Health Multimedia: Lifestyle Recommendations Based on Diverse Observations

Nitish Nag
University of California, Irvine
Irvine, USA
nagn@uci.edu

Vaibhav Pandey
University of California, Irvine
Irvine, USA
vaibhap1@uci.edu

Ramesh Jain
University of California, Irvine
Irvine, USA
rcjain@uci.edu

ABSTRACT

Managing health lays the core foundation to enabling quality life experiences. Modern multimedia research has enhanced the quality of experiences in fields such as entertainment, social media, and advertising; yet lags in the health domain. We are developing an approach to leverage multimedia systems for human health. Health is primarily a product of our everyday lifestyle actions, yet we have minimal health guidance on making everyday choices. Recommendations are the key to modern content consumption and decisions. Cybernetic navigation principles that integrate health media sources can power dynamic recommendations to dramatically improve our health decisions. Cybernetic components give real-time feedback on health status, while the navigational approach plots health trajectory. These two principles coalesce data to enable personalized, predictive, and precise health knowledge that can contextually disseminate the right actions to keep individuals on a path to wellness.

KEYWORDS

cybernetics; navigation; quantified-self; objective-self; precision medicine; predictive health; prevention; event mining; context awareness; mobile health; persuasion; personalized

ACM Reference format:
Nitish Nag, Vaibhav Pandey, and Ramesh Jain. 2017. Health Multimedia: Lifestyle Recommendations Based on Diverse Observations. In *Proceedings of ICMR '17, Bucharest, Romania, June 06-09, 2017,* 8 pages.
https://doi.org/http://dx.doi.org/10.1145/3078971.3080545

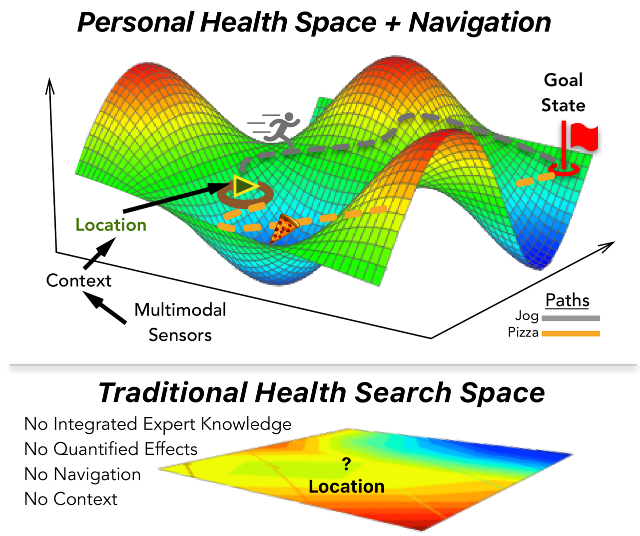

Figure 1: Increased dimensionality (both time and space) of personalized health space can provide actionable navigation towards health goals. As a simple example, although jogging may have a higher effort at the moment (hence the uphill initial slope), it informs the user how this action will reach the goal node. By having the pizza, it may be easier in the moment, but the user will have to climb out of a sink in the map to reach their goal.

1 INTRODUCTION

Inaccessibility and delayed dissemination of medical knowledge for lifestyle guidance persists in all health care systems, while global health burdens continue to rapidly rise. Health care is yearning for a major transformation, from reactive and hospital-centered to proactive prevention; from doctor-centered to patient-centered; from disease-centered to wellness promotion; and a shifting focus from temporary fixes to long-term solutions [23]. Commonly, physicians focus primarily on medical methods to manage health when a patient becomes ill. True health outcomes result from actions taken in every moment and place, not just medical intervention during sickness. Future advancements in health must continuously sense individual needs and rapidly provide the relevant resources so corrective actions ensure health stability.

Health is essentially a product of our genome and lifestyle [28]. Lifestyle is the primary controllable aspect of our health. Unfortunately, lifestyle associated diseases continue to lead the rapid rise of chronic diseases across the globe. Monitoring lifestyle factors has been difficult, and computational power with effective methods to address the needs of each individual have been limited. Trying to produce changes in routine lifestyle habits is also a tremendous psychological hurdle. Media platforms such as food analysis from images or accelerometer based activity sensors have started collecting data about lifestyle factors. These isolated efforts forgo the integration of other media sources and lack synchronization across time and space. Imagine how Google Maps on mobile phones alerts the driver of an upcoming turn. This is the product of real-time

ICMR '17, June 06-09, 2017, Bucharest, Romania

ACM ISBN 978-1-4503-4701-3/17/06…$15.00
https://doi.org/http://dx.doi.org/10.1145/3078971.3080545

multimedia synchronization to power recommendations. Synergistic and synchronized *multimedia* technologies will play a major role in advancing progress to solve health issues.

Current health search systems place the burden of querying on the uninformed user, without understanding any health context. A new era of health multimedia will usher expert knowledge with computational power to drive dynamic recommendations, alleviating the user of querying for answers, and provide the right actions at the right time for users to best manage their health.

1.1 Navigation Systems

A fundamental basis of capturing actionable knowledge rests in maps. Physical maps represent spatial worlds, where as a textbook maps the knowledge of a certain domain. Search engines capture the knowledge of virtual structures. Traversing this spatio-temporal knowledge is the fundamental problem in navigation. We start from an initial node and usually wish to end up at a goal node (Figure 1). In the instance of health, we want to construct the health map of each individual based on knowledge from biological science, but also from sensors that capture the dynamic life history of the person (Figure 5). Multimedia health systems must elucidate the current location of an individual on their personal health map, highlight the optimal goal destination, and illuminate the path from the current status to the goal (Figure 1). Two *key* differences separate this new paradigm from traditional search. First, understanding the context of the individual from multimedia signals maps the location of the user in their health spectrum, similar to a GPS location of a car. Second, we must quantitatively model the effects of future actions on the user while calculating the ability of the user to execute a given action, similar to calculating the various routes and traffic for driving. Given the models of the car, roads, and traffic, we can predict when the car will arrive at the destination. Similarly, we need to be able to predict the course of an individuals health for future health systems.

1.2 Health Issues of the 21st Century

Our paper targets three significant problems in modern health care. First: Health systems largely react to problems, rather than avoiding problems through preventive measures. Second: Gold-standard medical practices depend on "evidence-based" medicine taken from population averages. A lack of individual contextual analysis results in compromised care and sub-optimal outcomes. Third: Access to medical guidance is limited due to poor information dissemination, and restricted physical time and space. When doctors give lifestyle suggestions to patients, they are hard to translate into everyday life decisions. Although the knowledge exists, it is not in an actionable form. Users also do not know what their context means in relation to how their choices are affecting their health (Figure 1). When a need or question arises for health advice, such as "What should I eat?" or "Should I take this medicine now or later?", there is a large delay to receive meaningful assistance. The difficulty of scaling physical systems, like hospitals and personnel, further limits high quality care. Each of these three issues are coated with a vast splattering of emerging media and data streams. Herein lies a great opportunity for multimedia systems.

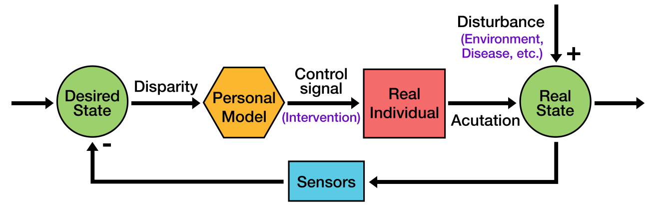

Figure 2: Cybernetic Control pairs the individual user and digital health assistance to enact real-world changes to optimize health.

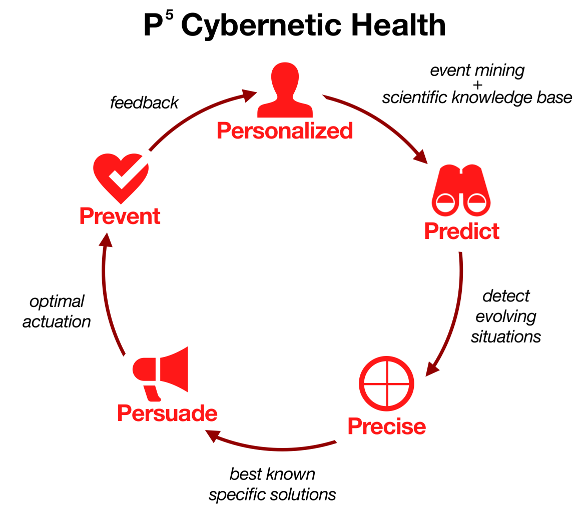

Figure 3: P5 Cybernetic Health coordinates the elements of personalized, predictive and precision medicine through persuasion techniques that result in disease prevention.

1.3 Cybernetic Principles

Cybernetic principles transformed the design of complex systems [32]. Continuous observations/measurements are a key component in closed loop feedback control systems. Airplanes, ovens, and other machines use these feedback loops to safely and efficiently operate. Imagine if the thermometer in an oven gave a reading once every year. How would the oven know to heat or turn off? The thermometer, heater, and other components of the oven must all be coordinated and continuously working for the machine to operate correctly. Similarly, the human body maintains homeostasis in a remarkable array of perturbations. Biological systems have an intricate play of real-time sensors and actuators within the body to do this. Cellular sensors collect personalized information based on the individual body changes. These signals affect outputs for corrective action. These cybernetic mechanisms keep the human body naturally stable (Figure 2).

Occasionally, this human system becomes unstable which results in deteriorating health. In the current medical ecosystem, there is a great delay between deterioration and detection, resulting in further destabilizing of the human biological system and degrading health. Our vision is to reduce this latency in health care through early detection, by continuously collecting sensor data that are specific to the individual, while enabling corrective actions to transform an unstable health condition back to full health stability. For example, anesthesiologists are beginning to develop closed-loop drug administration systems that will replace much of their own job during a surgery [22], and in type 1 diabetics, mechanical devices can replace the biological pancreas with a hormone pump and continuous glucose monitoring. Both of these examples are mechanical implementations of cybernetics for health. The next step needs to combine real world multimedia information sources into a virtual environment that patients can use to interact and understand their live health status.

Modern technology systems intimately link navigation and cybernetic systems. Airplane auto-pilot systems use navigational components to know the route, and mechanical cybernetic systems to help correct the course of travel en-route to the destination. This is only possible through perpetual monitoring and updates. Navigation systems, like Google Maps, became rapidly popular by providing expert knowledge with real-time personalized contextual information to guide travel. We believe this is the future of human health systems, and the time is ripe.

2 P5 CYBERNETIC HEALTH CONCEPT

Multimedia system integration into standard health practices paves a path for future leaps in addressing personalized and preventive health. Leading medical professionals have advocated for a deeper integration of technology into health care [27]. Future health care systems will intertwine lifestyle multimedia data with medical knowledge to develop a new paradigm that optimizes individual health. This is illustrated by research that shows patients make better lifestyle choices that would combat diabetes if given guidance [25]. Ultimately, transforming data and knowledge to actionable lifestyle choices is the most promising, effective, and attainable method to improve human health. We have developed a 5 component system to bring this vision to reality.

First, we use a multi-layer modeling system to understand how to build an increasingly accurate personalized model (Figure 5). Second, using this dynamic model connected with real time sensors allows us to predict evolving situations that an individual may encounter. Third, we use the predictions in conjunction with validated expert medical knowledge to give the most precise solutions to avoid emerging problems. Fourth, we effectively persuade the individual as an actuator in the system by optimizing their preferences, convenience, and specific health needs. Fifth, we give feedback on how the patient's actions have quantitatively affected their health. The realization of this personalized, predictive, precise, persuasive, and preventive system depends upon the coordination of available and future technologies. We call this above approach P5 Cybernetic Health (P5C) (Figure 3).

3 MULTIMEDIA RETRIEVAL IN THE P5C FRAMEWORK

Recommendations are the key to modern content consumption. Recommending effective lifestyle changes requires accurate models that represent the individual. This model will understand interactions

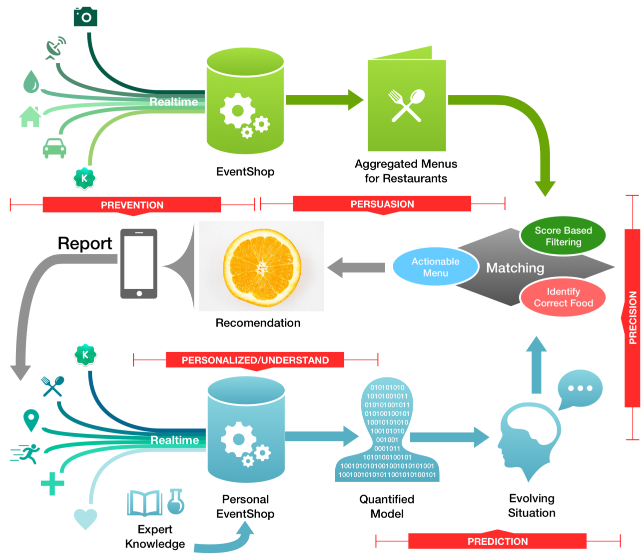

Figure 4: System Architecture of P5C: Quantified individual models allow for matching the individuals needs with the appropriate environmental resources to enable optimal recommendation systems.

of user events with various situations. Combining this understanding with live contextual situations establishes the real-time needs of the individual. We match these needs to accessible resources. These resources are presented as recommendations which move the person towards the optimal health state.

Multimedia work in this field has primarily focused on giving the user figures and statistics of past data. This is true for both hardware and software in personal health. Hardware such as the Fitbit, and health software like Apple's HealthKit only function to acquire and accumulate data. This does not fulfill the function of providing timely and personalized health advice in a predictive manner. Most importantly current digital health mechanisms are rudimentary in detecting context for each individual. Additionally, recommendation engines that are used in health applications ignore mechanisms to maintain retention and trust of the user. Users quickly get notification/alert fatigue from poor recommendations. To sustain users, applications must give users autonomy, cater to their desires and convenience, autonomously track data, while also informing them in an encouraging manner. Furthermore, many lifestyle data parameters are gathered through manual mechanisms. For example, popular nutrition tracking apps ask users to manually enter information, placing high burden on the user. This further causes a loss in user retention, while having poor fidelity of input values. Users desire low data entry burden, with high functionality to help reach their goals. All of these problems open vast opportunities for multimedia research and application.

Different components of P5C enable us to accomplish this task by creating a layered architecture where different layers are responsible for completing the cybernetic navigation of health (Figure 4).

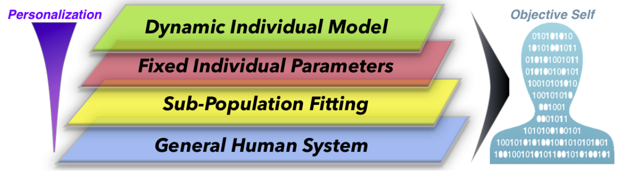

Figure 5: We integrate increasingly personalized knowledge layers derived from multimedia data fusion to build a comprehensive health map of each individual.

3.1 Personalized Model : Objective Self

An objective self threads all pieces of life data to make a true universal multimedia map of the human life. This becomes a hub for integrating sensors and methods to enhance the personal multimedia model. In the future, as computer science, medicine and other fields progress, there needs to be an effective method to combine these into increasingly high resolution health map layers.

The first step in making effective recommendations is maintaining a personicle, a complete and quantified log of the events in a person's life[14]. The personicle records all user activities on a time line and we can use this to model the user's current status. We refer to this model of the person as Objective Self, which allows us to determine the future state of the person based on the interactions between their current state and possible future events. Objective Self contains the user's current state along with their preferences. In our map analogy, this would represent the current location of the user and the topography of the map, which represent the user preferences for different types of events (Figure 1). High cost events are further away from the users preferences, while low cost events are aligned with their preferences. This cost is mapped by the slope of the health topography, with higher cost mapped as a higher slope (Figure 1). The Objective Self model of the person would allow us to find the quantitative effect of future events which are possible from the given state. Creation of Objective Self requires the integration of various life logs which are maintained as a part of the P5C system.

In health care applications, the most effective changes in a person's life can be brought about by inducing changes in controllable actions such as lifestyle, food habits, or environment. The emotional response of the user to different activities or events are relevant for letting us estimate the user's preferences. We maintain a log of the activities, dietary habits, environment, and emotional state of the user using primary sensors (such as mobile phones and wearables). Future health multimedia logging must be done passively without placing any overhead on the user or interfering with their usual activities.

Events in our daily lives and can be logged passively to create the activity log of a person [2, 6, 19, 20]. Personalized life log raw data is then extracted to semantic-level activities in real-time. These real-time semantic events allow for intuitive analysis of a dynamic human life. Realistic systems may never completely capture everything due to the lack of standard data formats and the complexity of life [24]. Therefore, the range of semantic activity also needs to be refined to a number of meaningful universal activities. Quantitative information of time, frequency, intensity, stress, enjoyment, and other affective states are most meaningful to medical researchers [17]. We currently target 17 standard semantic level activities including: socializing, relaxing, prayer, eating, exercising, home events(watching TV, preparing food, sleeping, housework), shopping, conversations, computer/e-mail/Internet usage, working, and commuting. We can extract additional information about events using existing expert systems and use these as features in the Objective Self models. For example, from an exercise event which represents jogging for 30 minutes, we can get the amount of calories and fat burned using expert systems such as Wolfram Alpha[33]. We can obtain similar features and measures for other activities as well and use expert systems to generate tags about the nature of these activities such as "strenuous" or "resting".

Nutritional intake can be captured by analysis of photos, transactions, or conversations of food to create a quantitative nutrition diary. Using pictures, open source deep-learning techniques by Google and Clarifai for food recognition in photos [4] can give us the semantic information about nutrition. Once we have the name of the dish and ingredients, we can look up the item in publicly available food databases such as USDA Food Composition Databases [30] and obtain the quantity of each nutritional group in the food item as well as features, tags, and relevant flavors related to the ingredients. Nutritional parameters for each person are then analyzed based on their personal health status [11, 34] to give them the most relevant food suggestions.

Environmental factors continuously affect the health of every individual. The ability to measure the local environment of each individual gives us insights into how they are being affected by factors the user is otherwise unaware of. Air and water pollution have been shown to increase the risk of asthma and diabetes [5] While public data regarding the quality of the environment is readily available, it is not incorporated and tracked at the individual level. We are using an open source software platform, called *EventShop* to ingest and assimilate different data streams [21, 26]. The real time synchronization of personal activities/events and the environmental factors allows P5C to respond to changing situations in a swift and effective manner.

The Objective Self could be viewed as the person's coordinates on a health map (the dimensions are represented by the health parameters which we are trying to estimate and optimize such as weight, blood glucose etc.), and the set of probable events obtained from the personicle represent the possible paths for the user. The desired state for the user can be represented in the same map, and the problem of reaching the goal node would be similar to finding the path with the least amount of work /resistance. In our analogy, the resistance is represented by the user's preferences and the prevalence of the event sequence in the personicle. It is possible that the shortest path to the destination may not be the best or the optimal path due to the terrain, similarly any event which has rarely appeared in a person's event history or is against their preferences is unlikely to be an actionable recommendation, even if it has the most positive impact on health. Thus in the precision layer of our system, we need to recommend an action by optimizing the health of the user under the constraints placed on them by their history and preferences. We can define the user's preferences by attaching tags with every event in the personicle, which are representative of the event's parameters. For instance, a food event may have tags related to the taste and ingredients of the dish as well as describing the image concepts obtained from deep learning based image recognition systems. Prevalence of a tag is an estimate of the user's preference of the event aspect represented by the particular tag. We can use use these preference scores to rank the set of possible events by matching the tags associated with the events and their respective preference scores.

It is vital that the system should allow for seamless integration of latest developments in AI and new sources of media in creating the Objective Self model. The event based processing model allows us to do that as long as we can convert the new source of data into an event stream, thus allowing us to adapt with changing modalities of data while preserving the existing infrastructure.

3.2 Prediction : Personalized Probability of Events

We have talked about viewing the recommendation problem as a navigation system, and how it is analogous to suggesting optimum routes considering terrain and distance. The first step in identifying the optimum path or activity, is to find out all the possible activities and estimating their influence on the health of the user so that we can identify desired activities. This is similar to comparing the available paths in the map and finding out the ones which bring us closer to our destination (Figure 1).

We use data driven analysis and pattern mining algorithms to find event patterns in a personicle and predict events that are most likely to appear next in the sequence. Event relationship operators formulate compound events and compute co-occurrences which are then tested with a new set of data [13, 15, 16]. This framework extends traditional complex event processing [3] significantly by including space, multiple event streams, with point and interval events to enable real world data analysis. Every event has a set of constraints as prerequisites which must be fulfilled for the event to occur. Given an environment, we filter out events whose prerequisites are not satisfied (meaning, they would not be executed). This gives us the set of next possible events based on the person's history. Based on how these events affect user's health, we recommend actions which induce a positive change.

Merged event streams in a personicle [14] produce a stream of time-indexed events, e.g., high fat meal eaten at 3pm Monday / 40 minutes of exercise at 2pm on Saturday / high blood glucose level at 5pm Thursday. Statistical models [12] are used to identify recurring patterns in a sequence of events. By fitting such models we can identify sequences of events that have high co occurrence with an adverse medical event, allowing prediction of health status in the future. Fitting the model to an individual's event stream data is a challenge that may require weeks of observations. Bayesian hierarchical models [10] can be used to leverage information from a population of users to give upfront meaningful analysis until the data from a single user is sufficient. This approach provides an intuitive data-determined degree of synergistic sharing between individual and population information. Parameters of models that fit separate individuals can be described by a population distribution where recurring patterns are shared while some remain unique to each individual [12].

We need to have a calibration period to train the models which can predict the effects of an event on the person's health. We need to create features which represent the inherent health aspects of the events, for example a food event may have features representing the nutritional content of the dish, and we may want to find an association between different dishes and the person's blood sugar levels. As described in the previous section, these features may come from deep learning and AI based systems or expert knowledge based systems and will be used to build models which can estimate the relevant health parameters. For example, in section 4, we talk about an application (HealthButler) for managing type II diabetes. In the calibration phase, the user may be required to measure their blood glucose levels at periodic intervals of time, we can use the features obtained for events (nutritional values, calories burned etc.) to build a model for estimating the change in health parameters (blood glucose level) due to these events. Once we have these models, we can use them to predict the effects of different events on the user's health to give healthy recommendations.

Aside from this, we can use the event history of the user to identify anomalies in their behavior. Some of these anomalies may represent medically significant behavior changes, for example if we can identify sudden mood shifts using an emotion log, there is a high likelihood of hyper or hypoglycemia. This activates the system to provide emergency relief services and information. We also use the personalized data to predict developing insulin resistance over time. An accumulation of low activity, high fat and sugar foods, with associated lethargy indicates the individual is not on track to improve their insulin sensitivity.

The prediction layer enables us to find the next set of events and how these may affect the user's health, the next step is to match these events with the user's preferences and obtain a ranked list of actionable items which the user is likely to follow.

3.3 Precision : Constraint Optimization with Knowledge and Context

A lack of data on personal lifestyle in relation to biomarkers has been a struggle in the quest to provide precise treatments for patients [31]. President Barack Obama also began the Precision Medicine Initiative to follow various cohorts of patients to understand what constitutes better treatments for different people [8]. Researchers are also trying to link genetic factors to diabetes outcomes, but they are confounded in their research due to a lack of quality lifestyle data [29].

By predicting the probability of physiologic events, we can send the most appropriate control signals to the user to take corrective action to maintain their health status before it starts becoming unstable. Algorithms should incorporate various factors such as the severity of adverse events along with the likelihood of occurrence. To prevent alarm fatigue, we dispatch the control signal only after the threshold of maintaining optimal health is crossed. Otherwise the control signal of actions are given as choices/recommendations to help guide the user towards better health goals. Precise diagnostic tools, medications and other medical interventions are also suggested to the physician as to reduce waste of resources and ensure better outcomes. Most importantly, giving the most precise treatment for an individual relies on the generation of actionable interventions [31].

Generating precise recommendations refers to filtering out and ranking the available events based on the context of the user and expert knowledge. Once the next set of possible events have been generated using event mining and the environment log, we can then match constraints placed on these events with the context of the user which is represented by their Objective Self. This includes matching the event parameters with the preferences of the user and using the appropriate domain knowledge. As mentioned earlier in section 3.1, the events have tags associated with them which are related to different aspects and parameters of the events. The Objective Self model stores the preferences of the users by assigning a weight to the tags which are present in the user's personicle and also can include the emotional response of a user to these events. Assigning a utility/preference score to these events by using a similarity metric (eg. Cosine similarity) to maximize recommendations for the events which align with user's preferences. Continuing our map analogy, this step is finding out the work required along the different routes (terrain is represented by the user's preferences), thus it gives us the most convenient route to reach our goal (Figure 1).

3.4 Persuasion : Enabling User Execution

The goal of our system is to induce gradual habit changes via suggestions which cater to user's preferences and push sustainable long-term positive habits and health. Executing this health navigation in the real world is only done when the user takes action on a given recommendation. Persuasive multimedia systems in advertising, social media, news, and other platforms are constantly under development to push people to take a certain action. Concepts from persuasive technology generate recommendations that are most suited for the individual[1, 9]. We use Fogg's Behavior Model [7] to optimize three factors which strongly determine behavior: motivation, ability and trigger.

Motivation captures the individual's willingness to follow through with a suggestion. The list of recommendations are ordered on the basis of the preferences of the user which are determined from Objective Self. This ensures that the recommended actions have the highest motivation weight.

Ability encompasses the resource availability in performing the suggested behavior. This factor is an interplay of the individual's environment and their intrinsic capability, and must ensure those constraints are met.

Triggers are instances of events that lead to accomplishment of a task[7]. These can be optimized by using the current spatio-temporal context of the user to filter the list of recommended actionable items. For example, a notification about healthy food options when the person is hungry. The efficacy of a trigger is maximized when it is generated according to the context of the user and the current situation. For example, a recommendation suggesting a cup of coffee may be very relevant at 9 AM but it is unlikely to be effective at 7 PM, even if the user prefers coffee and has no obvious constraints placed.

Thus by combining the right triggers with context, we can maximize the effectiveness of positive health actions.

3.5 Prevention : The Ultimate Goal

P5C systems prevent adverse outcomes through addressing several points. First, multimedia experiences (such as through an app) give direct feedback to quantitatively understand personal health status. Second, user ownership of their total integrated health multimedia data is key to empower patient participation in health and driving user engagement. Third, Multimedia systems can inform the individual of risks regarding their choices. In the USA, calories are printed on menus of large franchise restaurants by law. This allows consumers to have direct basic information on what they are consuming, but remains a rudimentary form of information dissemination. Similarly, the practice of having cancer warning signs on cigarette packages or labeling alcohol warnings for pregnant women are designed to inform the consumer of their choice. Future systems will use augmented reality systems to inform users in how real world items or actions will quantitatively affect their health. Systems derived from the P5C concept must focus on informing the user in a personalized fashion for *all* health related decisions. Ultimately, the prevention of deteriorating health conditions keeps the individual in the steady state of optimal health: original cybernetic goal.

4 HEALTHBUTLER APPLICATION

To illustrate the delivery of P5C, a type 2 diabetic named "Bruce Uberschweet", uses the HealthButler (HB) mobile app system to optimally manage his health condition. This will include analysis to better control blood sugar and reduce drug dependency through improved metabolism [18]. Bruce is looking for lunch on his commute to work on Monday and knows that HB always gives him the fastest access point to tasty and nutritious food. An intelligent recommendation engine takes into consideration his real-time personal tastes, logistical convenience, and current health needs (including blood sugar) to provide him a curated list of specific dishes that he can easily pre-order. He can also clearly see how each dish affects his diabetes so he can feel empowered to choose what is good for himself (Figure 8). After attending a wedding on Sunday, HB predicts a rising insulin resistance based on his life log data, and gives him immediate actions to take in order to address the worsening condition. It tells him when to take his blood sugar, gives him a clinical mono-filament test with a tutorial video that he can do at home, and simplifies booking an urgent appointment with his doctor to change his medication dose (Figure 8). These are two examples of how HB is actively engaged in predicting Bruce's health status from his health data (Figure 7), merging in with his daily life in an unobtrusive and useful way. Bruce can actively see how his external world and internal body are interacting through HB. Figure 6 shows the flow of information from the smart phone to the logs and how the raw media is converted to event streams and stored as logs.

5 RESEARCH OPPORTUNITIES

P5C paves new research pathways in which multimedia recommendation systems can take advantage of health data to benefit users.

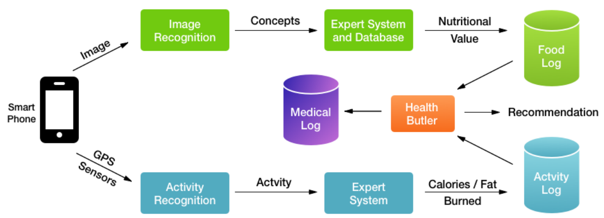

Figure 6: Flow of information in the Health Butler Application, with respect to activity and food event streams.

Figure 7: Left: Aggregate summary of health status Center: FoodSense Log. Right: MedicalSense Log.

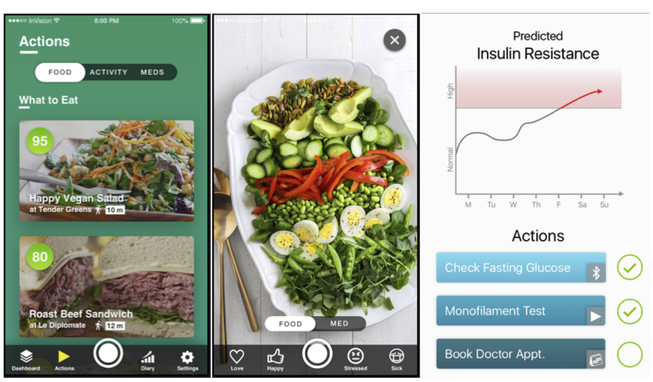

Figure 8: Left: Nutrition guidance is catered to the user's preferences, needs, and available resources. Center: Easy one-touch food and mood tracking. Right: real-time health status is shown, with direct actions to take for help.

These applications can range from acute, sub-acute, and chronic disease health management to preventive medicine (such as vaccines) and wellness. These opportunities also allow open more doors for collaborative efforts across various fields.

5.1 Diverse Multimedia/Multimodal Data

Health applications already use traditional multimedia data sources: text, images of different types, video, and audio. The more exciting and challenging part is the use of sensors that create data streams ranging from GPS and accelerometers to glucose measurements, heart-reate monitoring, blood oxygen level and many more. The list of different multimodal measurements will keep increasing. This offers a great opportunity to the multimedia research community to address time-sensitive data streams of diverse modality and combine their semantics to solve the most important problem for humans, health. Developing P5C to universally digest and synchronize various sensor data streams will continue to be important for human health and will be a rewarding challenge for multimedia technology.

5.2 Personal Logs and Relevant Applications

Life logs and related research areas have attracted multimedia researchers. Much activity in this area has remained limited to video sensors and determination of simple activities. As discussed above, now we need to associate everyday activities to different measurements related to behavior and health data obtained from diverse sensors. Such logs are essential for building models that will be used for personalization and prediction. Since we are dealing with diverse sensors, synchronization and semantics become a very important issue in analyzing these logs. Data is collected from different sensors at different frequencies. Synchronizing them requires the understanding of their semantics as well, which becomes a new challenge.

5.3 Multimodal Retrieval and Recommendation Approaches

Data from various streams need to be analyzed to detect events during the daily activities of a person. All these data streams will be stored 24/7 for the lifetime of a person. Since this data needs to be analyzed to build models of the person as well as to understand current events for predictive purposes, it is important to develop techniques for retrieving different events in these streams in context of various queries. Moreover, in most cases, the result of such queries is not limited to presenting those events, but to relate those events to personal models for predictions based on medical knowledge and provide appropriate actions as recommendations. As augmented and virtual reality platforms become more ubiquitous, future health multimedia systems can start overlaying health information directly on objects in our real world. For example, if there is a food dish with peanuts at the grocery story, it will be highlighted with a warning color in augmented reality glasses. Facebook and other companies are already working on general systems like this, but need health data to effectively power such a system. This is a new challenge for multimedia research where results from a retrieval are continuously converted to recommendations based on the situation of the user.

6 CONCLUSIONS

Cybernetic navigation principles lay the foundational bricks of building health systems that are responsive in keeping individuals at optimal health over the course of a lifetime. P5C is based on the most fundamental principles in cybernetics while incorporating the ability to seamlessly integrate both present and future technological advancements. The most important concept in this paper is how synergistic health multimedia fusion will fuel the recommendation engines of the future. These actionable recommendations enable the human actuator in cybernetic health to form a true closed-loop feedback system.

Real human health progress relies on an interdisciplinary effort between hospital clinicians, engineers, computer scientists, bioscience researchers amongst other fields. P5C systems will sow the seed of this interdisciplinary approach within future health systems will accelerate multimedia research progress into real systems that benefits patients.

REFERENCES

[1] S. M Arteaga, M. Kudeki, and A. Woodworth. 2009. Combating obesity trends in teenagers through persuasive mobile technology. *ACM SIGACCESS Accessibility and Computin* 94 (2009), 17–25.
[2] G. Biegel and V. Cahill. 2004. A framework for developing mobile, context-aware applications. In *Pervasive Computing and Communications, 2004. PerCom 2004. Proceedings of the Second IEEE Annual Conference on*. IEEE, 361–365.
[3] A. Buchmann and B. Koldehofe. 2009. Complex event processing. *IT-Information Technology Methoden und innovative Anwendungen der Informatik und Informationstechnik* 51, 5 (2009), 241–242.
[4] Clarifai 2017. Clarifai Food Recognition. (2017). http://blog.clarifai.com/.
[5] I. C Eze, M. Imboden, A. Kumar, A. von Eckardstein, D. Stolz, M. W Gerbase, N. Künzli, M. Pons, F. Kronenberg, C. Schindler, and others. 2016. Air pollution and diabetes association: Modification by type 2 diabetes genetic risk score. *Environment International* 94 (2016), 263–271.
[6] D. Ferreira, V. Kostakos, and A. K Dey. 2015. AWARE: mobile context instrumentation framework. *Frontiers in ICT* 2 (2015), 6.
[7] B. Fogg. 2009. A behavior model for persuasive design. *Proceedings of the 4th International Conference on Persuasive Technology - Persuasive '09* (2009), 40:1—-40:7.
[8] J. E Fradkin, M. C Hanlon, and G. P Rodgers. 2016. NIH Precision Medicine Initiative: Implications for Diabetes Research. *Diabetes Care* 39, 7 (2016).
[9] T. Fritz, E. M Huang, G. C Murphy, and T. Zimmermann. 2014. Persuasive Technology in the Real World: A Study of Long-Term Use of Activity Sensing Devices for Fitness. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (2014), 487–496.
[10] A. Gelman, J. B Carlin, H. S Stern, and D. B Rubin. 2014. *Bayesian data analysis*. Vol. 2. Chapman & Hall/CRC Boca Raton, FL, USA.
[11] S. Goyal, P. Morita, G. F Lewis, C. Yu, E. Seto, and J. A Cafazzo. 2016. The Systematic Design of a Behavioural Mobile Health Application for the Self-Management of Type 2 Diabetes. *Canadian Journal of Diabetes* 40, 1 (2 2016), 95–104. https://doi.org/10.1016/j.jcjd.2015.06.007
[12] K. Heins and H. Stern. 2014. A Statistical Model for Event Sequence Data.. In *AISTATS*. 338–346.
[13] L. Jalali. 2016. Interactive Event-driven Knowledge Discovery from Data Streams. (2016).
[14] L. Jalali, D. Huo, H. Oh, M. Tang, S. Pongpaichet, and R. Jain. 2014. Personicle: Personal Chronicle of Life Events. In *Workshop on Personal Data Analytics in the Internet of Things (PDA@ IOT) at the 40th International Conference on Very Large Databases (VLDB), Hangzhou, China*.
[15] L. Jalali and R. Jain. 2015. Bringing deep causality to multimedia data streams. In *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 221–230.
[16] L. Jalali, H. Oh, R. Moazeni, and R. Jain. 2016. Human Behavior Analysis from Smartphone Data Streams. In *International Workshop on Human Behavior Understanding*. Springer, 68–85.
[17] D. Kahneman, A. B Krueger, D. A Schkade, N. Schwarz, and A. A Stone. 2004. A survey method for characterizing daily life experience: The day reconstruction method. *Science* 306, 5702 (2004), 1776–1780.
[18] R. W McGarrah, C. A Slentz, and W. E Kraus. 2016. The Effect of Vigorous- Versus Moderate-Intensity Aerobic Exercise on Insulin Action. *Current Cardiology Reports* 18, 12 (12 2016), 117.
[19] H Oh. 2015. *A Life Event Detection System Using Real Time Heterogeneous Context Monitoring and Formal Concept Analysis*. Master's thesis. University of California, Irvine, Irvine,CA.
[20] H. Oh, L. Jalali, and R. Jain. 2015. An intelligent notification system using context from real-time personal activity monitoring. In *Multimedia and Expo (ICME), 2015 IEEE International Conference on*. IEEE, 1–6.
[21] S. Pongpaichet, V. K Singh, M. Gao, and R. Jain. 2013. EventShop: recognizing situations in web data streams. In *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 1359–1368.
[22] J. Rinehart and C. Canales. 2015. Closed-Loop Pharmacology in Anesthesia and Critical Care. *International Anesthesiology Clinics* 53, 2 (21 2015), 91–101.
[23] M. Sagner, A. McNeil, P. Puska, C. Auffray, N. D Price, L. Hood, C. J Lavie, Z. Han, Z. Chen, S. K. Brahmachari, and others. 2016. The P4 Health Spectrum âĂŞ A Predictive, Preventive, Personalized and Participatory Continuum for Promoting Healthspan. *Progress in Cardiovascular Diseases* (8 2016).
[24] A. J Sellen and S. Whittaker. 2010. Beyond total capture: a constructive critique of lifelogging. *Commun. ACM* 53, 5 (2010), 70–77.
[25] D. Sherifali, V. Viscardi, J Bai, and R M. U. Ali. 2016. Evaluating the Effect of a Diabetes Health Coach in Individuals with Type 2 Diabetes. *Canadian Journal of Diabetes* 40, 1 (2 2016), 84–94.
[26] M. Tang, P. Agrawal, and R. Jain. 2015. Habits vs environment: What really causes asthma?. In *Proceedings of the ACM Web Science Conference*. ACM, 30.
[27] E. Topol. 2015. *The patient will see you now: the future of medicine is in your hands*. Basic Books.
[28] E. J Topol. 2012. *The creative destruction of medicine: How the digital revolution will create better health care*. Basic Books.
[29] Type 2 Diabetes Genetics 2017. Type 2 Diabetes Genetics. (2017). http://www.type2diabetesgenetics.org/.
[30] USDA Food Composition Database 2017. USDA food database. (2017). https://ndb.nal.usda.gov/ndb/search/list.
[31] W. M. Valencia and H. Florez. 2017. How to prevent the microvascular complications of type 2 diabetes beyond glucose control. *BMJ (Clinical research ed.)* 356 (1 2017), i6505.
[32] N. Wiener. 1948. *Cybernetics: Control and communication in the animal and the machine*. Wiley New York.
[33] WolfromAlpha 2017. Wolfram Alpha. (2017). https://www.wolframalpha.com.
[34] F. G Wood, E. Alley, S. Baer, and R. Johnson. 2015. Interactive Multimedia Tailored to Improve Diabetes Self-Management. *The Nursing clinics of North America* 50, 3 (9 2015), 565–76.