traders. Also, parts of the data are specific to a product (e.g. liking lamb chops), while some are more generic (e.g. liking any fair trade product). We discuss product variations in Section 5.2.

Bagozzi's framework does not explicitly mention the consumer's socio-demographic category, although part of this information is implicitly contained in the consumer's social identity. We see such data, especially the one relating to lifecycle processes, as the *main driver of needs*, particularly, basic ones (food, clothing, shelter, healthcare, education). Our imaginary scenario of Section 2.1 also highlights the role of *relationships* in anticipating consumers' needs. Some of these relationships are implicit in the lifecycle concept, namely in the nesting stages where the consumer makes purchases (or purchasing decisions) for the benefit, and on behalf of dependents.

At any given point in time, consumers may be in a particular lifecycle stage. They can go from one stage to the next after some time has elapsed, and/or after some event has occurred. For example, the typical family lifecycle has a stage for young adults with preschool-age children, followed by a stage for school-age children. For a given child, the duration of the "young adults with preschool-age children" stage is 4 to 5 years. Similarly, the family lifecycle has an "unattached young adult" stage, followed by a "newly married adult, no children" stage. Marriage or otherwise attachment indicates the transition from one stage to the next. To the extent that a CEM system will typically *not* have access to court marriage certificates, some events will have to be *inferred* from *property value changes*. For example, when teenage or young adult children move out, they start having basic household needs (furniture, telecom services, etc.). Such an event can be detected by the change of address, from their par-

[12] Ignoring, for the time being, the fact that it is the customer's significant other who likes lamb chops.
[13] Which may be *characterized*, if not *defined*, by the customer's membership to, or militancy within, various organizations and associations. In our example, our shopper Chris is member of www.equiterre.org

ents' home address to an outside address. The reverse move back home may also be of interest to marketers, as the disposable income of the returning children increases. Fig. 6 shows a customer profile metamodel. Starting with the upper half, it shows that **Consumer**s have properties (**Property**) and values for those properties (**PropertyValue**). The association between a **Consumer** and a **PropertyValue** is characterized by a *confidence level*, reflecting our confidence in the value. For example, a value that was entered by the consumer when signing up for a service is more reliable than a value inferred by data mining techniques.

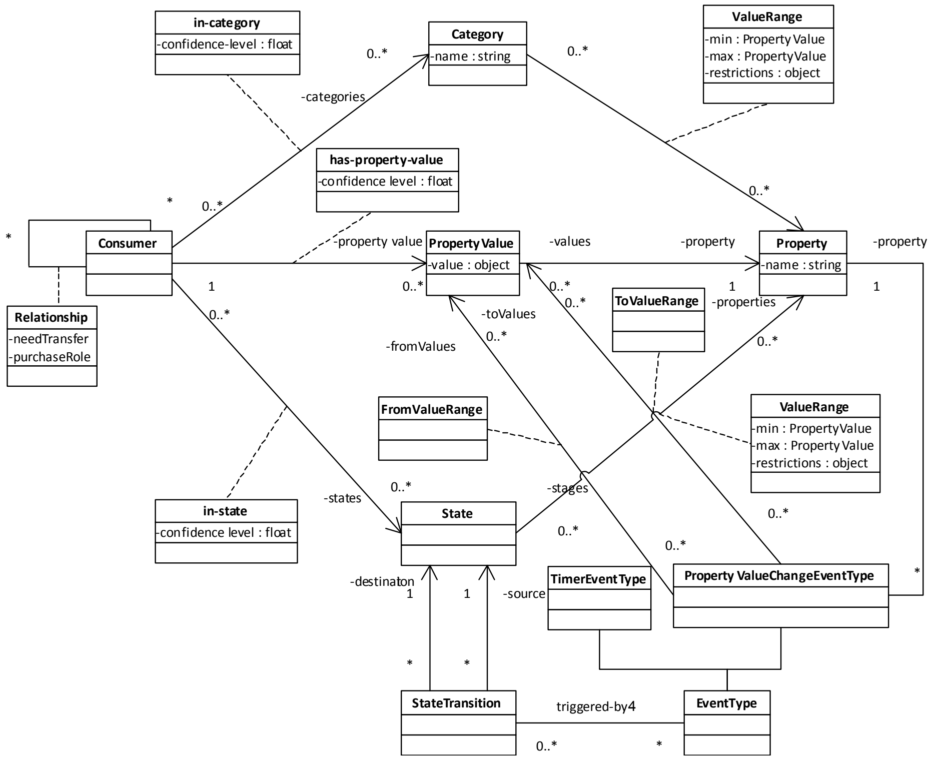

Fig. 6. A customer profile metamodel taking into account decision factors, demographics, and lifecycle data.

Consumers belong to **Category**(ies), with some confidence level. For example, we may not know for sure that Chris is a socially responsible consumer (a **Category**), but we know that she constantly accesses the web site of equiterre.org, and posts comments on their articles. From which we can infer with 80% confidence level that she belongs to the category of socially responsible consumers.

Category(ies) are also characterized by instances of **Property**. The relationship between a **Category** and a **Property** (the association class **ValueRange**) shows which property values are admissible, or expected, for this category. For example, for the property **PreferredTransportationMode**, socially responsible people may be known to prefer public transportation or self-propelled vehicles. Value ranges for properties of categories provide *default property values for members of that category*. For example, we may have no value for the **PreferredTransportationMode** property for Chris, though she is a member of the socially responsible consumer category, hence we can infer that she prefers public transportation.

Similar to the concept of **Category**, we have the concept of **State**. A consumer can be in different *states*, which correspond to qualitatively different consumption behavior in terms of needs, attitudes, etc. This accounts for the family lifecycle theory, where states represent lifecycle stages. It can also be used to represent finer concepts such as diaper stages, discussed in the scenario of Section 2.1 **StateTransition**s are triggered by events (**EventType**), which can be property change events (**PropertyValueChangeEventType**, as in getting married), or timer events (**TimerEventType**), as in going from size 3 to size 4 diapers after a fixed number of months. A property value may change with a worthwhile event in general, or if the initial value of the property equals a specific value (e.g. moving out of the parents' home, signalling need for furniture), or if the destination value equals a specific value (changing marital status, or moving to a significant other's house, or back to the parents house), hence the association classes **FromValueRange** and **ToValueRange** from **PropertyValueChangeEventType** and **PropertyValue**.

The **Relationship** association class from/to **Consumer** represents personal relationships of consumers, to the extent that they are relevant to consumption behavior. In our fictional shopping scenario, Chris gets lamb chops because her mate likes them. She was also told about size 4 diapers because the system suspected that she cared for a toddler, who had been size 3 for X months. Some relationships *transfer* needs: if I care for a baby, the baby needs diapers, then I need diapers, and I will be attentive to diaper adds, for example. The message sent to me may be adapted to my role in the purchasing process: whether I am an advisor, a decision maker, or a doer. We use the attributes needTransfer, and purchaseRole to represent these distinctions.

The model in Fig. 6 is more of *a metamodel* or a *high-level ontology* than a customer profile data model per se. An actual customer profile model can be constructed by *instantiating* the classes shown in Fig. 6. For example, we need to figure out which properties are worth representing about a customer depending on the product or service being sold, which consumer categories are worth considering, which lifecycle stages are relevant to the kind of product we sell, etc. This will be part of the framework instantiation discussed in Section 7.

5.2 Product Data

We showed in the previous subsections the kind of data that a company needs about its customers to assess their needs and desires (e.g. demographic and lifecycle data), and understand their attitudes, biases, emotions, and values that come into in the process of satisfying those needs. This information is used in identifying those, among its products, that best help address the needs and desires of those customers, and positioning those products in a way that appeals to customers, considering their (known) attitudes, biases, emotions, and values. To take full advantage of our knowledge about consumers, we need a comparably rich representation of the products and services sold by the company.

In our grocery shopping scenario, we knew the function of our products (nourishment), their composition (sodium content for crackers, but also mercury content for big fish), their assortments (a Shiraz, to go with lamb chops), and their production process (faire trade or not). Different kinds of products have different facets. Fig. 6 partially illustrates what our data could look like. We will comment on the most important aspects of the model.

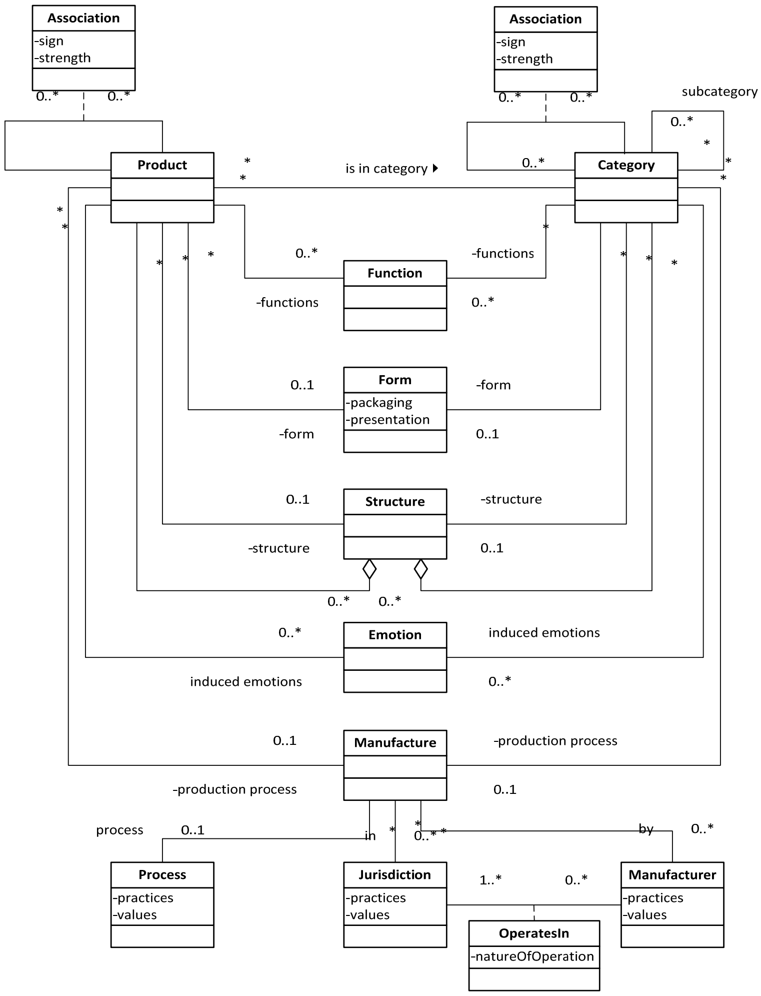

Fig. 7. An example of product representation for rich CEM functionality

Products are represented using five facets: 1) *function*, which represents the objective utility of the product (e.g. nourishment, transportation, cleaning product), 2) *form*, which represents things related to packaging (e.g. bulk, by unit, six-pack) and presentation (visual aspects), 3) *structure*, which can be used to represent the ingredients or nutritional information of a food item, or the major functional components of a device (e.g. a car having a V6 turbo engine, all-wheel drive, 17 inch wheels), 4) *emotions*, which represents, in a shorthand form, the emotional function of the product (the emotions it elicits), and 5) *manufacture*, which provides information about the manufacturing/construction process of the product. As a first-cut model, we include the process itself, the jurisdiction where the product was produced, and the manufacturer, all of which being information that environmentalists, fair traders, and various social activists might care about. A conscientious consumer may refuse to purchase the product of an environmentally unfriendly process (e.g. fish captured by bottom trawling), or manufactured in a jurisdiction known for poor environmental or labor regulations, or by a manufacturer operating in such a jurisdiction, or headquartered in a jurisdiction known as a tax haven. Finally, the association class **Association** represents positive ("goes with") or negative ("does not go with") associations between products, and the strength of the association ("goes really well with"), as in "the *Bleu de Brest* blue cheese goes really well with the 2011 Australian Shiraz".

In the same way that consumers are members of categories (market segments), products belong to categories. Each product category is characterized by the same five facets (*function*, *form*, *structure*, *emotion*, and *manufacture*). Categories hold *default information* known about classes/sets of products, which can be overridden for specific products. The categories are organized along specialization hierarchies using the subcategory relationship. We can have different categorizations, based on function, country of origin, packaging, etc. Finally, similar to the associations between products (goes with/does not go with), we have associations between categories: red wines go well with meats, white wines go well with fish, and polka dot shirts do not go well with pinstripe suits.

5.3 Other Kinds of Data

Different types of data may be relevant to interactions between seller and customer, with the purpose of capturing the needs of customers, their appreciation of a particular product or product category, information about products, such as technical product reviews, product comparisons, as well as promotional material, as described in Section 4.4.2 Good experience management depends on soliciting (customer → seller) or presenting (seller →customer) the right information at the right time. In this section, we illustrate the kind of representations of some of this data that would support fine-grained matches between products, customers, and context.

Let us take the example of promotional material. We showed in Section 4.4.2 how the message conveyed by promotional material can be matched to the customer and the situation at hand for maximum impact. To be able to realize this match automatically, we need to index the promotional material with a precise description of the message.

Fig. 8 shows a sample model for representing promotional material. The model is broken in two parts separated by the dashed line. The upper part reproduces a subset of the models shown in Fig. 6 and Fig. 7 that embodies our marketing knowledge, namely:

1) Certain consumer categories (**ConsumerCategory**) have a need for specific functions (**Function**).
2) Those functions are supported by certain products (**Product**) and product categories (**Category**).
3) Those products or product categories induce **Emotion**s.

Recall from Section 4.4.2 that in order to re-enforce/strengthen customers' desire of a particular product or function (the so-called *anticipated positive emotions*), one can show them a person they can identify with, who is using the product or the function, and who is experiencing those *positive emotions*. Hence it involves an **Actor** (as in comedian in a video), who is identified as belonging to a particular category. A video clip (or some other promotional material) would be indexed with the action taking place, which is a **ProductUsage** of the target **Function** or the target **Product**, and should be shown to experience the target **Emotion**. Thus, for promoting the latest smart phone, I could have two different clips: one featuring college students exchanging using their phones social media, and another clip showing business people/professionals using the office productivity tools of the phone. I would then send the appropriate clip to the appropriate customer by matching on customer category.

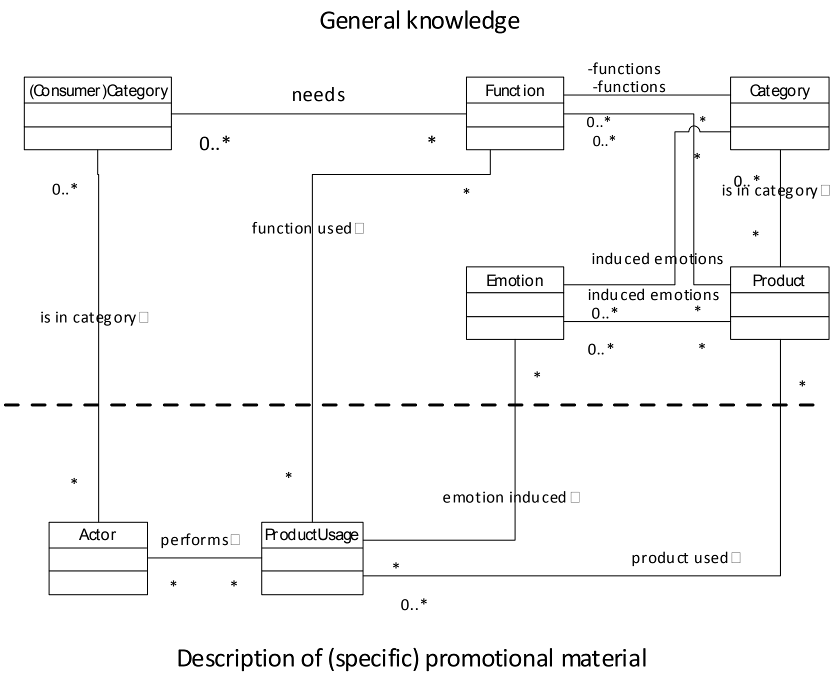

Fig. 8. A representation of promotional material

6 Computational Intelligence for CEM

6.1 Data Mining Techniques for E-Commerce

Data mining techniques have long been used in electronic commerce, for a variety of usages. Some of the first usages included the use of navigation traces for various purposes, including knowing more about incoming traffic (geographic location of users/potential customers, referring pages), optimizing the design of websites, and *recommender functionalities*, suggesting pages to look at from the pattern of traces, or suggesting products, based on customer profiles, on purchase history, or basket market-like analyses, of the kind we find on amazon.com (see e.g. for a survey of the early uses [23]).

As B2C customers started leaving *textual* traces on the internet, *text mining* techniques found many additional customer *relationship* management applications, including intelligent routing of users queries left on service portals, *sentiment analysis* to discover how customers *feel* about a company's products or actions in terms of a unidimensional polarity, and *opinion mining* for finer analyses of customers' views towards products and actions (see e.g. [16] for a thorough treatment of the various mining flavours).

The advent of social media has opened up a whole new set of possibilities and challenges:

1) On the positive side, whereas the textual traces left on customer service portals were typically limited to *actual* customers *after* they had made purchasing decisions, companies can now find out what is said about them and their products on other media, by non-customers, or not-yet customers *before* they make a decision. Companies are then capable to act upon, and influence some of the *earlier* steps of the purchasing process (see e.g. [21])
2) On the negative side, the proliferation of social media poses many operational challenges:
    a. CRM/CEM specialists now need to collect the information from *many* different sources (specialty blogs, portals, Facebook posts, Twitter tweets, etc.).
    b. We are no longer able (or less able) to trace the opinions or sentiments that are expressed on the web to individuals that we can reach out to address their needs or change their perceptions[14].
    c. We are not always able to manage identities across different social media. For example, if we find three similar complaints posted on three different media, under three different pseudonyms, do we have three different service failures, or is it the same person blasting us everywhere.

[14] If an unhappy customer leaves a complaint on the company's portal, the company is able to connect with that customer and remedy the situation.

In the next section, we look at the different potential usages of data mining techniques within the context of CA-CEM.

6.2 Data Mining for CEM

Data mining techniques can be used in at least three areas:

1) Categorization refers to both the *identification/specification* of categories, as in identifying or characterizing a particular consumer group (e.g. DINKs), and the assignment of individuals (a particular consumer) to categories, e.g. recognizing that the customer in front of me is a DINK.
2) Opinion/belief mining tries to figure out what the customer thinks or feels about things or issues.
3) Feedback assessment refers to the ability of sellers to automatically assess, and act upon, the feedback left by customers on their product ownership experience, or on their *customer* experience, i.e. their service interactions with the seller.

We discuss below the main issues raised by these three applications, and their operational implications on our CA-CEM development framework.

6.2.1 Categorization

Categorization is a very important aspect of CEM. The gist of CEM is knowing the customer, i.e. associating them with a consumer group that has identifiable needs, desires, and consumption patterns. As we saw in Fig. 6, a lot of information about the consumer will be inferred from the knowledge of her/his membership in specific consumer categories, as opposed as being specifically given/known for that consumer. Categorization raises two related issues: how to discover/codify categories, and how to categorize/classify new instances of individuals. The two issues are related, and will be discussed jointly. We discuss below two types of categorization, depending on how much prior marketing knowledge was have about the market we serve.

Categorization through an explicit codification of marketing knowledge

We codify marketing knowledge in the form of identifiable consumer groups along with socio-demographic properties/property value ranges. Those properties can be defining (for example for an age group, age is defining) or characteristic (for example for residents of a particular neighborhood, income level or ethnic origin are characteristics).Knowledge representation theory may dispute the distinction.

Although definitional properties may be difficult to obtain from customers (age, income), we can predict them from characteristic properties directly observed. These predictions can be made by machine learning based algorithms trained on customer data sets for which we have categorization and characteristic properties. We can also imagine a tool that enables us to tag the properties of a model as defining versus characteristic, thus triggering real time analyzes as data is inserted.

New instances get categorized based on whichever combination of properties are available. The system can then predict the value of other properties from the definition of the category with different confidence levels, which would determine the strength of the inferences that we draw from those inferred property values.

Categorization through unsupervised learning

The idea here is to use unsupervised learning techniques to discover new consumer categories. Clustering algorithms are fed customer's data. The quality of the categories will depend on the richness of the data. This is somewhat related to the earlier discussion about *definitional* versus *characteristic* properties. Among the relevant issues is whether the data that is being captured is discriminating for the classification of new instances.

Once instance data (consumer or product) has been clustered, we can compute for each cluster the value distribution of each property for the instances to get a better characterization of the cluster. Thus, if a customer is found to belong to a certain category, and is missing some property values (e.g. yearly income), we can infer values for those properties based on the property value distribution for that category[15]. For example, if I know that a customer is a DINK, and I know the family yearly income distribution for DINKs, I can translate that knowledge into income level probabilities, or fuzzy income values. This, in turn, will determine the strength of the inferences I draw from those income values.

Clustering algorithms tend to be computationally intensive. If customer categorization functionality is used by checkout/cash register clerks for cross-sells[16], then we need to make the categorization algorithm fairly quick. This could mean many things including using incremental clustering algorithms, or making an approximate and temporary category assignment, to be refined later on in batch mode during off or low traffic hours.

One of the major problems with automatic clustering algorithms is that the resulting clusters/categories are not labeled in a way that is meaningful to a marketer. Within the context of a CA-CEM development framework, we need to provide analysts/CA-CEM application designers with a set of tools that can perform such computations through minimal configuration.

6.2.2 Mining Beliefs and Values

Section 3 showed that *social identity*, and its corollary, *second-order moral values and self-evaluative standards* influence key purchasing decision steps, namely, *goal*

[15] In fact, the property value *is/can be mapped* to the property value distribution of the category.

[16] For example, the cashier just scanned cheese, and they know that the consumer is a DINK, they may suggest them a matching bottle of wine *in the mid-to-expensive* price range.

desire, *behavior desire*, and *implementation intention*. In other words, who we are (or we think we are), and our values/what we believe in, influence both *what* we desire*,* and *how* we go about fulfilling that desire. Chris, our mystery shopper of Section 2.1, is a "fair trader", and as such, was told about the latest fair trade certified arrivals. However, how can we tell if a customer is environmentally conscious, or a fair trader, or someone who cares not only about the price of products they purchase, but also the labor practices of the companies, or jurisdictions that produce them?

There are many strategies for inferring such opinions or beliefs. For example, a shopper whose basket regularly includes items that are certified organic, or simply picked up from the organic/natural foods section, can be safely assumed to be health conscious, and *probably* relatively well. Idem for fair trade coffee. As more and more food store chains carry kosher or halal sections, we can even guess consumers' religions. Thus, basket market analysis techniques can be used if we have a history of purchases, and appropriate product labeling.

Alternatively, we can mine consumers' opinions from the traces they leave on the web. Again, it is difficult to devise a general strategy. It depends on the kind of opinion that is sought, and on the traces that are available. Roughly speaking, the search for opinions can be decomposed into two sub problems: finding out if a customer *cares* about a particular issue (*topic*), and finding out *how* they feel about that issue. To answer the first question (*topic*), we can check issue-specific social media. For example, any article or comment you enter on www.climatecentral.org shows that you care about climate change. The same can be said if the site is listed on your Facebook page, or if you subscribe to #ClimateChange hashtag through your Twitter account. Idem if the phrase 'climate change' occurs in your posts *anywhere* on the web. Generally speaking, the more *diffuse* the topic of the medium, the more sophisticated the opinion mining technique is needed.

Once we know that you care about a particular issue, discovering *how* you feel about it becomes slighter easier, thanks to *sentiment analysis* techniques. For example, to find out whether you believe in human-induced climate change or not, I can look-up your posts on www.climatecentral.org and do a simple polarity analysis[17]. However, there are cases where simpler cues can go a long way towards mining opinions. For example, if you use Monsanto and Genetically Modified Food (GMF) in the same sentence, then your position on the subject of GMF is easy to guess.

From an operational point of view, one of the issues we need to address is whether our CA-CEM application should keep updating users' profiles incrementally, and if so, what should be the triggering event for the incremental re-evaluation. Should it be a new Facebook post, a nightly web crawl for my customers, a purchase, or should I run a batch process on my customer database to compute their environmental opinions?

[17] Knowing the *orientation* of the site alone does not tell us whether a frequent commentator shares the editorial line of the site: a regular commentator may express systematically contrarian opinions.

Despite advances in semantic opinion mining (see e.g. [10], [21], [22], [16]), opinion mining remains an art, because the best (or the only) available strategy to use in each case depends on or greatly benefits from the selection of *sources* from which to mine such opinions (e.g. *available* media specialized in the topic of interest), and the exploitation of information or constraints specific to that opinion (e.g. the *public image* of Monsanto, a *specific company*, in the agricultural food industry). Within the context of our framework, we can only provide a tool-box that opinion miners can use. It is probably something that should be done off-line, e.g. disconnected from the transactional databases/systems.

6.2.3 Assessing Consumer Feedback

The model of the purchasing process shown in Fig. 2 of Section 3 showed that the purchasing process ends with a post-purchase feedback whereby the consumers express, in one way or another, whether they have attained the desired goal that triggered the whole process in the first place. Naturally, companies are very interested in what their customers have to say about the products they purchased, and the quality of the *service* they got during the purchasing process, where the latter has been steadily gaining in importance in building customer satisfaction and loyalty.

There are many ways of getting this feedback. Post-service customer surveys are fairly common nowadays, where customers are given incentives to fill out customer surveys. While such surveys are undoubtedly very valuable because of their granularity and focus, they suffer from a number of problems:

1) A low return rate, despite the incentives
2) They do not address image and perception issues in the general public to the extent that they focus on actual customers.
3) The artificially accommodating/agreeable reviews that all but the most irate consumers end up entering at the insistent requests of customer service personnel.

Thanks to the internet and social media, consumers are now leaving unstructured, unscripted, and uninhibited product and service reviews all over the place. This is a very active research area (see e.g. [10], [28], [21], [22], [16]). Thanks to knowledge base combining lexical, semantic, and sentiment/affective knowledge, recent opinion mining techniques are able to go well beyond *polarity* to explore *emotions* and *intensity,* and even identify *fake* product reviews (see e.g. [21], [16]). As we saw in Section 3, the *anticipated positive emotions* and *anticipated negative emotions* are important deciding factors in goal/behavior desirability. Emotions are not only important indicators of product appreciations, but they are very valuable measures of the success or failures of *service encounters* (see e.g . Cook et al., [12]).

Within the context of our framework, we need to provide resources and tools to CA-CEM applications designers that they can use to mine textual/unstructured product or service reviews left on the seller's portal. Anything beyond the data left on the seller's own web site offers the same opportunities and challenges described in Section 6.2.2.

6.3 Approximate Reasoning to Deal with Incomplete Data

As mentioned earlier, the personalisation embodied in experience management centers around selecting the products and services appropriate for the consumer's needs and desires, and presenting them to the consumer in an appealing way. Both CEM functions rely on the company's knowledge of its products and its customers. The models of Fig. 6 and Fig. 7 show fairly elaborate representations of customers and products, respectively. It is very likely that the company will not have this level of information about its products or customers. Furthermore, for many products and purchase types, much of this information will not be relevant.

We outline four strategies for dealing with incomplete data:

Instance data versus class data: For both products and consumers, we made a distinction between individuals (**Consumer** and **Product**) and classes/categories (**Category**, in both Fig. 6 and Fig. 7, and State in Fig. 6). If we do not have information about an individual, we look it up in the category. Thus, if I do not know that lamb chops go with the 2011 Shiraz, at least I know that meat goes with red wine, and lamb chops being meat, we know that we can recommend red wine, but not the 2011 Shiraz specifically. Thus, we have graceful degradation in the recommendation quality.

1) Category inheritance: If some information is missing from a category, we can look it up in its super-categories. Marketers may know something about transportation needs of DINKs in general, but not specifically about under thirty DINKs.
2) Because of the above, our confidence in the information about a consumer or a product depends on the length of the inference path (from instance to category, and up in the category hierarchy).
3) Use a scoring approach, whereby the different facets contribute positively or negatively to a <product , consumer> match. Consider two customers C_1(needs: detergent, values: environment) and C_2(needs: detergent, values: none), and three detergents D_1(function: detergent, manufacturing.process: contains phosphate), D_2(function: detergent, manufacturing.process: does not contain phosphate), and D_3(function: detergent, manufacturing.process: NA). Assume that a match on function counts for 100, and a match on manufacturing process counts for 20 (or -20). The additive scoring approach would yield the following:
    a. For customer C_1, match(C_1, D_1) = 100 -20 = 80, match(C_1, D_2) = 100 + 20 = 120, and match(C_1, D_3) = 100 + 0 = 100. Thus, D_2 is the best match
    b. For customer C_2, match(C_2, D_1) = match(C_2, D_2) = match(C_2, D_3) = 100 + 0 = 100. In other words, customer C_2 does not care. They may use other criteria to select (e.g. cost).

Note that, when the customer does not care, the company's *own values* could come into play, to position one product higher than another.

7 Instantiating the Framework for a CEM Scenario

Section 4 laid the foundation for CEM by outlining the ***range*** of data and functions that *may be used* to support CEM functionality within the context of a generic e-business system. In this section, we outline steps for selecting and customizing the generic functionality presented in Section 4 to handle a particular CEM scenario.

7.1 Specifying the Relevant Purchase Scenario

In Section 3, we presented the generic structure of a purchasing process as a goal-directed behavior, and identified the various factors that influence or are relevant to the purchasing process. The process described in Fig. 2 may be an overkill for grabbing a carton of milk, or may not reflect the reality of impulsive purchases where one goes window-shopping and ends up making purchases, or walks into an electronic store to buy a USB key and comes out with a laptop.

Thus, the first step in using our CA-CEM framework is to specify the purchasing scenario(s) appropriate for the business by answering the following questions:

1) Which steps of the process in Fig. 2 are relevant? This depends on the type of purchase (see e.g. [25]).
2) Which factors are relevant to those steps? This depends on the type of need (e.g. food versus entertainment), and the type of product (basic versus luxury).
3) When and where do the various steps occur? This is important for many reasons: when, and how much time does the consumer have to make a decision, and what is the best way to position the products to the customer. This also depends on the kind of store/products (e.g. Sharper Image versus Home Depot), and the sale channel used (in-store vs on-line).

A given business can enact several purchase scenarios, depending on the product and the channel, and a CA-CEM system needs to support all the scenarios. Furthermore, the system may *need* to guess which type of purchasing process is in progress, e.g. by using the browsing pattern (physical or on-line) in the store.

The outcome of this step will be a set of processes, each one of which representing a subset of the process of Fig. 2.

7.2 Specifying the Desired Level of Functionality

The model of Fig. 3 shows that all the steps of the purchasing process and all the influence factors are amenable to CEM intervention/functionality. Based on the process specified in section 7.1, and when and where each step takes place, a retail business needs to figure out for each process step:

1) Whether intervention at a particular process step brings value: a grocery store may not need to intervene at the ***Goal desire*** step, as a consumer needs no enticements to buy a carton of milk.

2) Whether it has the means to intervene. This is an issue of opportunity and physical means. If you drop lamb chops in the shopping cart, and I have no way of finding that out (physical limitation) until you reach the cashier, it is probably too late (opportunity) to suggest an appropriate wine.

3) What is the best modality. Intervening at various steps of the purchasing process can easily become intrusive. Should I "push" unsolicited specials to the consumer (the Australian wine special), or wait until s/he asks for wine recommendations, or suggest a wine when s/he pick a product that goes with wine?

The answers to these questions can help a brick-and-mortar retailer decide which and where to put sensors to get the best cost-effective CEM solution. Again, this approach contrasts with "stick an RFID tag to everything that moves, and see what you can do with it".

7.3 Build the Relevant Data Schemas

CEM functionalities rely on our knowledge of purchasing processes, products, and consumers. We saw in Sections 5.1 and 6 the *range* of information we can manipulate about consumers and products to support CEM, and how default information about products and consumers can be gleaned from the categories to which they belong.

The desired level of functionality specified in the previous sections enables us to figure out what/how much data about consumers and products we need. We now have to specify the structure of the data, and populate it for categories (for consumers and products) and states (for consumers).

The models shown in Fig. 6 and Fig. 7 represent ontologies, i.e. *metamodels* of the actual data models needed by our CA-CEM functionality. In this step, we need to *instantiate* those ontologies to specify the actual data schemas needed by our system. For example, the model in Fig. 6 shows that a **Consumer** can have a bunch of **Property**(ies), and that each **Category** is characterized by a range of *values* for those properties. In this step, we need to specify the properties that we are interested in. For a **Consumer**, such properties include your typical CRM data such as age (or birth date/birth year), address, payment methods, profession, perhaps income bracket, marital status, and transaction history. To this, we add significant relationships, social identity(ies), and values. Consumer likes for specific products or product categories can be represented by ternary relationships.

Most of these **Property**(ies) may be used to characterize **Category**(ies) (e.g. DINKs, generation Y) and **State**s (married people, various family lifecycle stages), but some are exclusive to instances (e.g. name, transaction history). When a property applies to a **Category** (or **State**), it is represented by a *range* of values. For example, DINKs, may have an income bracket of [100 k, ∞[, like sports cars, while generation Y have an age range of [18 - 35]. Entering value ranges for these properties is crucial, since learning of a consumer's membership to a particular **Category** or **State** can help us infer a lot of information about their needs and desires. ***Those values embody the marketing knowledge that the business has about its customers or potential customers***.

A similar process needs to take place for the product data. We need to first specify which properties are worth representing. Then, we populate information about product *categories* (see the model of Fig. 6). That is where we specify that red wines go with meats, that big fish tend to have high mercury contents, or that Absurdistan has poor labor laws and lax environmental regulations.

7.4 Populate Instance Data

The product instance data is a one-time deal for the products sold by the business. Idem for promotional material (see Section 5.3).The consumer instance data, however, is a work in progress. We show below the lifecycle of a consumer record:

1) Creation. Businesses learn about new or potential customers in many ways: a) when customers go through the cash register the first time around, b) when they access anonymously the web site of the company - in which case they may be known by IP address, c) when they create accounts on the company's web site, or post comments with an ID from another portal, and d) by purchasing consumer lists. Depending on the channel, the business will have more or less information.
2) Category/state assignment. This is the process through which a consumer is identified with a category (a generation Y), or with a state/lifecycle state (young married couple in early childbearing phase). This information can be entered explicitly (filling out a questionnaire when applying for a store credit card), or inferred. For example, someone who buys baby formula or diapers regularly is assumed with a high-level of certainty to be in the early childbearing phase. This assignment allows to infer other data.
3) Property value updates. Property values, or our confidence in their values, may be regularly updated through explicit entries or accumulation of evidence. For example, if someone has bought baby formula three times in the past two weeks, s/he is *probably* in childbearing stage (confidence level of 70%) - although s/he could just be staying with someone who is. If s/he keeps at it week after week, our confidence level goes up. Generally speaking, consumer classification rules need to kick-in each time information about the customer is updated.
4) State change. Recall, from the model in Fig. 6, that state changes can be triggered by two kinds of events, timer events, and property-value change events. A preschool toddler will remain in this state for a maximum of 5 years. A teenager who changes home address, from his/her parents' home to an outside address has probably moved into young adulthood. Thus, *each time a consumer property value is updated, we run statement assignment rules*.

We mention the possibility that category data be updated from consumer instance data. For example, our marketing department came up with a category (market segment) based on age an income. We later realize that all of the consumers of that category live in a particular area. The zip code could become a *characteristic feature* of this category, i.e. its value used as a default for consumers of the category, but it would not be used as a criterion for membership (i.e. it is not *definitional*).

8 Discussion

CEM aims at personalizing a customer's interactions with a company around the customer's needs and desires [29]. Researchers and e-business visionaries have been fantasizing about the kind of personalization afforded by ubiquitous computing (e.g. the Internet of Things). While all the technical ingredients for such fantastic scenarios exist today (sensor technology, middleware, mobile computing, social network analysis, data mining techniques), there are no guidelines to help either e-business software vendors, to figure out which functionalities to provide in their software to support CEM functionalities, or e-business software clients to figure out which CEM functionalities to implement, let alone *how* to implement them, given the type of products they sell, the channels through which they sell those products, and their software and hardware capabilities.

Consumers interact with companies to acquire products or services that satisfy their needs and desires: a *purchase* (et al., 2007b). To personalize such interactions, we need to understand them. Toward this end, we relied on studies of *consumer behavior* from marketing and social psychology. Such studies identified the steps undertaken by a consumer in a purchase process, and the factors that influence their decision making in those steps (Section 3). For the purposes of CEM, these steps are potential interaction opportunities between the company and the consumer, and the influence factors represent relevant data that the company could use to personalize those interaction (Section 4.1). This enabled us to identify an *experience management pattern* that could be applied to various steps of the purchasing process. However, we showed that a blind application of the pattern to all the steps of the purchasing process described in Section 3 would not be appropriate (Section 4.2). We argued that the application of this pattern is best considered within the context of *service design* (Section 4.3), and showed an example of the analyzes a service designer could make to instantiate this pattern (Section 4.4). While Section 4 explored the design vocabulary for *customer experiences*, Section 5 presented the kind of representation that is required for consumers and products to be able to match products to consumer's needs/desires, preferences, biases, and attitudes. Section 6 discussed how to deal with incomplete customer or product data. Given such an infrastructure, we proposed the first elements of a methodology enabling e-business software users to specify the CEM functionalities they wish to support (Sections 7.1 and 7.2), set-up the data infrastructure (Section 7.3), and populate the consumer data (Section 7.4).

This work is at an early stage. In this paper, we focused on the foundations for *a methodology*, *for designing and implementing CA-CEM functionalities*. The proposed approach gives us a frame of reference to study the multitude of issues raised by CEM functionalities within the context of e-commerce software. We are currently testing the theory on a specific purchasing scenario provided by an e-commerce suite vendor. It will enable us to validate the theory, and implement a first cut of some of the software artifacts of our framework.

References

1. Ajzen, I.: From intentions to actions: A theory of planned behavior. In: Action Control. pp. 11–39 Springer (1985).
2. Ajzen, I., Fishbein, M.: Understanding attitudes and predicting social behavior. Prentice-Hall (1980).
3. Ajzen, I., Madden, T.: Prediction of goal-directed behavior: Attitudes, intentions, and perceived behavioral control. J. Exp. Soc. Psychol. (1986).
4. Armitage, C., Conner, M.: The theory of planned behaviour: Assessment of predictive validity and'perceived control. Br. J. Soc. Psychol. (1999).
5. Bagozzi, R. et al.: Network analyses of hierarchical cognitive connections between concrete and abstract goals: An application to consumer recycling attitudes and behaviors. (1996).
6. Bagozzi, R.: The self-regulation of attitudes, intentions, and behavior. Soc. Psychol. Q. 55, 2, 178–204 (1992).
7. Bagozzi, R. et al.: The social psychology of consumer behaviour. Open University Press (2007).
8. Bagozzi, R., Dholakia, U.: Antecedents and purchase consequences of customer participation in small group brand communities. Int. J. Res. Mark. (2006).
9. Bagozzi, R., Warshaw, P.: Trying to consume. J. Consum. Res. 17, 2, 127–140 (1990).
10. Cambria, E. et al.: SenticNet 2: A Semantic and Affective Resource for Opinion Mining and Sentiment Analysis. FLAIRS Conf. (2012).
11. Constantine, L.: Human Activity Modeling: Toward A Pragmatic Integration of Activity Theory and Usage-Centered Design. Human-Centered Softw. Eng. 27–51 (2009).
12. Cook, L.S. et al.: Human issues in service design. J. Oper. Manag. 20, 2, 159–174 (2002).
13. Fazio, R.: How do attitudes guide behavior. In: Handbook of motivation and cognition: Foundations of Social Behavior. The Guildford Press (1986).
14. Fishbein, M., Ajzen, I.: Belief, attitude, intention and behavior: An introduction to theory and research. (1975).
15. Hill, A. V. et al.: Research opportunities in service process design. J. Oper. Manag. 20, 2, 189–202 (2002).

16. Liu, B.: Sentiment Analysis Mining Opinions, Sentiments, and Emotions. Cambridge University Press (2015).

17. Mylopoulos, J. et al.: From Object-oriented to Goal-oriented Requirements Analysis. Commun. ACM. 42, 1, 31–37 (1999).

18. Ostrom, A.L. et al.: Moving Forward and Making a Difference: Research Priorities for the Science of Service. J. Serv. Res. 13, 1, 4–36 (2010).

19. Palmer, A.: Customer experience management: a critical review of an emerging idea. J. Serv. Mark. (2010).

20. Patricio, L. et al.: Multilevel Service Design: From Customer Value Constellation to Service Experience Blueprinting. J. Serv. Res. 14, 2, 180–200 (2011).

21. Poria, S. et al.: Enhanced SenticNet with affective labels for concept-based opinion mining. IEEE Intell. …. (2013).

22. Poria, S. et al.: Towards an intelligent framework for multimodal affective data analysis. Neural Networks. (2015).

23. Sarwar, B. et al.: Analysis of recommendation algorithms for e-commerce. In: 2nd ACM conference on Electronic commerce. pp. 158–167 , Minneapolis, MN, USA (2000).

24. Schwartz, S.: Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. Adv. Exp. Soc. Psychol. 25, 1, 1–65 (1992).

25. Solomon, M. et al.: Consumer behavior: buying, having, and being. Pearson (2009).

26. Teixeira, J. et al.: Customer experience modeling: from customer experience to service design. J. Serv. Manag. 23, 3, 362–376 (2012).

27. Verhoef, P.C. et al.: Customer Experience Creation: Determinants, Dynamics and Management Strategies. J. Retail. 85, 1, 31–41 (2009).

28. Vinodhini, G., Chandrasekaran, R.: Sentiment analysis and opinion mining: a survey. nternational J. Adv. Res. Comput. Sci. Softw. Eng. 2, 6, (2012).

29. Weijters, B. et al.: Determinants and Outcomes of Customers' Use of Self-Service Technology in a Retail Setting. J. Serv. Res. 10, 1, 3–21 (2007).

30. Wells, W., Gubar, G.: Life cycle concept in marketing research. J. Mark. Res. (1966).