IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received September 8, 2020, accepted September 28, 2020. Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

*Digital Object Identifier 10.1109/ACCESS.2020.3028777*

# Generating Recommendations From Multiple Data Sources: A Methodological Framework for System Design and Its Application

**FEDERICA CENA[1], AMON RAPP[1], (Associate Member, IEEE), CATALDO MUSTO[2], AND GIOVANNI SEMERARO[2]**
[1]Department of Computer Science, University of Turin, 10124 Turin, Italy
[2]Department of Computer Science, University of Bari Aldo Moro, 70121 Bari, Italy
Corresponding author: Cataldo Musto (cataldo.musto@uniba.it)

**ABSTRACT** Recommender systems (RSs) are systems that produce individualized recommendations as output or drive the user in a personalized way to interesting or useful objects in a space of possible options. Recently, RSs emerged as an effective support for decision making. However, when people make decisions, they usually take into account different and often conflicting information such as preferences, long-term goals, context, and their current condition. This complexity is often ignored by RSs. In order to provide an effective decision-making support, a RS should be ''holistic'', i.e., it should rely on a complete representation of the user, encoding heterogeneous user features (such as personal interests, psychological traits, health data, social connections) that may come from multiple data sources. However, to obtain such *holistic recommendations* some steps are necessary: first, we need to identify the *goal* of the decision-making process; then, we have to exploit common-sense and domain knowledge to provide the user with the most suitable suggestions that best fit the recommendation scenario. In this article, we present a methodological framework that can drive researchers and developers during the design process of this kind of ''holistic'' RS. We also provide evidence of the framework validity by presenting the design process and the evaluation of a food RS based on holistic principles.

**INDEX TERMS** Theoretical framework, design methodology, recommender systems.

## I. INTRODUCTION

Recommender systems (RSs) are *"systems that produce individualized recommendations as output or have the effect of guiding the user in a personalized way to interesting or useful objects in a large space of possible options"* [9]. Recently they emerged as an effective means to support users in making decision [39]. Recommendations are usually related to simple low-risk decision-making processes, like products to buy, music to listen to, or movie to watch; however, these systems can be also exploited to make more complex high-risk decisions, in order to *e.g.,* be healthy, save money, and optimize time [23].

When people make choices, they commonly focus on different aspects, pertaining to both the domain of the choice and other (even apparently unrelated) domains [22]. For example, when a person chooses *what to eat*, she may first consider her preferences in the food domain (*"I'd like a cake"*), and then other aspects, such as her current mood (*"I feel sad"*), her past experiences (*"When I feel sad usually sweets help me"*), constraints related to her health condition (*"But I have to control my glucose intake"*), long-term goals (*"I'd like to lose weight for the summer"*), contextual factors (*"there is no good pastry shops in the nearby and I have no time to search for other places"*) and social acceptability (*"maybe my boyfriend could be disappointed if I don't eat with him"*). Much of this information may be related (*e.g.,* mood and sugar intakes), or in conflict (*e.g.,* the momentary desire of eating sweets and its long-term consequences), and thus some kind of cost-benefit evaluation among these aspects need to be performed by the individual [37], even when decisions are unconscious and thus she is not aware of the underlying process [14], [25].

In this perspective, in order to be an effective support for decision making, a RS should act similarly to the *way* people make decisions: i.e., it should manage a greater amount of

The associate editor coordinating the review of this manuscript and approving it for publication was Siddhartha Bhattacharyya.

information about the user, as well as take into account how different aspects of her life impact on the decision-making processes.

Going back to the aforementioned example regarding food recommendation, a recommendation algorithm for this domain should take into account *at the same time* the user's food preferences, her health data (e.g., she is overweight with a heart condition), her current mood (e.g., she is sad, and sweets might turn her day to the better), how much physical activity she has done today (e.g., she ran for 1 hour), and her general goals (she wants to lose weight).

However, current recommender systems do not act in such a way yet [12], because they typically consider only the user's preferences in terms of ratings [39] in a single target domain: they do not pay attention to potentially relevant aspects of the user's life that do not closely pertain to the recommendation domain (*e.g.,* the user's mood when recommending food), excluding them from the recommendation process. This problem is connected with the fact that despite the heterogeneity of the data currently available, most RSs exploit a single source of information about the user. In other words, current RSs miss to consider the complexity of the user's decision making process, which takes into account multiple, and even conflicting goals, constraints, contextual factors, etc. proposing an oversimplified way of giving advice. However, this way of providing suggestions may be particularly relevant in *complex domains*, such as food or health recommendations, where a misleading advice could not only be annoying but also harmful [49].

To address this point, a RS should encode and reason over a larger set of evidences and information about the user, thus considering the individual as a real person in all her complexity, rather than modeling the user as a simple list of ratings. This implies to gain as much knowledge as possible about the user. To this end, the recent advancements in Internet of Things technology and the spreading of mobile devices and social web sites allow to collect a huge amount of personal data, supporting, in principle, the creation of a sort of total, *holistic* representation of the individual [13]. Some small advances in this direction in RS research have been provided by contributions in the area of cross-domain recommenders [10], which exploit the user's preferences in one domain to infer her preferences in another domain. Nonetheless, such recommenders do not completely implement a *holistic recommendation process* since they merely use the user's preferences in one specific domain to fill the knowledge gap in another domain [24], [34]. Therefore, they do not exploit *together* different kinds of information pertaining to the user's life, which may potentially affect the recommendation process.

By following these insights, in this article we introduce the concept of *Holistic Recommendation* (RecHol in the following), that is to say, a suggestion that is obtained by considering a comprehensive representation of the user, as well as of the recommendation task itself. To model the user we exploited what we called *holistic user profiles*, extensive user models that encode in a *single user profile* information about the user's interests, affects, psychological states, physical states, social connections and behaviors [11], [29]. Such a representation is obtained by gathering rough data from diverse data sources, such as social networks, smartphones, wearable devices, and environment sensors, and by reasoning over these data in order to populate the different facets of the profile.

However, gathering a huge set of user data and encoding them in a *holistic user profile* are not enough to obtain *holistic recommendations*. To fully realize this paradigm, a recommendation strategy needs to be defined. A holistic recommender system shall first identify the *goal* of the decision-making process, that is to say, the user's *sphere* of life that should be prioritized through the recommendation process (*e.g.,* to maximize healthiness, or amusement, etc.). Given such a goal and the information encoded in a *holistic user profile*, common-sense and domain knowledge need to be collected and used to identify the most suitable suggestions that best fit the recommendation scenario (*e.g.,* if a food recommender system has to maximize the healthiness of the suggestions, the recommendation algorithm should give more importance to the user's health condition rather than to her food preferences).

Holistic recommendations draw inspiration from different related areas, such as context-aware recommendations [3], cross-domain recommendations [10] and knowledge-based recommendations [8]. In particular: *(i)* the contextual situations of the user and of the recommendation task are taken into account, as in context-aware recommender systems, *(ii)* a huge number of (cross-domain) features about the user are encoded in the profile; *(iii)* rules and reasoning strategies are used to adapt and instantiate the user profile on the ground of the requirements and the goal of the recommendation setting, as in knowledge-based recommender systems. The combination of these elements can provide a recommendation framework with *the necessary information* for enhancing the decision-making process, making it closer to the way human beings make decisions.

However, significant differences between holistic recommendations and these approaches exist. In particular:

- Differently from original CARS, which mainly consider the influence of context on user's preferences expressed as ratings, we try to take into account the influence of the contextual factors and the features that are encoded in the profile of the user (for example, how physical activity impacts mood or physiological states as sleep, etc.) and to reason over this information to further adapt the user model on the ground of these evidences.
- Differently from cross-domain recommenders that mainly focus on user preferences in different domains and do not consider other user characteristics as the current mood of the user, her goals, her behaviors and so on, we propose a recommendation strategy which is based on the acquisition and the processing at the

same time of several heterogeneous data points about the person.
- Differently from classical knowledge-based recommender systems, which are focused on a particular domain or scenario, we encode general common-sense knowledge that also considers the impact and the role of different spheres of life (e.g. health, free time, etc.). This is a completely new different research direction.

To summarize, exploiting a variety of personal data has several advantages for recommender system design: i) it allows users to make more informed decisions, since in this way the recommender can act as a decision support system similar to the way people make decision, considering different aspects of the user's life and their possible conflict and interactions; ii) it allows RS designers to tackle complex domains, like food and health, which commonly have wide impacts on many aspects of the user's life, where considering only user preferences on a single domain is not enough to generate useful recommendations.

The contribution of this article is twofold. First, we outline a methodological framework that can guide researchers and practitioners in the process of designing RSs based on holistic user profiles. Second, we present an application of this framework by describing a food recommender system that has been designed on the basis of the holistic principles: a preliminary evaluation validates the framework proving its effectiveness.

The article is structured as follows: Section II provides the background, introducing the related work, while Section III outlines the notion of holistic user profile. Next, Section IV provides the designers with a reference model to build a recommender system based on holistic user profiles. Section 4 presents an application of the model, showing how to implement a *food recommender system* based on the methodology we propose and validating the approach through an experimental study. Finally, Section 6 concludes the work discussing issues opened by our framework.

## II. STATE OF THE ART

In the following, we situate our approach within related research areas, presenting relevant related work.

### A. CROSS-DOMAIN RECOMMENDER SYSTEMS

RecHol exploits information about the users coming from different and heterogeneous sources. From this perspective, holistic recommendations have a relation with cross-domain recommendations, where the preferences of a user in a domain are used to infer her preferences in another domain. For example, a cross-domain recommender can use the user's preferences on music to infer her preferences in the book domain [24]. However, some important differences between the two approaches exist. Differently from cross-domain recommenders that mainly focus on the user's preferences and do not consider other user's characteristics, like her current mood, goals, and behaviors, RecHols propose a recommendation strategy that is based on the acquisition and processing of several heterogeneous data points regarding the person. Moreover, the goal of the two approaches is different as well: cross-domain recommenders aim to reuse existing knowledge (*e.g.,* ratings) to solve the cold start problem at the beginning of the interaction with the system when the amount of users' ratings is not sufficient to generate the recommendations ( [24], [34]). Differently, RecHols aim to acquire as much as possible knowledge about the user with the goal of providing a comprehensive modeling of the different aspects that characterize the person.

### B. KNOWLEDGE-BASED RECOMMENDER SYSTEMS

Knowledge-based Recommender Systems (KBRSs) [8], [50] typically exploit knowledge about users and items to pursue a knowledge-based approach in order to generate a recommendation. Such knowledge (both common-sense and domain-specific) is typically encoded as *rules*, and it is used to filter out, or to give priorities to, particular classes of items.

These systems have the main advantage of not being affected by the *cold-start* problem, which is common in both content-based and collaborative recommendation methods. Even though KBRSs did not receive particular attention from scholars, because the information and knowledge necessary to execute them are particularly hard to collect, encode and maintain, recent work showed their effectiveness in both health [19] and academia [43] domains.

Nonetheless, this class of systems is particularly interesting for our goals. Indeed, we have used the principles of KBRS (in particular, reasoning over the users' profile by applying specific rules) to implement a goal-oriented recommendation methodology, which gives more importance to a particular sphere of the user's life with respect to the other ones.

### C. CONTEXT-AWARE RECOMMENDER SYSTEMS

The injection of contextual data and contextual factors in the recommendation process can significantly improve the overall accuracy of the suggestions. As shown by several works [1], [4], [21], [38], Context-Aware Recommenders (CARs) usually outperform merely content or collaborative-based recommendations.

RecHol is also based on factors such as the current situation of the target user or the current context ((*e.g.,* time, place) in which the recommendation is delivered. However, differently from CARs, which mainly consider the influence of context on the user's preferences expressed as ratings ([6], [45]), RecHol takes into account the influence of the contextual factors and the features that are encoded in the user profile (for example, how physical activity impacts mood or physiological states, such as sleep, etc.) to reason over this information and further adapt the user model on the ground of these evidences.

By contrast, most of current CARs do not implement a holistic user model nor a holistic recommendation strategy [52]. For instance, InCarMusic [5] is a Context-Aware music recommender offering music recommendations to the passengers of a car. The system is composed of a tool for acquiring context relevance subjective judgments, in order to quantitatively estimate the dependency of the user preferences on a candidate set of contextual factors (driving style, road type, landscape, sleepiness, traffic conditions, mood, weather, day time). Then, the users rate items under certain contextual conditions and, based on these, a predictive model is built. This has the goal of predicting the user's ratings for items under target contextual situations, extending a classical matrix factorization.

Reference [7] presents a social recommender system in the tourism domain able to identify user preferences and information needs, suggesting personalized recommendations related to POIs in the surroundings of the user's current location. The system exploits information from social networking, user reviews, and local search Web sites; it uses a neural network to match context-user features with POIs. It employs a richer contextual description that besides traditional physical and environmental factors focuses on the classification of basic human activities or scenarios.

Reference [36] presents a recommender of mobile applications that traces context information that can be gathered implicitly based on sensors equipped on devices (e.g. location, acceleration, and noise level), geographic information systems (e.g., detecting if somebody is near a lake or inside a building), meteorological services (e.g., temperature) or social networking sites (e.g., detecting nearby friends). The logger also records explicit context information that is announced by a user, e.g., a description of her current situation. Secondly, it also traces the application usage, i.e., when an application is installed, used or deleted.

Reference [51] presented a context-aware point-of-interest (POI) recommender system by combining community-based knowledge with association rule mining to alleviate the cold start problem. This is basically a hybrid model that combines four key factors: 1) community created knowledge, 2) ontologies, 3) association rule mining, and 4) an innovative scoring function based on probability metrics.

Reference [57] introduced a novel context- aware group recommendation for the point-of-interest generation. The novelty of the work lies in considering the location and the intra-group influence when making group decisions. By taking into account the importance of location in POI recommendations, the authors employed distance-based pre-filtering and distance-based ranking adjustment to improve recommendation satisfaction.

### D. HEALTH RECOMMENDERS

Health domain has a high level of complexity and involves different aspects of a person [18]. In this perspective, health may require a holistic approach [15]. The need to know different heterogeneous information about the user and her life, as well as to manage an enormous quantity of data and their relations [49], [55], makes health recommenders an ideal application of a holistic-based recommender.

However, most of health recommenders currently available do not implement a holistic user model nor a holistic recommendation strategy. Usually, they aim to deliver trustworthy relevant information to end-users, improving their safety [40], or providing lifestyle change recommendations: e.g., by suggesting how to adopt more healthy habits [16], by improving users' eating behavior [33], [41], [46], [47], by recommending physical activities while taking into account the user's health state [17], or by preventing not healthy behavior, such as alcohol consumption [56].

The BlueMedics system [40], for instance, suggests relevant information about interactions between different medicines, in order to avoid health risks, as well as detailed explanations that help the patient better understand the link between the recommended content to her profile. To do so, the system maintains a user model called PHR (Personal Health Records) for each patient in the system. The PHR is continuously automatically updated from official medical data sources and contains: up-to-date information about patients and their associations with various medical entities, such as medications, allergies, immunizations, various clinical conditions (e.g., pregnancy, diabetics, chronic diseases), treatment plans, etc.

Toledo *et al.* [47] present a food recommender which provides tailored food intake advice according to users' physical and physiological data, as well as other personal information. The paper presents a general framework for daily meal plan recommendations, incorporating the simultaneous management of nutritional-aware and preference-aware information.

HyperRecSysPA [17] is a system for recommending physical activities for hypertensive patients based on a hypertensive user profile model. The model is composed of 32 elements divided into three groups, which were used in the modeling of user profiles within the system for generating HyperRecSysPA recommendations.

Reference [16] presents a lifestyle change recommender which exploits a new method, the Intrapersonal Retrospective Recommendation, for generating recommendations that use only the user's personal history. The key idea behind this approach is that recommendations can be based on what behaviors worked and did not work for the individual in the past. They consider the count of the number of times the user performed the behavior, the total amount of food or time exercising, and the total calories burned or consumed to be part of the implicit rating of an item. They provide suggestions in relation to food and exercise.

u-BabSang [33] is a context-aware food recommendation system for well-being care applications. It provides individualized food recommendation lists at the dining table, and it is based on dietary advice. The system creates a user's profile with physiological signals and environmental information around the dining table in real time. In other words, u-BabSang recommends appropriate food for each individual's health in real time.

The GlycoRec system [46] aims to support diabetes patients in managing their disease. It supports decisions and gives individualized recommendations based on the patient's behavior, physiology and treatment history. Individualized advice includes i) estimation of nutritional characteristics

**TABLE 1. Comparison among examples of state of the art recommender systems with respect to RecHol features.**

| system | domain | HUM features | context features |
|---|---|---|---|
| [40] | health | PHR (personal health records) | personal health context |
| [16] | behavior change | medications, allergies, immunizations, clinical conditions (e.g., pregnancy, chronic diseases), food habits, sport habits | day, time, place |
| [33] | food | gender, clinical history, body constitution, physiological signals (e.g., heart rate, skin temperature, sweat on palms) | environmental information (e.g., intensity of illumination, noise level, temperature) |
| [46] | food | age, weight, co-morbidity, glucose, stress level, | place, activity level, production information |
| [56] | behavior change | consequences of drinking, motivation to change, dependence to alcohol, risk factors, frequency of drinking, demographic features, effective features | social interactions, location, time of drinking |
| [17] | activity | risk profile (blood pressure, disease, weight, age) | intensity, duration |
| [47] | food | preferences, physical data | nutritional facts |
| [5] | music | sleepiness, mood, driving style, preferences | road type, landscape, traffic conditions, weather, day time |
| [7] | POIs | preferences, user activities (working, travelling) | physical context, environmental context |
| [36] | apps | preferences | location, acceleration, noise level, weather, friends, application usage |
| [57] | POIs | user location, personal user influencek | season, time, weather, number of people, price |
| [51] | POIs | preferences, places visited | day, time period, weather |

such as carbohydrate content and glycemic index of meals, ii) recommendations about insulin application based on glucose level, activity and food intakes, iii) warnings if blood glucose levels are at risk of leaving the target range.

Reference [56] propose a user model specific for alcoholics which exploits the target-behavior related features. In fact, they proposed a user model composed of eight different groups of features: consequences of drinking (from the subjective and objective points of view), motivation to change, dependence on alcohol, risk factors, frequency of drinking, demographic features, affective features, and contextual features (e.g., social interactions, location, and time of drinking).

Table 1 compares some examples of state-of-the-art recommenders with respect to RecHol features. Even though more and more recommenders are starting to increase the kinds of user and context features collected, no one uses them to provide holistic recommendations.

To summarize, the novelties of our approach are the following ones: i) we propose a recommendation strategy based on the acquisition and the processing of different heterogeneous data about the person, ii) we consider how the contextual factors and the user model's features influence each others (such as how physical activity impacts on mood or sleep, etc.); iii) currently, none of the existing health and food recommenders implement our approach.

## III. FUNDAMENTALS OF HOLISTIC USER MODELING

Holistic User Models (HUMs) represent the main pillar of our framework, which aims to provide users with holistic recommendations. A HUM is intended as *a digital representation of the person merging heterogeneous footprints spread by the users in their on-line (purchases, generated content, social connections, etc.), and real-world behavior (GPS data, daily activities, food)* . This is a mandatory requirement for developing holistic RSs: in absence of such a representation it is not possible to trigger holistic recommendation strategies, since the availability of a rich and comprehensive representation of the person is necessary to start the RecHol process.

A HUM can be seen as an extension of Lifelong User Model, since it aims to build a complete depiction of the person putting together several aspects of her life in a unique and comprehensive representation [53]. The main difference between HUMs and Lifelong User Models lies in the generality and in the abstraction of the representation, since HUM aims to build a domain-independent model that can be (potentially) used in several recommendation scenarios, while Lifelong User Models are targeted to the learning domain. This Holistic User Model, albeit similar, at a first sight, to Google profiles,[1] contains a wider range of information, such those related to emotions and mood, and, in principle, aims to model all the different aspects of the user's life (i.e., all her life spheres, as we will see in the next Sections). It also accounts for the relations among these aspects. This said, a main contribution of this article is related to explain how to effectively exploit a HUM for recommendation purposes, by pointing out all those passages that need to be considered for designing holistic recommender systems.

More concretely, HUMs are inspired by the model proposed by Cena *et al.* in [11], where real-word data coming from environmental and wearable sensors are used to model the user. With respect of such conceptualization, HUM also includes information coming from the web (social connections, interactions, textual messages like posts, comments and tags) in order to create a more complete picture of the user.

In particular, we can state that a HUM is split in eight different facets, each of which aims to describe a different aspect of the life of the individual: *demographics, interests, knowledge and skills, affects, psychological traits, behaviors, social connections, physical states*.

As shown in [30], a HUM is built in two steps: *(i)* rough data about the person are gathered from social networks, smartphones and wearable devices; *(ii)* these data are processed through natural language processing and machine learning techniques and are used to populate the facets that

[1] https://policies.google.com/privacy?hl=en#infocollect

**TABLE 2.** Features in the holistic user model (adapted from [11]).

| User data | Short Term | Long Term |
|---|---|---|
| demographics | address, job, marital status, etc. | name, date of birth, etc. |
| knowledge and skills | particular skills | general capabilities |
| preferences | opinions<br>short term interests | believes<br>long term interests |
| affects | emotions, moods | emotional disorder |
| psychological traits | cognitive states<br>(level of attention,etc.) | cognitive skills<br>(orientation, etc)<br>personality traits |
| behavior | tasks, activities | habits |
| physical states | physiological parameters<br>(blood pressure, etc.) | chronic diseases |
| social connections | encounters | user's social network |

compose a HUM. For the sake of simplicity, we do not discuss here the computational tasks required to acquire and process personal data and to map them to the *facets* of a holistic user profile. To better understand how a HUM is built, we can refer to [30]. To our aims, it is sufficient to assume that a HUM of a user has been built and is available. Table 2 summarizes the groups of features (split into *short-term* and *long-term* features) that are typically encoded in a holistic user profile.

It is worth to note that the above presented list includes *all* the facets that can be ideally encoded in the HUMs. Clearly, a holistic recommendation strategy can be triggered even if just a *subset* of the input is available.

## IV. DESIGNING A HOLISTIC RECOMMENDER SYSTEM

HUMs can provide designers and practitioners with a huge amount of personal data that can be useful to build personalized systems and applications. Once a HUM is defined, several tasks are required to generate holistic recommendations. In particular, some of these tasks are *modeling* tasks that regard the definition of the *goal* of the process and the identification of the *contextual factors* that influence the recommendations. Conversely, other tasks trigger some computations, such as running algorithms or acquiring data.

To this end, in this section we introduce a *methodological framework* supposed to act as a *reference model* for a designer who wants to implement a holistic recommender system that takes the best out of the personal data available. This model allows designers to take in considerations all the steps needed for creating a RS that exploits information coming from heterogeneous sources. In so doing, the model points out those opportunities and criticalities that may arise when recommendations are designed for addressing multiple, intertwined, life domains, which may have goals or priorities that conflict each other (e.g., staying healthy and finding satisfaction in eating) or should consider specific "constraints" coming from particular situations (e.g., the impossibility of following the diet when the user is at a restaurant with her friends").

The methodological model we propose is composed of five main steps that outline the pipeline to generate *holistic recommendations*:

1) **Scope Definition:** definition of the *scope* of the recommender, i.e., domain, object, goal, context;

**TABLE 3.** Designers checklist.

| Dimensions | Questions |
|---|---|
| **Scope definition** | What is the *domain* of the recommendation task?<br>What is the object of the recommendation?<br>What is the *goal* of the recommendation task?<br>Which *contextual dimensions* impact on the process? |
| **Spheres Modeling** | Which *spheres* can we define?<br>What is the *main sphere/secondary sphere*? |
| **User Modeling** | Which *facets* of a holistic user model are needed?<br>Is there any *domain-specific feature* or requirement? |
| **Reasoning** | What kind of *knowledge* do we acquire from the main sphere?<br>What kind of *knowledge* do we acquire by analyzing user behavior?<br>How does such *knowledge* impact on the representation of the user?<br>How does such *knowledge* impact on the recommendation process? |
| **Recommending** | Which *algorithm* can we run to obtain recommendations?<br>What is the *best style* to present recommendations?<br>When is it *better* to provide the recommendation? |

2) **Sphere Modeling:** identification of the *spheres* that describe the life of the individual and definition of the *priorities* among the spheres, based on the recommendation goals;
3) **User Modeling:** acquisition of a (holistic) user model and identification of the current contextual information;
4) **Reasoning:** codification of *common-sense* and *domain specific knowledge* and reasoning over the user model and the recommendations;
5) **Recommendation:** generation of the recommendations, on the ground of the current contextual situation, the characteristics of the user and the reasoning process.

Given such a high-level pipeline, we introduce here a *checklist* for a designer who intends to build a RecHol. As shown in Table 3, such a checklist is built to elicit the right questions that we have to ask ourselves at each step of the design and implementation processes.

In the following, we thoroughly discuss the checklist as well as the pipeline to build a holistic recommender.

### A. SCOPE DEFINITION

The first and mandatory step to design a RecHol obviously regards the definition of the *scope* of the recommendation, i.e., the domain of the recommendation (*e.g.,* travel, food, sport, etc), the nature of the *objects* to be recommended (*e.g.,* book, movies, meals) and the final goal (e.g, to improve user daily activity).

Moreover, another basic aspect to be designed pertains to the definition of the *context*. Indeed, context plays a key role in every decision-making process, and recommender systems are no exception. Accordingly, the designer has to define: *(i)* the contextual factors, that is to say, the factors that influence the recommendation process (*e.g.,* companionship, mood, kind of dinner, etc.); *(ii)* the contextual dimensions, namely, the values that each factor can have (*e.g.,* companionship=friends, companionship=family, etc.).

It is worth to note that RecHols do not provide multi-domain recommendations. Instead, they are designed for a single-domain recommendation scenario.

### B. SPHERES MODELING

A distinguishing trait of a holistic recommendation strategy is represented by the definition of the *spheres*. The *spheres* can be seen as a set of ''shells'' representing the life of a person.

Generally speaking, the spheres model fundamental domains of the user's existence and bring along some *common-sense* knowledge which is exploited to better drive the recommendation process. In other words, each sphere can be seen as a view on a particular aspect of the user's life. On the basis of literature about sociology of life course (e.g., [26]), we preliminary identified four spheres, namely the *''health''* sphere, which models all those aspects related to the user's physical health and well-being, the *''free time''* sphere, which models all the activities she carries out in her spare time, the *''work''* sphere, which concerns the aspects related to the user's job, the *''family''* sphere, which relates to the user's intimate dimension.[2] It is worth to note that spheres are different from simple contextual factors, since a context is traditionally meant as a *momentary situation* [2]. Rather, we consider the spheres as ''macro-contexts'' of a person's life, as a set of elements and factors that can influence the whole decision-making process.

Then, given the *goal* of the recommendation task, the designer shall use this information to define a *priority* mechanism among the spheres. In particular, the designer should decide the primary sphere and the secondary ones, that is to say, the main sphere representing the field of application of the recommendation (i.e., its goal), while the secondary spheres addressing those aspects that have an impact on the decision-making process. Two examples can be seen in Table 4:

In *Example A*, the user is interacting with a food recommender system that will suggest her a recipe to prepare. In case the designer explicitly sets the *health sphere* as the primary sphere, the recommendation strategy will be influenced by the health-related features of the profile of the target user, as her weight, her medical condition, allergies and so on. Moreover, since the free-time sphere is set as the secondary sphere, elements such as food preferences, previously visited restaurants, and food habits will be also taken into account to provide the user with a suitable recommendation. In this case, we can assume that by prioritizing the *health* sphere the RecHol will suggest a healthy recipe. However, by setting a different main sphere (*e.g.,* free time) we can imagine that more importance to the food preferences of the individual, in comparison to health-related aspects, will be given.

**TABLE 4. Examples of different RecHols.**

| | Example A | Example B |
|---|---|---|
| Domain/Objects | food recsys | tourism recsys |
| Goal | to maximize healthiness (*e.g.,* lose weight, consider medical conditions | to maximize amusement (*e.g.,* discover new places or popular events) |
| Primary Sphere | health | free-time |
| Secondary Sphere | free-time | health |

Similarly, as for *Example B*, we can imagine a tourism recommender system that aims to maximize the user's amusement. In this case, preferences and previous behaviors of the user will be the main elements of the profile that will drive the recommendation process. Along these aspects, secondary elements such as her health (*can he walk a lot? can he stay in crowded places?*) or the features related to the family sphere (*who is currently with the user? does she has children?*) will be considered. This will lead to the identification of a suitable recommendation that matches both the preferences of the individual and the requirements coming from the goal of the recommendation task.

This formalization allows to preliminary highlight one of the characteristics of RecHol: thanks to the *explicit* definition of the goal of the recommendation strategy, several *different* implementations of the same recommender system can exist at the same time. Indeed, a different goal will induce a different *internal logic* of the recommendation task and, in turn, a different suggestion.

### C. USER MODELING

Afterwards, it is necessary to gather information about the target user who will receive the suggestions. Accordingly, a mandatory element is the acquisition of a representation of the individual that is called *holistic user model* (HUM) (see Section 2).

A HUM can be seen as a list of elements[3] that follows this structure:

```
facet-name(feature, value,
relevance, timestamp),
```

where

- `facet_name` describes the name of the facet;
- `feature` encodes the name of the feature;
- `value` stores the value of the user feature;

[2] According to the goal of the recommendation strategy, other spheres can be easily instantiated by following the guidelines we provide.

[3] A HUM can be made available in different and more sophisticated formats, such as JSON and XML. For the sake of simplicity, we here assume that a HUM is exposed as a simple list of elements.

- `relevance` describes how much the feature is relevant in a specific recommendation task;
- `timestamp` provides temporal information about when the evidence was collected or inferred.

As for the structure of the profile, `facet-name` can be valorized with one of the names of the facets we introduced in Section 2. Then, each `feature` can be encoded as a simple keyword or can be defined on the ground of more sophisticated formalism, such as controlled vocabularies and ontologies. The `value` of the feature can be a numerical weight or a categorical value, while the `relevance` of the feature is a numerical or categorical score.

An example sketching some examples of features encoded in a HUM follows:

```
demographics(name, "John", medium, 1-jan-2020)
demographics(weight, 90kg, medium, 1-jan-2020)
preference(sushi, 0.9, medium, 1-jan-2020)
preference(u2, 0.7, medium, 1-jan-2020)
affects(mood, positive, medium, 1-jan-2020)
affects(mood, negative, medium, 31-dec-2019)
psychological-aspects(extroversion, high, medium, 1-jan-2020)
behavior(running, low, medium, 1-jan-2020)
behavior(running, low, medium, 31-dec-2019)
behavior(drinking-alcohol, high, medium, 1-jan-2020)
physical-state(heart-pressure-max, 130, medium, 1-jan-2020)
physical-state(heart-pressure-max, 125, medium, 31-dec-2019)
```

It is worth to note that, thanks to the modeling of the *timestamp*, we can also take into account the *temporal dimension* of the evidences collected or inferred in the user profile. Such a timestamp is useful to understand how a specific feature is evolving over time and can be used to model a specific aspect of the user in a large time frame.

As shown in the checklist, once a HUM is available, a RecHol designer needs to take into account the following aspects:

- defining the relevant *facets*;
- defining (if any) *domain-specific* features or *requirements*;

The first aspect is particularly relevant, since in certain recommendation settings, or in specific domains, some of the *facets* encoded in the HUM may not be relevant. As an example, a particular instance of a music recommender system may not consider as relevant all the features related to the *physical state* of the person, or a job recommender system may ignore the *behaviors* or the *physical activity* of a particular user.

Therefore, the first design choice regards the selection of the relevant portions of the HUM that will be used throughout the process. Of course, this is not a mandatory step, since a designer may gather and exploit, in a specific recommendation setting, the *entire* HUM.

The second step regards the identification of some *domain-specific* features or requirements. Even if a HUM covers most of the relevant aspects regarding the life of the individual, it may happen that something related to the specific domain of the recommendation is not encoded in the profile. For instance, a food recommender system may probably need some information about the *cooking skills* of the person or about particular requirements (*e.g.,* vegan or lactose-free food). In this case, the role of the designer is to preliminary identify such features. If these features are not covered by the information already available in the HUM, their value can be explicitly asked to the user when she interacts with the recommender system, or can be implicitly inferred from the available data (*e.g.,* some food requirements can be inferred by analyzing food habits). These domain-specific features can be understood as a further list of elements to be encoded in the profile, following the same structure we previously presented. For instance, we can store the following elements:

```
food-recsys(cooking-skill, high, high, 1-jan-2020)
food-recsys(vegan, yes, high, 1-jan-2020)
```

As shown in the example, *food-recsys* is used as facet name, since the extra features stored in the profile are related to the specific setting. Moreover, the *relevance* of the features is automatically set as *high*, since we can assume that these features are highly important to drive the recommendation process.

### D. REASONING

The reasoning step is one of the distinguishing traits of our workflow to generate *holistic recommendations*. This step borrows concepts from the area of knowledge-based recommender systems, and mainly relies on the definition (and the application) of a set of *rules* on the different entities involved in the recommendation task, that it to say, *items*, *context* and *user profiles*.

The reasoning step can be further split into two parts: a *knowledge encoding* phase and a *knowledge exploitation* phase. Intuitively, the goal of the first phase is to acquire and encode general knowledge about the goal of the recommendation task, about the requirements of the specific application domain and about the characteristics of the user; while the second phase regards the application of such rules to the current holistic user model and to the current catalogue of items. By referring to the checklist presented in Table 2, we can state that the *knowledge encoding* phase covers the first two questions, while the *knowledge exploitation* phase concerns the third and the fourth questions.

### 1) KNOWLEDGE ENCODING

At this stage, the designer should identify and encode how the different entities (contextual dimensions, user profile, features of the recommended items) involved in the process *interact* each other. In our framework, this interaction is called *'relation'*.

Basically, this task can be carried out by collecting a set of evidences that describes such relations. This can be done by exploiting several sources such as knowledge bases, documents or personal and background knowledge. By and large, we can identify three different groups of evidences we can encode in this phase: *general background evidences*, *domain-specific evidences*, *user-specific evidences*.

The first group includes general background knowledge that is directly inherited from the previously selected main sphere. As an example, if the *health* sphere is the main sphere, evidences such as *"a longer sleep improves mood"* or *"fat food increases weight"* can be easily collected and can be encoded in a knowledge base. Next, the second group relies again on the main sphere but it is directly related to the application domain of the recommendation. By considering food recommendation as a use-case scenario, we have *"eating sweets improves mood but increases weight"* or *"people overweight should reduce their caloric intakes"* as evidences. Finally, user-specific evidences describe particular relations that are dependant on the user. For example, we can mention the relation between eating sweet food and mood: for one user eating sweets may improve her mood, while for another one this behavior may worsen it.

It is worth to note that some *conflicts* between domain-specific evidences and user-specific evidences can occur. In this case, it is preferable for the designer to give a higher priority to the information that is directly inferred from users' behaviors.

To summarize, the output of this step is a set of relations forming the *knowledge base* that guides the recommendation process. Formally, we can define a dependence between two features as a *rule* having the form `if X then Y`, where `X` can be a feature encoded in the user profile, a contextual setting or the main sphere; while `Y` can be a feature of the (potentially) recommended items or a feature encoded in the profile of the user as well. In particular, five different kinds of relations may be encoded:

- `Profile` → `Profile`
- `Context` → `Recommendation`
- `Profile` → `Recommendation`
- `Sphere` → `Profile`
- `Sphere` → `Recommendation`

In all these cases, if the condition expressed in the left part of the rule is *true*, then the score of a certain feature encoded in the profile is updated or the relevance (or the suitability) of a particular recommendation is adapted.

Rules such as *"Increasing sleep time has a positive impact on mood"* fall in the first case, while the second group includes rules such as *"It is preferable to avoid horror movies when you are with children"* or *"It is preferable to avoid carbohydrates at dinner"*. The third group of rules influences the recommendations on the ground of specific elements encoded in the profile, such as *"People overweight should avoid fat food or food with high calories"*.

Finally, as for the last two groups of rules, we can state that the main sphere can affect the importance of a specific feature that is encoded in the profile or that describes the recommendation. As an example, in a *music recommendation* scenario, *some* of the demographics data about the user (*e.g.*, age, nationality) may be labeled as highly important while other features (*e.g.,* height and weight) may be considered as secondary.

Clearly, in order to complete this process, the designer needs to hold a general knowledge about the structure of HUMs and about the features that are stored in the profiles. Similarly, a selection of relevant classes of features can be done for other facets, such as the *preferences* of the user. In the previously mentioned *music recommender system* scenario, we can assume that all the music-related features are maintained in the representation at the expenses of the others. The definition of *how* this selection process can be carried out is out of the scope of this article. However, we can state that exogenous knowledge bases as ontologies or domain vocabularies can be used to easily carry out this step.

### 2) KNOWLEDGE EXPLOITATION

The knowledge exploitation phase of the reasoning step concerns the concrete *application* of the rules to the holistic user profile or to the recommendation list. In particular, we can state that when a specific rule is matched this can lead to: *(i)* the introduction of a new feature in the HUM (or the adaptation of the value for an already existing feature); *(ii)* the update of the relevance score for some of the features, according to the goal of the recommendation task inherited from the main sphere.

As for the first case, if we have collected the evidence that the user has drunk alcohol, for instance, we can introduce a new element in the user profile encoding her current mood. As for the second case, the computational task carried out in the workflow concerns the update of the evidences stored in the profiles. As an example, given *health* as the main sphere, we expect that the relevance score of features such as the *weight* of the person is updated to *high* or *very high*, while the score of other features such as her *name* or her preference for *U2* is updated to *low* or *very low*.

To conclude, we can state that the application of the rules in the reasoning step leads to the generation of a *new* holistic user profile, which encodes in the representation of the individual the output of the reasoning.

However, it is worth to note that further reasoning is carried out in the recommendation step. In particular, all the rules having something related to the object of the recommendation in the right part are used in the recommendation algorithm to further tailor the list of the suggestions to the characteristics of the user as well as to the main sphere. Details will be provided in the next section.

### E. RECOMMENDING

The recommendation phase is the *core* of the overall workflow.

As shown in the checklist, a main question concerns the definition of the strategy that is used to generate the recommendations. The range of techniques that are typically used to provide users with *context-aware recommendations* can be also used in this scenario. It is worth to note that to trigger such a context-aware recommendation algorithm a *context acquisition* step is needed.

Once the context has been acquired, the computational tasks to obtain the recommendations can be carried out. As for this aspect, *post-filtering* techniques [35] represent the most promising and suitable way to generate holistic recommendations. Differently from *pre-filtering* techniques, which preemptively filter out the evidences (e.g., user preferences) that are not collected in the current contextual situation, the goal of post-filtering context-aware recommendation algorithms is to use the contextual information as a *weighting factor* to re-rank the original recommendation list on the ground of the contextual setting.

This is particularly suitable for our goals, since we can generate a (context-aware) recommendation that relies on the preferences of the user and we can use the remaining reasoning rules to increase or decrease the relevance of the recommendation on the ground of the main sphere. As an example, to maximize the healthiness of the recommendation, a specific rule that significantly increases the score of healthy recipes can be applied. Similarly, the score of fat recipes (especially if the target user is overweight or has high blood pressure) can be lowered, given the goal of the recommendation task. More details about a post-filtering recommendation strategy based on holistic user profiles will be provided in the next section.

Finally, the designer should decide those that are typically called *Affordances*,"opportunities for action", i.e., how it is opportune to recommend:

- *style*: what is the style of the recommendation (for example, a simple information, or a suggestion, or a prescription)
- *timing*: when it is better to provide the recommendation
- *format*: in which format it is better to provide the recommendation (textual, graphical, audio)

This choice, which is another distinguishing aspect of our workflow, can be done by exploiting the personal characteristics of the individual that are encoded in the profile.

## V. A STEP-BY-STEP EXAMPLE OF THE APPLICATION OF THE FRAMEWORK: THE DESIGN PROCESS OF A HOLISTIC FOOD RECOMMENDER SYSTEM

In this section, we move the discussion from the *conceptual* level to the *concrete* level, by showing how to apply the framework we surfaced above to the design of a holistic recommender system in the food domain. We chose food as an example of a complex, health-related domain, where considering only user preferences is not enough to generate accurate recommendations, and where thus a holistic approach may improve the recommendation's outcomes

**TABLE 5. Scope definition for a holistic food recommender system.**

| Question | Answer |
|---|---|
| *What is the domain of the recommendation?* | Food |
| *What is the object of the recommendation?* | Recipes |
| *What is the goal of the recommendation task?* | To provide healthy food recommendations |
| *Which contextual dimensions impact on the process?* | Time of the day, companionship and the location |

The goal of this section is to show *step by step* how our methodological framework can be exploited to simplify the deployment of a holistic recommender. To this aim, we will outline how the design of the system followed RecHol principles. For each step, we display in a table the questions of the designer's checklist (see Table 2) pertaining to that step and give them an answer based on the food recommender we designed.

#### 1) SCOPE DEFINITION

As we have seen in the previous Section, scope definition is the first step that needs to be faced when designing a RecHol. The designer primarily chooses the domain and the object of the recommendation, the goal of the recommendation task and the contextual dimensions that impact on the process. Table 4 precisely describes the choices we made, pointing out those questions we had to answer to start the design process of the holistic food recommender.

By explicitly defining the goal of the recommendation task, as well as the contextual dimensions that impact on the process, the designer is forced to think of the overall structure of the recommendation method *before* the system is implemented. This is certainly one of the advantages coming from the adoption of the checklist we designed for the development of holistic recommender systems, since all the main factors that influence the recommendation task are put in the foreground in order to be considered by the designer and the developer of the algorithm.

#### 2) SPHERES MODELING

The second step of RecHol methodological framework is the modeling of the spheres, which is extremely important because it induces the designer to thoroughly consider the macro-factors that may influence the recommendation process. In our application example, we first defined "health" and "free time" as those spheres involved in the recommendation process. Then, we set "health" as the main sphere and "free time" as the secondary sphere. This means that the recommendation strategy of our system will maximize the healthiness of the suggested recipes, taking seriously into

**TABLE 6. Spheres modeling for a holistic food recommender system.**

| Question | Answer |
|---|---|
| *Which spheres can we define?* | Health and free time |
| *What is the main sphere and the secondary sphere?* | "Health" as the main sphere and 'free time' as the secondary sphere, in order to firstly take into account the user's health-related characteristics preferences and secondly consider her preferences |

**TABLE 7. User modeling for a holistic food recommender system.**

| Question | Answer |
|---|---|
| *Which facets of a holistic user model are needed?* | Demographics (sex, age), affects (mood), health data (height, weight, amount of sleep, stress level), behaviors (amount of physical activity), interests (food preferences). |
| *Is there any domain-specific features or requirement?* | Frequency of home-cooked meals, recipe website usage, cooking experience, restrictions (vegan, lactose-free, low-nickel, etc.) |

account the health-related characteristics of the user; while elements such as food preferences and previously visited restaurants, which pertain to the "free time" sphere, would be also considered, but would have a secondary importance in the recommendation process.

It is worth to notice that the whole recommendation process is deeply impacted by the choice of the main sphere. By selecting health as the main sphere, we prioritized health-related issues over users' preferences about food. By forcing the developer to think about the spheres involved in the recommender system and their priority, the methodological framework makes the design of the recommendation algorithm easier.

### 3) USER MODELING

The third step of the methodological framework consists in acquiring a HUM of the user and deciding which facets of the profile need to be exploited. Moreover, we need to decide whether domain-specific information that may not be encoded in the HUM is needed. Therefore, during the design process of our holistic food recommender, we first selected all those facets contained in the HUM of our potential users that were fundamental to provide health-based food recommendations. Then, we also included some task-specific and domain-specific features, such as cooking experience and food restrictions, which are certainly relevant to the specific recommendation task.

Table 6 points out all the facets we considered in the HUM as relevant to recommend food items aimed at supporting the user's health. It shows how our methodological framework encourages the designer to carefully consider all those kinds of data that may enrich the recommendation process. Furthermore, it makes her reflect on those features that are currently not included in the HUM of the system's potential users but that could be potentially relevant for the recommendation task.

**TABLE 8. Reasoning mechanisms encoded in a holistic food recommender system.**

| Question | Answer |
|---|---|
| *What kind of knowledge do we encode from the main sphere?* | Since health is the main sphere, we needed health-related information. In particular, we encoded some general rules about the impact of nutrients and food features on the health state of the user. |
| *What kind of knowledge do we encode by analyzing user behavior?* | We analyze user behavior to infer more specific rules about the influence and the impact of food characteristics on the user herself. As an example, we infer that a specific user has an increase in the mood when she eats sweets. |
| *How do the rules impact on the representation of the user?* | We introduced some weighting scheme to increase/decrease the weight or the relevance of a feature if a specific situation happens. For instance, we increase mood (e.g., from neutral to positive) if the user has a high daily intake of sweets and she is "sensitive" to sweets. |
| *How do the rules impact on the recommendation process?* | We introduced some weighting scheme to increase/decrease the weight or the relevance of a recommendation in case a specific situation happens. As an example, we decrease the relevance score of an item if the user is overweight, or we filter out a meat recipe if the user is vegetarian. |

### 4) REASONING

The fourth step of our methodological framework relates to the need to trigger some reasoning mechanisms in order to update the representation of the user on the ground of the current contextual situation, as well as to increase or decrease the score of potential recommendations.

Therefore, we first encoded general rules about the impact of food features on the user's health, as health was the main sphere of our recommender. Then, we considered all those aspects of the user's behavior that could affect her eating habits, such as the possibility that her mood could be influenced by specific foods (e.g., sweets). Finally, we defined those rules in charge of updating the user profile depending on the circumstances, as well as those that regulate how the recommendations are delivered on the basis of the user's current situation. In Table 7 we summarize the reasoning strategy we implemented in our food recommender prototype, by pointing out the questions we had to answer during the design process.

As we may see, all the *knowledge* we encoded for our holistic food recommender is directly inspired by the *health goal* we defined at the beginning of the design pipeline. The methodological framework we outlined in the previous section allowed us to focus on all those "rules" that could be relevant for the recommendation goal: it suggested that we first collect a set of evidences describing how food and health are related (general background evidences and domain-specific evidences), as well as what kinds of user's features could influence her eating behavior (user-specific evidences); then, that we encode those relations that could have an impact on the recommendation process. For instance,

**TABLE 9. Recommendation strategy for a holistic food recommender system.**

| Question | Answer |
|---|---|
| *Which algorithm can we run to obtain recommendations?* | We can draw on the literature about context-aware recommender systems. The most promising direction is to implement a post-filtering strategy that exploits the rules encoded in the knowledge-aware component of the reasoning mechanism to re-rank the recipes. |
| *What is the best style to present the recommendations?* | We will show the recipe's name, an image and the ingredients or the nutrients. In case, it is possible to provide the user with preparing instructions. |
| *When is it better to provide the recommendation?* | This aspect is not covered by our prototype, since we aimed to design a web application that will provide users with food recommendations on request. |

we encoded rules that lower the relevance score of sugary drinks if the user is overweight.

The introduction of these reasoning mechanisms is a characteristic that distinguishes holistic recommender systems from other recommendation paradigms, since the knowledge-based part of the pipeline works together with the context-aware component of the recommender in order to meet the overall goal of the recommendation strategy.

**Algorithm 1** Holistic Food Recommendation Process

**Require:** Holistic user model $hum(u)$ for user $u$
**Require:** Dataset of Recipes $R = \{r_1 \ldots r_n\}$
**Require:** Food Knowledge encoded as rules $K = \{k_1 \ldots k_m\}$

```
/* filtering phase */
for all r_i ∈ R do
    if FILTER(r_i)==FALSE then
        add r_i to candidate recipes
        R'=R' ∪ r_i
    end if
end for
/* ranking phase */
for all r_i ∈ R' do
    score(r_i) ← popScore(r_i)
    for all k_i ∈ K do
        if hum(u) matches k_i then
            /* apply rule k_i by calculating holistic(u, r) */
            score(r_i) ← holistic(u, r)
        end if
    end for
end for
Rank recipes r_i ∈ R' according to score(r_i)
return top-1 recommendation (recipe with the highest score)
```

### 5) RECOMMENDING

The last step refers to the recommendation itself, i.e., how to tune the recommendation algorithm to obtain the final recipe suggestion. In table 8 we give an answer to the questions posed by our methodological framework which focus the designer's attention on decisions regarding the recommendation strategy to choose and the way of presenting it, that is, what kind of information is needed to be displayed to the user and in which form. The food recommender system will thus implement a post-filtering recommendation strategy and show information about recipe name, ingredients, image and possibly preparing instructions.

The information we provide in Table 8 represents the output of the whole process, that is to say, a *food recommendation* that matches the preferences of the user by meeting the goals and the contextual constraints of the recommendation task.

## VI. IMPLEMENTATION OF THE HOLISTIC FOOD RECOMMENDER SYSTEM

In the following, we give some technical details about our food recommender system, which has been designed by following the steps presented above. The workflow is depicted in Figure 1. In a nutshell, it is based on three main components: a Profiler, a Filter and a Ranker.

The recommendation process starts with a user asking for a specific recipe (*main course*, *second course*, or *dessert*). Next, in the first step of the workflow, we acquire a Holistic User Model of the target user. In particular, we acquire only the personal information that is reported in Table 6 and we explicitly ask the user to provide some domain-specific information. As for the HUMs, we exploit a public endpoint exposed by Musto *et al.* [30], a platform for building holistic user models, while a screenshot showing some of the domain-specific questions we asked to the user is reported in Figure 2.

Once the profile is built, the Filter generates a preliminary set of candidate recipes by filtering non-compliant recommendations. This step is carried out by analyzing the user's food restrictions and cooking experience, and subsequently removing recipes from the list of candidates which contain ingredients that a user wishes to avoid (e.g., lactose, meat), or that are too complex to prepare. It is worth to note that the Filter component implements a basic part of the reasoning mechanisms we designed in our food recommender systems. As reported in Table 7, a food recommender system should filter out non-compliant recipes, and this component carries out this step.

Next, once the filtering is completed, the *real* recommendation process comes into play. The set of candidate recommendations is still large and requires further re-ranking, which is done by the Ranker component. By referring to the previous checklist, this component implements some of the principles of the RecHol *reasoning* strategy together with some intuitions about the calculation of the *recommendations*.

Formally, given a user $u$, the goal of this component is to assign to each recipe $r$ a $score(r, u)$, to rank all the candidate recipes and to identify the top recipes that best match the user, in terms of her characteristics, contextual setting

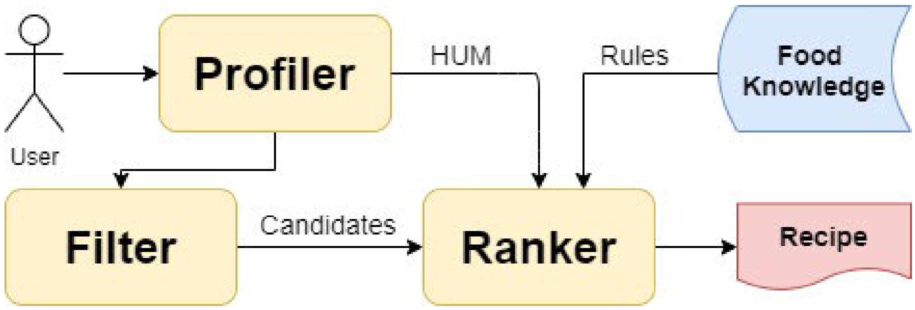


**FIGURE 1.** **Workflow of our Holistic food recommender system.**

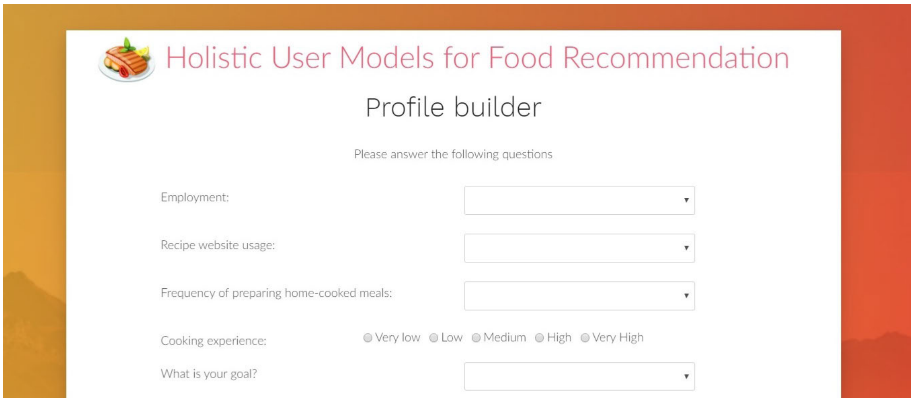


**FIGURE 2.** **Domain-specific features to be encoded in our food recommender system.**

and constraints. To do so, we propose a scoring mechanism that combines three different factors: a *preference* score, a *popularity-based* score and a *reasoning-based* score. Formally, the scoring formula can be defined as:

$$holistic(u, r) = pref(u, r) * popScore(r) * reasoning(u, r)$$

where $pref(u, r)$ is a preference score that is based on previous food preferences of the user. For instance, if a user has previously liked one of the ingredients of the recipe $r$ the preference score will be higher than 1 (thus, it will act as a boost factor); otherwise, it will be lower than 1. If the user did not provide any food preference or her HUM does not contain information about food consumption, this factor will be equal to 1 (and it will not impact on the overall score). Next, $popScore(r)$ quantifies the popularity of the recipe, and is implemented as follows ($count(r)$ is a rating counter):

$$popScore(r) = avgRating(r) * log_{10}(count(r))$$

Finally, the knowledge-aware part (i.e., $reasoning(u, r)$) of the scoring formula is a modifier that increases or decreases the score of a recipe by using some general knowledge about food choices. By referring to the literature about context-aware recommender systems, we could state that $reasoning(u, r)$ acts as a *weighting factor*, as it commonly happens in post-filtering context-aware recommender systems.

As stated in the previous section, such knowledge is encoded as *rules*, having the form *factor* → *modifier*. If the left part of a rule is satisfied, the modifier is applied to the recipe or to the characteristics of the user.

**TABLE 10.** **A Selection of rules encoded in our knowledge-aware recommender.**

| Facet | Factor | Modifier |
|---|---|---|
| 2*Demographics | *BMI>25* | low-calories recipes $\uparrow$<br>high-calories recipes $\downarrow$ |
| | *BMI<25* | high-calories recipes $\uparrow$ |
| Mood | *bad* | recipes with sugar $\uparrow$ |
| Behaviors | *>30 min/day activity* | high-protein recipes $\uparrow\uparrow$<br>high-calories recipes $\uparrow$ |
| Health | *stress=yes* | high-sodium recipes $\uparrow$ |
| Health | *sleep=low* | high-magnesium recipes $\uparrow$ |
| Health | *depression=yes* | high-fat recipes $\downarrow$ |

Although we cannot discuss all the details of our knowledge-based scoring formula, we emphasize that these rules are based on common-sense knowledge about food choices. For example, our knowledge-aware recommender awards a lower score to recipes that are high in calories, if the user has a high Body Mass Index. Moreover, we consider insights from recent studies about the link between user factors and food consumption, such as the relation between stress and the amount of salt in recipes.[4]

Table 9 reports some of the rules we encoded in our recommender system, listing user aspects and factors, and the corresponding modifiers.

[4]https://oklahoman.com/article/feed/687315/did-you-know-salt-reduces-stress

Once the scoring function has been identified for the most relevant recipes for the user, these can be returned to her. In the next session we will discuss the outcomes of a preliminary experiment that aims to evaluate the effectiveness of this recommendation strategy.

It is worth to note that this implementation only represents a preliminary proof-of-concept prototype of a holistic food recommender system. Indeed, the current prototype is characterized by some limitations and there is large room for improvement. For example, we did not use reasoning mechanisms to analyze user behaviors and to infer further user features. To this end, machine learning techniques to extract patterns from data could be implemented. Moreover, as for the contextual information, we did not use information about the companionship of the user. This can be easily obtained by collecting several HUMs and by combining the single relevance scores of the users in a global relevance score that takes into account the characteristics of a group of person (as it happens in group recommender systems).

However, beyond these limitations, our proof-of concept prototype clearly shows the potential of our framework. As the goal of the recommendation task changes (e.g., food preferences can be prioritized with respect to healthiness), different rules may be encoded in the reasoning component, and this will lead our recommender system to a completely different behavior. Similarly, if more information about the user can be inferred or acquired, this will also change the recommendation returned by the system. Generally speaking, we can state that by exploiting our methodological framework it is possible to easily build a holistic recommender system that tries to take the best out of users' personal data.

## VII. EVALUATION OF THE EFFECTIVENESS OF THE HOLISTIC FOOD RECOMMENDER SYSTEM

In order to evaluate the effectiveness of our framework, we carried out a preliminary experiment in the *food recommendation* domain.

We asked users to interact with the platform we introduced in the previous Section and we provided each user with three different pairs of recipes. Next, each user was asked to choose the recipe she preferred the most. For each couple of recipes, one was obtained by running our *holistic recommendation strategy* while the other was obtained through a simple *popularity-based baseline*.

Recipes were sampled from a database of 4,671 recipes, which we share online.[5] The recipes were obtained from the popular Italian food community platform GialloZafferano,[6] and translated into English. The recipes contain information about their name, category, preparation difficulty, as well as their ingredients, macro-nutrients, calories, rating count, and average website rating. Moreover, they also include several binary tags, such as *vegetarian*, *vegan*, *lactose-free*, and *low-nickel*.

Our experiment was carried out by recruiting a sample of 200 participants on Amazon MTurk, who were rewarded with 0,5 USD for a HIT, which took them on average five minutes to complete. As previously explained, each user interacted with the food recommender system deployed online[7] by providing information about their gender, age, BMI (5-point scale), recipe website usage (4-point scale), cooking experience (5-point scale), and mood (i.e., 'good', 'neutral', or 'bad'). In addition, users also provided information about sleep length, stressed and depressed feelings (yes/no), dietary goals (lose or gain weight, or none), and dietary constraints (e.g., vegan, low-nickel). This information was necessary to feed our holistic user profiles according to the structure we defined in the previously presented checklist.

Next, we ran two different strategies to provide users with food recommendations. Each user was given three pairs of recipes, where each pair represented a different part of a meal: main courses (i.e., mostly pasta dishes), second courses (i.e., mostly meat-based dishes), and desserts. For each comparison, we presented one recipe returned by our holistic recommendation strategy and one returned by running a popularity-based baseline. An example of the presentation of the recipes is shown in Figure 3. For each pair of recipes, we asked users which of the two recipes they preferred the most, or whether they preferred none of them. In addition, we also inquired about their underlying motivations for choosing a recipe (if any), presenting four propositions about the chosen recipe on 5-point Likert scales: "It seems savory and tastier", "It helps me to eat more healthily", "It would help me to lose/gain weight", and "It seems easier to prepare".

Through this experiment we tried to evaluate: *(i)* whether users preferred recipes suggested by our holistic recommendation strategy to those obtained through a popularity-based baseline; *(ii)* which users' factors and motivations influenced the choice of a recipe.

Given that this article focuses on the conceptual model that allows the construction of holistic recommender systems, a thorough discussion of the outcomes of this experiment is out of the scope of the current paper. For a more detailed overview we suggest to refer to [31].

In this article we provide an overview of the findings of the experiment. Through a sequence of two-sample t-tests we found that users were more likely to choose holistic recipes (49.2%) over popular recipes (33.0%) for the second course ($t(190) = 2.51$, $p = 0.013$). Conversely, as for main courses (54.0% (popular) vs 32.5% (holistic)) and desserts (54.5% vs 36.1%) popular recipes were preferred ($t(190) = -3.27$, $p < 0.01$ and $t(190) = -2.70$, $p < 0.01$ respectively). These mixed results suggested that the holistic user model did not entirely outperform the popular baseline, which, nonetheless, is hard to beat in food recommendation scenarios [48], and a more thorough analysis of *personal factors* that influence users' choice is needed.

[5] https://tinyurl.com/recipes-uniba

[6] https://www.giallozafferano.it/ricette-cat/

[7] It can be found at: http://90.147.102.243:8080/foodrecsys/

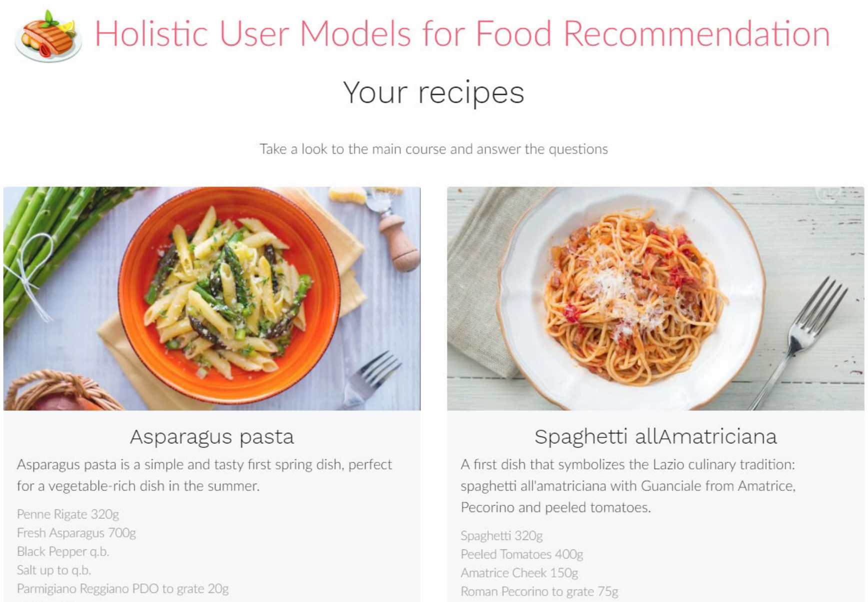


**FIGURE 3.** **Two recipe recommendations returned by our holistic food recommender system.**

To this end, we carried out a statistical analysis based on logistic regression to make emerge user factors and motivations that influence such choices. An interesting outcome of the experiment regards the role of *mood*, one of the novel features we encoded in the model thanks to the holistic user profile. According to our analysis, users who reported to be in a good mood preferred the holistic recipe ($R^2 = .88$ $p < 0.05$). This is a notable finding, as mood is a common factor in the music recommender domain [44], but not in food.

Moreover, a very interesting outcome emerged from the analysis of *users' motivations*. Indeed, data showed that users with health-related goals were more likely to choose the holistic main course ($R^2 = .98$ $p < 0.001$) while *taste* was one of the main reasons behind the choice of popular recipe. This is a very notable (and fundamental) outcome, since our recommendation strategy was designed with the aim of rewarding healthier recipes and allowing users to make healthier food choices. Such a goal, which guided the whole design of the recommendation process, was confirmed by the results we obtained, since users with interest in healthiness were actually supported in achieving their goals and further confirmed the effectiveness of our design.

Finally, we also investigated the impact of the ingredients of the recipe in the decision-making process. In this case, results showed that the recipes returned by our holistic recommendation strategy were preferred if they were *low-carb* and *high-protein*, while *popular* recipes were more likely to be chosen if they contained more carbs and more saturated fat. This finding is in line with the health-related motivations to choose a holistic recipe we previously discussed.

To sum up, this preliminary experiment provided us with mixed result. Recipes returned by our holistic recommendation strategy were the most selected ones only for the second courses. This suggests that the set of information we encoded in user profiles as well as the knowledge we encoded in the reasoning component of the pipeline need to be refined and improved. However, as already shown by Trattner *et al.* [48], it is necessary to reemphasize that popularity-based baseline is often hard to beat, especially in the food recommendation scenario.

However, the underlying motivations of user choices (such as taste, health, and easiness to prepare the recipe) do signal that a holistic user model can appeal to users who wish to purse healthy food choices. This is perfectly in line with the overall design of our recommendation strategy and confirms the validity of the conceptual model we introduced in this article.

## VIII. DISCUSSION AND CONCLUSION

In this article we have introduced the concept of holistic recommendation (RecHol), namely a set of suggestions generated by exploiting a comprehensive representation of the user, which relies on personal information coming from different heterogeneous data sources. Then, we presented a methodological framework that can guide the designers in the process of designing holistic recommenders. Finally, we validated the framework presenting how it can be used to design a holistic-based food recommender.

As future work, we plan to assess the validity of the holistic recommendations by testing the food recommender presented in Section VI with real users.

Here, we discuss the main challenges opened by the proposed approach.

*Privacy preservation.* Holistic recommendation adds new complex dimensions to the problem of computing privacy-preserving models [42]: the interplay between lifestyle aspects, purchasing behaviour, and contextual properties, for instance, could reveal private habits and preferences. Here, the most important issue concerns the definition of leakage or attack models and the related countermeasures leveraging the uniqueness of the relations between multiple dimensions. Moreover, we need to improve the user's awareness of her privacy, by enhancing her perception of the trade-off between the accuracy of the recommendations given and the amount of private sensitive information that may be disclosed.

*Ethical issues*. In the light of the characteristics of holistic-based recommendations, a lot of ethical issues arise, especially in relation to the health domain. The user empowerment and engagement are here essential, and we should, as designers, encourage the user's reflection on the recommendations. In a preliminary empirical analysis, [20] show that presenting uncertainty to the user might help her reflect on the recommendation. Moreover, it is crucial to provide accurate recommendations, since wrong suggestions, especially in the health domain, can be harmful for the user [15].

*Explainability*. Explainability is the ability of an algorithm to be interpretable [27]. In holistic recommendation, a way to achieve explainability consists in enabling the user to visualize her models following a *scrutable* approach [54]. Open issues here relate to how to make such a complex User Model scrutable. However, it is not feasible to present all the data to the user, since they could cause information overload. This problem is strictly related to the granularity of the collected data. Data to be visualized could be changed in format according to the specific application that is used by the individual, and/or the user's features (e.g., goals or expertise), and/or the specific context. In this case, the exploitation of the information gathered from the Linked Open Data cloud [28] can be useful to encode machine-readable and explainable features in the profile of the user.

A holistic recommender should gather and maintain vast knowledge. Thus, a second main issue is related to knowledge creation and management.

*Rules derivation*. How to derive rules is not trivial. An option could be asking the user about what rules are valid for her. However, this can bother the user, since it can lead to a combinatorial explosion of possibilities. Moreover, it is possible that the user herself is not aware of some of such relations. Alternatively, it is possible to learn relations in the form of rules among features extracted from the data, by exploiting some machine learning techniques for feature selection, or doing some statistical analysis [32]. This can provide results only if a huge amount of data is available: however, the results could be incomprehensible for the user. Otherwise, it could also be possible to inherit rules from similar users in a sort of stereotyping reasoning: for example the positive relation among drink alcohol and social life is more present among young people, and thus if the user belongs to this category, we can add such rules in her model.

*Conflict management* The huge availability of heterogeneous data can significantly increase the information (and the knowledge) about the user, but it is not simple to decide how to proceed when tension between the spheres arises. For example, a priority mechanism could be created, giving more power to the user or to an external authority such as the physician, in case of health recommendations.

## REFERENCES

[1] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin, "Incorporating contextual information in recommender systems using a multidimensional approach," *ACM Trans. Inf. Syst.*, vol. 23, no. 1, pp. 103–145, Jan. 2005.

[2] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, Jun. 2005.

[3] G. Adomavicius and A. Tuzhilin, "Context-aware recommender systems," in *Recommender System Handbook*. Berlin, Germany: Springer, 2011, pp. 217–253.

[4] I. M. A. Jawarneh, P. Bellavista, A. Corradi, L. Foschini, R. Montanari, J. Berrocal, and J. M. Murillo, "A pre-filtering approach for incorporating contextual information into deep learning based recommender systems," *IEEE Access*, vol. 8, pp. 40485–40498, 2020.

[5] L. Baltrunas, M. Kaminskas, B. Ludwig, O. Moling, F. Ricci, A. Aydin, K. H. Lüke, and R. Schwaiger, "Incarmusic: Context-aware music recommendations in a car," in *Proc. Int. Conf. Electron. Commerce Web Technol.*, 2011, pp. 89–100.

[6] L. Baltrunas and F. Ricci, "Context-based splitting of item ratings in collaborative filtering," in *Proc. 3rd ACM Conf. Recommender Syst.*, New York, NY, USA, 2009, pp. 245–248, doi: 10.1145/1639714.1639759.

[7] C. Biancalana, F. Gasparetti, A. Micarelli, and G. Sansonetti, "An approach to social recommendation for context-aware mobile services," *ACM Trans. Intell. Syst. Technol.*, vol. 4, no. 1, pp. 1–31, Jan. 2013, doi: 10.1145/2414425.2414435.

[8] R. Burke, "Knowledge-based recommender systems," *Encyclopedia Library Inf. Syst.*, vol. 69, no. 32, pp. 175–186, 2000.

[9] R. Burke, "Hybrid recommender systems: Survey and experiments," *User Model. user-Adapted Interact.*, vol. 12, no. 4, pp. 331–370, 2002.

[10] B. S. Cantador and P. Cremonesi, "Cross-domain recommender systems," in *Recommender System Handbook*. Berlin, Germany: Springer, 2015, pp. 919–959.

[11] F. Cena, S. Likavec, and A. Rapp, "Real world user model: Evolution of user modeling triggered by advances in wearable and ubiquitous computing: State of the art and future directions," *Inf. Syst. Frontiers*, vol. 15, pp. 1–26, May 2018, doi: 10.1007/s10796-017-9818-3.

[12] F. Cena, A. Rapp, S. Likavec, and A. Marcengo, "Envisioning the future of personalization through personal informatics: A user study," *Int. J. Mobile Hum. Comput. Interact.*, vol. 10, no. 1, pp. 52–66, Jan. 2018, doi: 10.4018/IJMHCI.2018010104.

[13] F. Cena, A. Rapp, C. Musto, and P. Lops, "Towards a conceptual model for holistic recommendations," in *Proc. 26th Conf. User Model., Adaptation Personalization*, Jul. 2018, pp. 207–210, doi: 10.1145/3213586.3225248.

[14] A. Dijksterhuis, "Think different: The merits of unconscious thought in preference development and decision Making.," *J. Personality Social Psychol.*, vol. 87, no. 5, pp. 586–598, 2004.

[15] J. D. Ekstrand and M. D. Ekstrand, "First do no harm: Considering and minimizing harm in recommender systems designed for engendering health," in *Proc. Workshop Recommender Syst. Health*, 2016, p. 14.

[16] R. G. Farrell, C. M. Danis, S. Ramakrishnan, and W. A. Kellogg, "Intrapersonal retrospective recommendation: Lifestyle change recommendations using stable patterns of personal behavior," in *Proc. 1st Int. Workshop Recommendation Technol. Lifestyle Change*, Dublin, Ireland, 2012, p. 24.

[17] L. R. Ferretto, E. A. Bellei, D. Biduski, L. C. P. Bin, M. M. Moro, C. R. Cervi, and A. C. B. De Marchi, "A physical activity recommender system for patients with arterial hypertension," *IEEE Access*, vol. 8, pp. 61656–61664, 2020.

[18] M. Ge, F. Ricci, and D. Massimo, "Health-aware food recommender system," *In: Proc. 9th ACM Conf. Recommender Syst.*, New York, NY, USA, 2015, pp. 333–334. [Online]. Available: https://dl.acm.org/doi/abs/10.1145/2792838.2796554, doi: 10.1145/2792838.2796554.

[19] M. Gil, R. El Sherif, M. Pluye, B. C. M. Fung, R. Grad, and P. Pluye, "Towards a knowledge-based recommender system for linking electronic patient records with continuing medical education information at the point of care," *IEEE Access*, vol. 7, pp. 15955–15966, 2019.

[20] K. Herrmanny and A. Doganguen, "The impact of prediction uncertainty in recommendations for health-related behavior," in *Proc. Int. Workshop Health Recommender Syst.*, 2017, p. 14.

[21] M. Iqbal, M. A. Ghazanfar, A. Sattar, M. Maqsood, S. Khan, I. Mehmood, and S. W. Baik, "Kernel context recommender system (KCR): A scalable context-aware recommender system algorithm," *IEEE Access*, vol. 7, pp. 24719–24737, 2019.
[22] A. Jameson, "Choice architecture for human-computer interaction," *Found. Trends Hum.-Comput. Interact.*, vol. 7, nos. 1–2, pp. 1–235, 2013, doi: 10.1561/1100000028.
[23] A. Jameson, M. C. Willemsen, A. Felfernig, G. M. de, P. Lops, G. Semeraro, and L. Chen, "Human decision making and recommender systems," in *Recommender System Handbook*. Springer, 2015, pp. 611–648.
[24] B. Li, Q. Yang, and X. Xue, "Can movies and books collaborate?: Cross-domain collaborative filtering for sparsity reduction," in *Proc. 21st Int. Jont Conf. Artif. Intell.*, San Francisco, CA, USA, 2009, pp. 2052–2057.
[25] B. Libet, "Unconscious cerebral initiative and the role of conscious will in voluntary action," *Behav. Brain Sci.*, vol. 8, no. 4, pp. 529–539, Dec. 1985.
[26] K. U. Mayer and U. Schoepflin, "The state and the life course," *Annu. Rev. Sociol.*, vol. 15, no. 12, pp. 187–209, 1989.
[27] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," *Pattern Recognit.*, vol. 65, pp. 211–222, May 2017.
[28] C. Musto, P. Basile, P. Lops, M. de Gemmis, and G. Semeraro, "Introducing linked open data in graph-based recommender systems," *Inf. Process. Manage.*, vol. 53, no. 2, pp. 405–435, Mar. 2017.
[29] C. Musto, M. Polignano, G. Semeraro, M. de Gemmis, and P. Lops, "Myrror: A platform for holistic user modeling," *User Model. User-Adapted Interact.*, vol. 30, no. 3, pp. 477–511, Jul. 2020.
[30] C. Musto, G. Semeraro, C. Lovascio, G. M. de, and P. Lops, "A framework for holistic user modeling merging heterogeneous digital footprints," in *Proc. UMAP*, 2018, pp. 97–101.
[31] C. Musto, C. Trattner, A. Starke, and G. Semeraro, "Towards a knowledge-aware food recommender system exploiting holistic user models," in *Proc. 28st Conf. User Modeling, Adaptation Personalization*, 2020, pp. 1–5.
[32] A. Odiä, "Relevant context in a movie recommender system: Users' opinion vs. statistical detection," in *Proc. 4th Workshop Context-Aware Recommender Syst.*, 2011, pp. 1–6.
[33] Y. Oh, A. Choi, and W. Woo, "U-BabSang: A context-aware food recommendation system," *J. Supercomput.*, vol. 54, no. 1, pp. 61–81, Oct. 2010.
[34] W. Pan, E. Xiang, and Y. Yang, "Transfer learning in collaborative filtering for sparsity reduction," in *Proc. 24th AAAI Conf. Artif. Intell.*, 2010, pp. 210–235.
[35] U. Panniello, A. Tuzhilin, M. Gorgoglione, C. Palmisano, and A. Pedone, "Experimental comparison of pre- vs. post-filtering approaches in context-aware recommender systems," in *Proc. 3rd ACM Conf. Recommender Syst.*, New York, NY, USA, 2009, pp. 265–268, doi: 10.1145/1639714.1639764.
[36] M. Prinz and G. Bauer, "Contextualizing mobile applications for context-aware recommendation," in *Proc. Pervasive*, 2010, pp. 1–5.
[37] E. Quah and J. Haldane, *Cost-Benefit Analysis*. Evanston, IL, USA: Routledge, 2007.
[38] S. Raza and C. Ding, "Progress in context-aware recommender systems—An overview," *Comput. Sci. Rev.*, vol. 31, pp. 84–97, Feb. 2019.
[39] F. Ricci, L. Rokach, and B. Shapira, "Recommender systems: Introduction and challenges," in *Recommender System Handbook*. Berlin, Germany: Springer, 2015, pp. 1–34.
[40] H. Roitman, Y. Messika, Y. Tsimerman, and Y. Maman, "Increasing patient safety using explanation-driven personalized content recommendation," in *Proc. 1st ACM Int. Health Informat. Symp.*, 2010, pp. 430–434.
[41] M. Rokicki, E. Herder, and E. Demidova, "What's on my plate: Towards recommending recipe variations for diabetes patients," in *Proc. UMAP Workshops*, 2015, pp. 1–6.
[42] P. Samarati and L. Sweeney, "Generalizing data to provide anonymity when disclosing information (abstract)," in *Proc. ACM PODS*, Seattle, WA, USA, Jun. 1998, p. 188.
[43] H. Samin and T. Azim, "Knowledge based recommender system for academia using machine learning: A case study on higher education landscape of pakistan," *IEEE Access*, vol. 7, pp. 67081–67093, 2019.
[44] M. Schedl, P. Knees, B. McFee, D. Bogdanov, and M. Kaminskas, "Music recommender systems," in *Recommender System Handbook*. Berlin, Germany: Springer, 2015, pp. 453–492.
[45] P. Sitkrongwong, S. Maneeroj, and A. Takasu, "Latent probabilistic model for context-aware recommendations," in *Proc. IEEE/WIC/ACM Int. Joint Conferences Web Intell.*, Washington, DC, USA, Nov. 2013, pp. 95–100, doi: 10.1109/WI-IAT.2013.14.
[46] E. Stephan, H. Weibelzahl, K. Dominikus, M. Heckmann and J. Schildt, "Daptive recommendations for patients with diabetes," in *Proc. UMAP*, 2015, pp. 1–3.
[47] R. Yera Toledo, A. A. Alzahrani, and L. Martinez, "A food recommender system considering nutritional information and user preferences," *IEEE Access*, vol. 7, pp. 96695–96711, 2019.
[48] C. Trattner, D. Moesslang, and D. Elsweiler, "On the predictability of the popularity of online recipes," *EPJ Data Sci.*, vol. 7, no. 1, p. 66, Dec. 2018.
[49] A. C. Valdez, M. Ziefle, K. Verbert, A. Felfernig, and A. Holzinger, "Recommender systems for health informatics: State-of-the-art and future perspectives," in *Machine Learning for Health Informatation*. Berlin, Germany: Springer, 2016, pp. 391–414.
[50] V. Vijayakumar, S. Vairavasundaram, R. Logesh, and A. Sivapathi, "Effective knowledge based recommender system for tailored multiple point of interest recommendation," *Int. J. Web Portals*, vol. 11, no. 1, pp. 1–18, Jan. 2019.
[51] I. Viktoratos, A. Tsadiras, and N. Bassiliades, "Combining community-based knowledge with association rule mining to alleviate the cold start problem in context-aware recommender systems," *Expert Syst. Appl.*, vol. 101, pp. 78–90, Jul. 2018.
[52] N. M. Villegas, C. Sánchez, J. Díaz-Cely, and G. Tamura, "Characterizing context-aware recommender systems: A systematic literature review," *Knowl.-Based Syst.*, vol. 140, pp. 173–200, Jan. 2018.
[53] S. Warburton, *Digital identity and social media*. Hershey, PA, USA: IGI Global, 2012.
[54] R. Wasinger, J. Wallbank, L. Pizzato, J. Kay, B. Kummerfeld, M. Böhmer, and A. Krüger, "Scrutable user models and personalised item recommendation in mobile lifestyle applications," in *User Modeling Adaptation, Personalization*, S. Carberry, S. Weibelzahl, A. Micarelli, G. Semeraro, Eds. Berlin, Germany: Springer, 2013, pp. 77–88.
[55] M. Wiesner and D. Pfeifer, "Health recommender systems: Concepts, requirements, technical basics and challenges," *Int. J. Environ. Res. Public Health*, vol. 11, no. 3, pp. 2580–2607, Mar. 2014.
[56] U. Yasavur, R. Amini, and C. L. Lisetti, "User modeling for pervasive alcohol intervention systems," in *Proc. 1st Int. Workshop Recommendation Technol. Lifestyle Change*, 2012, pp. 29–34.
[57] Q. Zhu, S. Wang, B. Cheng, Q. Sun, F. Yang, and R. N. Chang, "Context-aware group recommendation for point-of-interests," *IEEE Access*, vol. 6, pp. 12129–12144, 2018.

**FEDERICA CENA** is currently an Associate Professor with the Computer Science Department, University of Turin. Her research interests include human–computer interaction and artificial intelligence. In particular, she is working on user modeling, personalization and ubiquitous computing, and semantic web for user modeling and adaptation. In the last years, she is mainly devoted in studying the implications of the Internet of Things for user model and personalisation. She has served as the Program Committee Member for the most important conferences in her field (ACM Hypertext, ACM UMAP, and ACM IUI) and she was, among the other, Programme Chair of ACM UMAP 2017, the Workshop and Tutorial Chair of ACM UMAP 2016 and ACM Hypertext 2014, and the Publicity Chair of ACM Hypertext 2014. She organized international workshops on social adaptive semantic web (UMAP 2010-11) and personal informatics and quantify self (Ubicomp 2015-16-17). She was an Editor of several special issues on semantic adaptive web for *ACM Transactions on Intelligent Systems and Technology* Journal, in 2013, personal big data for *ACM Transactions on Interactive Intelligent Systems* Journal, in 2016, gamification for the *International Journal of Human-Computer Studies*, quantified self on *Computer Journal*, and advanced personalized mobile service on *Mobile Information Systems* Journal, in 2017.

**AMON RAPP** (Associate Member, IEEE) is currently an Assistant Professor with the Computer Science Department, University of Turin, where he leads the Smart Personal Technology Lab at ICxT. Before joining the University of Turin, he researched interactive TV systems and ubiquitous technologies for Telecom Italia S.p.A. at the Research and Trends Department. He has authored more than 100 papers in international journals (e.g., TOCHI, *Human–Computer Interaction*, IJHCS, and *Computers in Human Behavior*) and peer-reviewed conferences (e.g., CSCW, CHI, ISWC/UbiComp, and UMAP). His scientific research is situated within the area of human–computer interaction. It focuses on the investigation of the effects of interactive and intelligent technologies on people's everyday lives.

**CATALDO MUSTO** received the Ph.D. degree, in 2012, with a dissertation on Enhanced Vector Space Models for Content-Based Recommender Systems. He is currently an Assistant Professor with the Department of Informatics, University of Bari Aldo Moro. His research interests include the adoption of natural language processing techniques for fine-grained semantic content representation in recommender systems and user modeling platforms. He was an Invited Speaker at the Workshop on Semantic Adaptive and Social Web (SASWeb) at UMAP 2012 and at the first Workshop on Financial Recommender Systems (FINREC 2015). He gave tutorial at UMAP 2016 and has published over 50 papers and served as a Reviewer or a Co-Reviewer in the Program Committee of several conferences in the area as *ACM Recommender Systems*, ECIR, UMAP, and WWW.

**GIOVANNI SEMERARO** is currently a Full Professor of Computer Science with the University of Bari Aldo Moro, Italy, where he teaches intelligent information access and natural language processing and programming languages. He leads the Semantic Web Access and Personalization (SWAP) "Antonio Bello" Research Group. He has been a Visiting Scientist with the Department of Information and Computer Science, University of California at Irvine, in 1993. From 1989 to 1991, he was a Researcher with Tecnopolis CSATA Novus Ortus, Bari, Italy. His research interests include machine learning, AI and language games, recommender systems, user modeling, intelligent information mining, retrieval, and filtering, semantics and social computing, natural language processing, the semantic web, and personalization. He is the author of more than 400 publications in international journals, conference, and workshop proceedings, as well as of two books, including the textbook *Semantics in Adaptive and Personalised Systems: Methods, Tools and Applications* publication by Springer. He is a member of the Steering Committee of the National Laboratory of Artificial Intelligence and Intelligent Systems (AIIS), National Interuniversity Consortium for Informatics (CINI) and the Steering Committee of the ACM Conference Series on Recommender Systems. In 2015, he was selected for the IBM Faculty Award on Cognitive Computing for the Deep Learning to Boost Cognitive Question Answering Project. He was one of the founders of Italian Association for Computational Linguistics (AILC) and on the board of directors, till 2018. From 2006 to 2011, he was on the Board of Directors of the Italian Association for Artificial Intelligence (AI*IA). He has been the Principal Investigator of the University of Bari Aldo Moro in several European, national, and regional projects. He regularly serves in the PC of the top conferences in his areas and is the Program Co-Chair of CLiC-it 2019. Among others, he has served as the Program Co-Chair of CLiC-it 2016 and ACM RecSys 2015 and the General Co-Chair of UMAP 2013. Since 2013, he has been the Coordinator of the Second-Cycle Degree Program in Computer Science at the University of Bari Aldo Moro. He is also the Coordinator of the first edition of the Master's in Data Science at the University of Bari Aldo Moro.