Published in final edited form as:

IEEE J Sel Top Signal Process. 2016 August ; 10(5): 962–974. doi:10.1109/JSTSP.2016.2565381.

Leveraging Multi-Modal Sensing for Mobile Health: A Case Review in Chronic Pain

Min S. Hane Aung

Department of Information Science, Cornell University, Ithaca, NY 14853 USA

Faisal Alquaddoomi and Cheng-Kang Hsieh

Department of Computer Science University of California, Los Angeles, CA 90095 USA

Mashfiqui Rabbi

Department of Information Science, Cornell University, Ithaca, NY 14853 USA

Longqi Yang [Member, IEEE], J. P. Pollak, and Deborah Estrin [Fellow, IEEE]

Jacobs-Technion Cornell Institute, New York, NY 10011 USA

Tanzeem Choudhury

Department of Information Science, Cornell University, Ithaca, NY 14853 USA

Abstract

Active and passive mobile sensing has garnered much attention in recent years. In this paper, we focus on chronic pain measurement and management as a case application to exemplify the state of the art. We present a consolidated discussion on the leveraging of various sensing modalities along with modular server-side and on-device architectures required for this task. Modalities included are: activity monitoring from accelerometry and location sensing, audio analysis of speech, image processing for facial expressions as well as modern methods for effective patient self-reporting. We review examples that deliver actionable information to clinicians and patients while addressing privacy, usability, and computational constraints. We also discuss open challenges in the higher level inferencing of patient state and effective feedback with potential directions to address them. The methods and challenges presented here are also generalizable and relevant to a broad range of other applications in mobile sensing.

Keywords

Activity monitoring; affective computing; audio sensing; behavioral signal processing; chronic pain; face expression; mobile health; mobile sensing; modular architecture; self-reporting; smartphones; survey; wearable technology

I. Introduction

The unprecedented spread of smartphone use across the world has led the development of new mobile systems for health measurement, analysis and intervention or *mHealth* [1],[2]. In this context, the smartphone can be viewed as a multi-sensor device which serves as a continuous monitor capable of informing both clinically-relevant inferences with efficacious and timely patient-feedback. Given the enormity of the number of potential applications

within *mHealth*, in this paper we focus on the area of pain management. This choice of pain as a case study will help to demonstrate open but surmountable challenges when aiming to implement effective *mHealth* systems; we expect that many of the implications for system design are transferrable to other patient-centric health applications.

Pain is a complex experiential phenomenon; it has a range of root causes and can impact and manifest in many domains such as emotion, cognition, socialization, function as well as behaviour [3]. Intensity is subjective and as such it is difficult to tangibly detect in any direct sense. The broadly accepted definition by the International Association for the Study of Pain is principally a generalized catchall and does little to clarify it as an objective signal: “*an unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage*” [4]. Moreover, recent commentaries state that clinical taxonomies for pain are overall not well ordered and some even illogical [5]. Having said this, two known components of pain: negative psychological states (e.g. anxiety and fear) and resultant *pain behavior* (e.g. limping, grimacing, avoidance) [6], [7] have been widely researched in observation-based behavior studies. The form and characteristics of different *pain behavior* have more explicit definitions often containing specific physical terminology [7]–[9]. A particular type of pain behavior that is regularly assessed in clinical practice is *pain interference*, in which activities are disrupted or avoided due to pain. From a signal processing perspective it is these outward manifestations of *pain behavior*, including interference, which are more tangible and convey the expressions of pain experience and serve as meaningful target for *mHealth* systems.

The emergent fields of behavioral signal processing (BSP) and affective computing (AC) [10], [11] have demonstrated that complex human behavior and psychological states can be inferred from multi-modal data, principally from audiovisual [12] and physiological sensing [13]. Such studies have made valuable contributions in the understanding of signal processing and recognition methodologies for inference of human behavior (for an overview of multimodal pain behavior recognition see [14]). However, there remains systemic challenges when using mobile systems, challenges that are not apparent in typical laboratory settings with specialized apparatus and control. In terms of modalities, highly informative internal bio-signals such as galvanic skin response, electroencephalography, electromyography and heart rate are not directly measured from broadly-adopted smartphones; and are still relatively poorly measured on wearables. Also for audiovisual data, artefacts relating to noise, motion, occlusion or crosstalk cannot easily be controlled and minimized in unconstrained real word situations. Furthermore, the reliability and completeness of data streams may also be less than ideal when using a mobile device, depending on the condition and habits of the user. On the other hand, these shortcomings only have critical implications for *spatiotemporally small scale behaviors* (e.g. a facial grimace or changes in vocal pitch or heart rate due to pain conditions); typically the performance of such systems have a high dependency on granularity and data quality. In contrast, for *spatiotemporally large scale behaviors* (e.g. the amount and pace of walking, or hours spent out of the house, on a daily or weekly basis) the aforementioned shortcomings are of far less concern due to data being drawn over greater space and timescales.

As mentioned, we focus on pain as a case study. The justification for this is in part that *spatiotemporally large scale behaviors* are especially relevant to chronic pain (CP). For example, the tracking of adherence to daily exercise regimes over the span of weeks or months is valuable information for chronic back pain rehabilitation [15]. Typically such management regimes are on-going and self-led with many of the required tasks done outside of clinical settings and away from health experts. For such situations, *mHealth* systems and/or smartphones are a natural solution as they are habitually carried in close proximity and are in constant use. This allows for measurements, inferences and feedback to be processed locally, continuously, and over long periods of time.

However, CP management is multifaceted; it's success is dependent on a multitude of physical, cognitive as well as socioeconomic factors [16], [17]. Therefore, if one aims to develop a *mHealth* system for CP management, there is a particular need to accommodate a range of patient generated data. In measurement terms, this means a combination of both *active* sensing which requires proactive action from the user (e.g. self-reporting, diary logging, interactive assessments) and *passive* sensing (e.g. geo-locating, activity recognition, audio). Following this, the correct and efficient parsing of informative descriptors from low level measurements is required. These descriptors serve as input to models designed to infer high-level behavioral patterns. Finally, these inferences can be used for the correct and timely issuance of feedback and/or interventions that are clear, persuasive and actionable by the patient or clinician; the effects of which can then be further measured and the process repeated. We outline this with a tripartite functional loop: *measurement–inference–management* as a basis for system design.

Although we address pain as an over-arching application domain, there still remains a wide range of sub-domains even within CP conditions. Each sub-domain has specific clinical and technical requirements and therefore a comprehensive discussion of all potential systems for every pain condition would be cumbersome, repetitive and out of the scope of this discussion. Therefore to efficiently illustrate requirements and challenges from a systems perspective, we highlight within the following (Section II) example cases of different sensing modalities, with inference to a relevant behavioral outcome and how this inference can be used for management; each following the *measurement–inference–management* process. Section II ends with examples of modular software platforms that are critical to take multiple sensing streams from measurement to meaning in a reusable and evolvable manner. Section III discusses open challenges for higher level inferencing, including temporal considerations, and an overview of potentially useful advances in pattern recognition. Also discussed are factors in improving the usability and persuasiveness from a user's perspective. Finally, we conclude in Section IV.

II. Signals and Signal Specific Processing

We focus on physical activity, acoustics, image capture, and self-reporting as four foundational data streams. While it is known that other signals are also indicative of pain and pain related emotions such as physiological data [18] and that other mobile devices are gaining in both penetration and range of sensing modalities (e.g. the Empatica E4 Wristband which measures blood volume pulse, heart rate variability, galvanic skin response and

peripheral skin temperature); the methodologies and approaches we review should provide a template for their handling and the integration of future data streams from mobile apps and wearable devices.

A. Physical Activity

The continued monitoring of physical activities away from clinical settings is essential for understanding progress in pain reduction, management, and healing, whether it be in the context of post-treatment follow up, or long term physiotherapy[19]–[21]. Such monitoring was traditionally achieved through specialized wearable devices [22], [23], or special-purpose instruments deployed in clinical environments [24]. However, activity tracking on broadly deployed smartphones has become a popular and affordable alternative. The broad adoption and inherent utility of smartphones for end-users means that for many day to day measurements, no additional devices need to be purchased, carried, or forgotten by the user [25]–[27]. Consumer wearables in the form of fitness trackers and smart watches have similar advantages and typically transmit data via the smartphone. The activity sensors embedded in wristbands and phones are largely identical, and so much of this section's discussion is applicable to both.

Smartphone-based activity tracking typically uses the continuous data collected by the phone's accelerometer to determine mobility states. The standard set of states are: sedentary (or stationary), walking, running, or in vehicle. In addition, the location information captured by the phone's GPS or derived from WiFi signals is often used to infer the locational context in which the activities were taking place. There is also the potential for further contextualization to the standard states from low level sensor data that has implications for pain management. The degree of incline of a particular route taken can be determined using either topographical information from maps or from accelerometer and barometer values. This informs the understanding of difficulty level and energy expenditure of a route; the type of surface on a high incline path whether it is stepped or a smooth slope will also have a bearing on difficulty and could potentially be inferred from the inertial and barometric readings. This further contextual information is important for musculoskeletal conditions where posture and speed are relevant factors. Also for speed and step count inference GSM traces can also be utilized [28]. However, such detections are currently not done with standard commercial systems.

Case Example: An Activity Based App for Rheumatoid Arthritis (RA): RA is characterized by episodes of flaring of joint inflammation. Identifying flares early and minimizing their intensity and duration is important for long term management because each episode of flare contributes to further joint damage. In day to day management patients struggle to understand what environmental and behavioral factors contribute to the triggering and duration of flares. While directly and continuously measuring pain is not currently possible, proxies for pain can be found in their physical function and daily activities [29]. In the following, we describe a four-stage inference algorithm from which raw sensor data generates high-level behavioral information that is useful in RA management. This algorithm has been deployed in an open-sourced activity tracking app, called *Mobility* [30]. The clinical validity of the derived measures is currently being evaluated within RA, as well

as in other orthopedic contexts such as surgical recovery. As the validation proceeds and is refined over time for various populations and disease states, there will be a need to modify and tune the specific derived measures, further motivating the need for modular architectures such as those described later.

Stage 1–Instantaneous mobility state inference: a decision tree generated using the C4.5 algorithm was developed to classify tri-axial accelerometer data collected over a short period of time into a set number of mobility states. For example, it is common to extract descriptive features from accelerometer data from a 1 second window sampled every minute to determine the mobility state for that minute. Commonly used features include variance of the acceleration magnitudes and the Fourier coefficients of the acceleration magnitudes [31]. The classifier is typically pre-trained with a large number of labeled training data collected from multiple users. More recent smartphone operating systems, however, have made the instant-based mobility state inference function available as a system application programming interface (API), and the use of these APIs are usually recommended as they are able to leverage special built-in hardware for more efficient continuous sensing (e.g. iPhone's M7 motion sensing chip) than an app running in the application-level can achieve.

Stage 2–Activity segmentation: the instantaneous mobility states inferred by Stage 1 are, however, not particularly useful in themselves. Durational context needs to be added, for example, compared to the information: *John was walking at 6:01 AM*, the information: *John walked for 15 minutes from home to subway station from 6:00 AM to 6:15 AM* is more meaningful in understanding the user's behavior. At this stage, a segmentation algorithm is developed to segment a series of mobility states into *activity segments*, each of which represents a period of time in which the user maintains a mobility state. However, one important requirement is to take into account the potential errors and uncertainty contained in Stage 1's output. For example, signals of being in a vehicle that slows down momentarily for a stop sign could be mistakenly classified as *sedentary* rather than its true state, *in vehicle*. The segmentation algorithm needs to take the uncertainty into account and infer the activities that are consistent with the user's real behavior.

A well-known technique that can incorporate this uncertainty is *hidden Markov models* (HMM) [32]. HMM is a generative probabilistic model in which a sequence of observable variables X is generated by a sequence of internal hidden states Z using an emission probability function. In this case, mobility states inferred in Stage 1 are taken as the observable variables X and assume they were generated by the user's true but unobservable activity states Z . Uncertainty in Stage 1's inference results is encoded in the emitting probability function $P(x/z)$, and the *Baum–Welch* algorithm is used to infer the maximum likelihood of activity states Z . Then, a sequence of consecutive activity states Z that have the same maximum likelihood state will form an activity segment.

Stage 3–Location association and correction: at this stage each activity segment is associated with the location data based on the collection time and conduct error correction. A location data point is composed of longitude, latitude, and sometimes includes accuracy. Both GPS-captured and WiFi-enabled location data are known to have ~10 to 40 meters of error [33], [34]. For different activities, different approaches are used to improve the location accuracy.

For non-stationary activities, such as walking or in vehicle, a *Kalman filter* is used along with a map-matching technique to snap location points to the streets that the user was most likely on [35]. For stationary activities, the mostly likely location at which the user stayed can be inferred by taking the median values of the latitude and longitude among all the location samples [36].

Stage 4–Summarization: finally, statistics for a users' activities that are known to be relevant to the well-being or to a certain disease are computed. Say *et al.* [29] shows that for RA management, the majority of rheumatologists they interviewed identified 1) time spent walking, 2) time away from the house, and 3) gait speed as most useful in defining the level of RA disease activity. Rheumatologists preferred a simple, visual format that demonstrates trends over days, weeks, or months, which can help them make better management decisions and ultimately result in decreasing permanent joint damage.

Based on these insights, the above-mentioned statistics are computed in daily, weekly, and monthly bases and are presented as a calendar with a color coded index per day derived from a weighted sum of the above three factors using weightings suggested by rheumatologists to indicate the importance each factor (see Fig. 1).

B. Acoustics

Nonverbal characteristics from speech can strongly convey a speaker's psychological and emotional state [37] as well as other more nuanced social expressions (e.g. sarcasm [38]). There has been extensive work in the understanding of descriptors that can be derived from modulation, spectral characteristics and also energy based measures that can be indicative of stress [39],[40], emotion [41] or even depression [42]. In the context of pain such indicators are valuable. Clearly, stressful utterances, interjections and exclamations can be indicators of acute pain episodes.

However, a more interesting contribution from leveraging audio data is in monitoring at a *larger temporal scale* which has implications for long term CP management. Underlying negative psychological states such as stress, anxiety or depression [43] are well understood to have negative consequences in CP management [44]; such is the case that psychosocial treatments such as *cognitive behavioral therapy* (CBT) are a prevalent course of action [45]. Particular manifest behaviors such as rumination and catastrophization [46] are of interest in this context and are often vocally expressed. In the following example we demonstrate how passively acquired audio streams from smartphone microphones can deliver such high level inferences, in this example by the logging of stressful situations. We must note here that since acoustic processing methods within smart-phone deployed systems are particularly prone to privacy breaches, this leads to the preference to locally process signals that contain intelligible data. This in turn then necessitates methods with low demand on the battery which would require the use of the least computationally expensive methods. This method in the following is designed to take into account of this trade-off. However, there remains a research opportunity in the utilization of more sophisticated detection techniques [47], [48] into a mobile application.

Case Example: Detection of Stressful Situations from Voice for CP

Management: *Stage 1–Noise detection:* the first stage is the detection of *noise* in the sense of any significant level of sound. This is an essential first stage in the passive monitoring to retain only non-silent segments of the data. Simple thresholding on the root-mean-square values over a window of raw audio data are applied to differentiate the segments of *noise* versus silence[49].

Stage 2–Human speech detection: following this a further classification between as human voice or *noise* caused by another source is required. To this end, three simple features:(1) number of auto-correlation peaks, (2) non-initial autocorrelation peak, and (3) relative spectral entropy to robustly detect presence of human voice in audio streams [50] are used (Fig. 2). Moreover, the same three features can be retained to compute important prosody information regarding pitch and speaking rate [51]. This is important to note in terms of privacy. Previous research shows that the reconstruction of verbal information requires at least pitch and 2 harmonics, thus in using the above three features verbal content remains anonymous but prosody can be inferred [51]. In addition to being privacy sensitive, these features are computationally efficient and can work robustly to recognize voice in acoustically noisy environments [50]. This method has been used in several embedded and mobile phone systems [30], [52] and deployed in systems across several real world studies. These studies have demonstrated the inference of face-to-face conversation quality [53], social networks [51], [54] as well as social isolation and depressive symptoms [43] which all have implications for long term pain management.

Stage 3–Speaker identifier: once human voice is detected, it must be determined whether the voice came from the phone user. In Lu *et al.* [55], a Gaussian mixture model (GMM) based universal background model of all speakers is precomputed, the voice of the phone carrier is opportunistically sensed from phone conversations, and a model for the phone carrier's voice is learnt that is different from the universal background model. Subsequently, the learnt model is used to detect the phone carrier's voice. At this stage features such as mel-frequency cepstral coefficients where the verbal content can potentially be reconstructed are needed. Therefore all processing is done locally and all privacy-violating features are discarded save for the output.

Stage 4–Stress detection: if a detected human voice is that of the phone carrier the detection of stress is invoked. The following characteristics typify stress: the distribution of spectral energy switches to higher frequencies. Also, speaking rate and variability in pitch increases whereas pitch jitter decreases. A recent study showed a non-linear speech production model, based on a Teager Energy profile, can capture stressfulness in noisy outdoors situations. A GMM based classifier utilized these changes as features for both indoor and outdoor application [56]. Also, to account for idiosyncrasies during stressful conditions. A small amount of individual data can be used to calibrate personalized models. It was found that classifier accuracy improved by 8% if 2 minutes of labeled data is provided by the individual. Similar to stage 3, privacy violating features are discarded once the classification output has been obtained.

Stage 5-Logging: the number of stress occurrences can be logged giving a passive indicator of stress frequency with respect to time and location. Such continuous psychological monitoring gives a rich source of information for CBT suggested for CP management. Since the log is updated only when the prior 4 stages are invoked on demand, the computational overhead is minimized.

In other health related applications, the use of acoustic data in this way to infer higher level information about a user is not without precedent. Using similar privacy sensitive acoustic features there have been attempts to infer mental state or well-being with acoustics. These systems use acoustic data within multimodal frameworks using synchronous data from other smartphone sensors [43], [57]–[59]. Going forward, this suggests that the use of acoustic data for complex applications such as pain management would be one component within a multimodal frameworks.

C. Facial Expressions

An important and observable modality that can reflect internal state is facial expression. The prevalence of imagers from our laptops to our lampposts has greatly increased the availability of facial imagery as a source of data about an individual's state. Studies show that facial expressions during acute pain had a high level of consistency and repetition even when elicited by different stimulating modalities [60]. Moreover, a finding by Kappesser & de C. Williams [61] showed that a specific pain face can be distinguishable from faces expressing other negative emotions.

There have been numerous works in the automation of pain face recognition by way of analyzing video imagery from BSP/AC studies. An early example is [62], where face shape features were used with artificial neural networks to classify images of subjects' faces in a normal mood versus images taken from a pain inducing task. Lucey *et al.* [63] publically released the widely used UNBC-McMaster shoulder pain dataset which contains videos of patients with shoulder pain and a temporally concurrent pain score based on the Prkachin and Solomon pain intensity score [64]. Several subsequent studies in automated recognition followed [65]–[67], which utilize a range of image features from active appearance models based features, tracked anatomical points [68], discrete cosine transforms and local binary pattern [69] features with support vector machines (SVM) to classify pain faces. Sikka *et al.* [70] addressed the drawback labelling pain expressions at a sequence level leading to temporal uncertainties in terms of onset and end. The authors proposed a multi-segment multi-instance learning framework to determine expressive subsequences within super-sequences that contain pain expression. The same authors extend the machine vision capability beyond adults by demonstrating successful recognition of pain expression in postoperative children [71]. Finally, interesting findings by Bartlett *et al.* showed machine vision methods outperforming human observers in distinguishing real versus fake expression of pain [72], [73], this has important implications in the reliability of future ubiquitous pain monitoring systems.

Although much work has been done in the wider machine vision communities, to our knowledge there are no smartphone based systems designed to detect facial expressions of pain and only a few smartphone based systems for face expressions in general. This is

primarily due to the difficulties acquiring usable imagery from mobile devices from real world situations. Motion artifacts, out of plane head rotations, various lighting conditions and occlusions add to the difficulty with smartphones. In addition, power and computational constraints add to the challenge of locally processing if needed and the temporal concurrence of when images are captured with the onset of pain experiences cannot be guaranteed. Additionally, the effect of acting for the camera also generates a further confound, though as mentioned above acted and natural pain expressions can be automatically differentiable.

However, all of the aforementioned difficulties are only problematic if we consider single images or videos in isolation. If multiple instances taken over a long periods are analyzed collectively a broader perspective can be gained and could lend support to the inference of *spatiotemporally large scale behaviors*. For example, studies have shown that factors such as social isolation[74] are linked to physical pain and can be detrimental to CP management [75] and factors into CBT [45]. A number of relevant quantities can be extracted from images that are descriptive of social isolation. Emotional expressions from the user's face[12] can be detected and augmented to further contextual factors such as scene type detection [76], the number of people in the image and even the overall mood of the group of people in the image [77]. These measures can be further augmented to other relevant data such as usage of social networking apps, number of calls or instant messages made or received to construct an index for social isolation. Also, the use of further data modalities such as audio which could aid in reducing ambiguities in situational context and could also aid in identifying the true number of people present. In principle, a longitudinal analysis of trends in the index of social isolation can be used to feedback into CBT. The staged processing described in earlier sections could also be applied to expression extraction from images. The onset of more widespread implementation is bolstered by emerging systems such as in Suk *et al.* who proposed an initial framework for real time locally processed mobile use demonstrated the successful classification of seven affective states [78]. The recently released Affdex software development kit could facilitate the implementation of such analyses along with new datasets that contain real world ecologically valid datasets such as the AM-FED dataset[79].

D. Self-Reports as Signals

Self-report data are somewhat different from the signals previously discussed not only because intentional action is required from the individual to contribute data, but also because the data generated exists at a higher semantic level from inception. As such, rather than discussing stages of processing of self-report data, we share insight on modern self-report techniques and considerations.

The goal of much work in sensing is to reduce or remove the need for user input in data collection systems. While great strides have been made in sensing and detecting patient state, we are not yet in a position to collect all of the data relevant to the pain patient and clinician using passive techniques. Some aspects of behavior, such as detecting whether a patient struggles to button their shirt in the morning, is technically possible but not broadly and affordably deployable given current technology, though we can imagine it becoming so

in the future. Then there are cases where we are able to sense the behavior but not fully understand the context without input, such as a case where a patient walks more slowly to work. In this case, it is not known whether this is because of pain, another ailment, or simply commuting with a friend or colleague who walks more slowly.

The current standards for reporting the level of pain one is experiencing are pain scales and maps. Pain thermometers and pain faces (e.g. Iowa faces of pain) remain the most widely used and do translate well to digital experiences [80]. There have been several adaptations and improvements on user input for these assessments [81]. Pain maps are basic digital maps of the human body that allow patients to document/describe the location and nature of pain. It should also be pointed out that the assessment of pain is quite complex (due to variability in individual pain thresholds, diagnosis, chronic versus acute, and so on), and the community is not fully in agreement on how best to do so [82].

Like other data streams discussed in this paper, we consider self-report to be largely if not exclusively temporal. Data are collected at a specific time and those data represent a certain time span. Unlike most sensed data, however, the time data is collected and the time span the data ‘cover’ are often substantially different in two important dimensions:

Timeliness of data.—Self-report data are out of necessity a form of recall, with the patient recounting experiences that have occurred in the past. While a goal of self-report is often to reduce the amount of time between data collection and the recalled experience, practically there is a great deal of variability in these time lags. For example, data may be requested from a patient about an event as it is detected, but the patient may not respond for several hours or even days.

Time span of the data.—A single self-report data point maybe reflect a narrow point in time, such as an emotional response to a stimulus or level of pain experienced. Such measurements are often referred to as measures of *state*. On the other hand, the patient may be asked to report data in such a manner that a longer time span is recalled, such as emotional state over the past year, or even irrespective of time span, such as a patient’s personality type. Such measurements are referred to as *trait*. [83].

Self-report originated as one of the main forms of data transfer between patient and clinician, but the advent of modern medicine relegated it to completing forms and sharing symptoms as precursors to testing. More recently, mobile devices have led to a resurgence of self-reported data through a class of techniques labeled ecological momentary assessment (EMA) in which patients are polled for information, typically using abbreviated versions of gold standard clinical and behavioral assessment forms, frequently and *in situ* [84]. The central argument for the use of EMA is that while the data may yet be subject to interpretation and distortion, the effects of recall bias and error can be mitigated by timeliness of surveillance [85]. To further address recall bias a method was proposed by Rahman *et al.* [86] where EMA queries are presented with the addition of contextual information about the time, location activity and acoustic state of the user during the time in question: *contextual recall*. This allows for the EMA activation time to be more flexible, the

results the study showed improved recollection of stress levels compared to when no context was given.

Self-report and EMA have continued to evolve beyond simply asking patients standardized questions. The large touch screens on current generation mobile devices have afforded more interactive, visual methods for assessing patients. These methods are often employed to mirror and simplify ‘pen and paper’ assessments rather than simply translate them to digital. The photographic affect meter (PAM), for example, asks patients to reflect on their emotional state by choosing from a series of emotionally representative images (see Fig. 3 left). Responses take a few seconds to complete and are calibrated to and correlate with the 20-item PANAS, widely considered the gold standard for measurement of emotional state [87].

Another approach to self-report has been to find more straightforward and answerable representations of complex or difficult behaviors or experiences. Pain level, for example, is notoriously difficult to measure reliably, so many clinicians focus on pain interference or activities of daily living (ADLs) to assess what patients are or are not able to do as a result of the pain they experience. This can serve as both an adequate proxy for pain levels and a more direct approach to treating the patient to return to normal and desired activities. In a variant of ADLs, a system called YADL (your activities of daily living) [88] was designed to allow patients to rapidly and intuitively sort through photos of common activities to document which are easy or difficult for them as a result of their pain. A daily version of YADL presents patients with only those images they have selected as difficult, prompting them to choose which have been hard for them on that particular day, or which they avoided because of anticipated pain (see Fig. 3 right). Such an approach allows for an improvement on the traditional means of assessing ADLs along with more fine-grained, frequent data on patient experience. YADL is currently undergoing validation, as with the mobility index mentioned above, the modular design of YADL will allow it to be tuned and adapted to particular populations and disease states.

An alternative modality for self-reporting which does not rely on visual prompts or stimuli are tangible interfaces. Such methods may for some be a more intuitive way to convey an internalized level of pain. By leveraging the gripping action as a natural gestural response when pain is experienced, a small lightweight squeezable stick *Keppi* [89] designed to be continuously carried by the user has been developed (see Fig. 4). The device consists of a conductive foam based, force-sensitive resistor covered in soft rubber with embedded signal conditioning, an ARM Cortex-M0 microprocessor, and Bluetooth low energy. Although still in a developmental phase an in-lab feasibility study showed participants were able to consistently map pressure to four stratified levels of potential pain as well as match a dynamic visual cue when given visual feedback of the pressure value. Further real world evaluation and miniaturization of this device with pain participants is in process.

E. Modular System Architecture

In the previous sections the contribution of individual signals and modalities were discussed. However, *mHealth* systems will be most effective in specific disease contexts and across contexts if they incorporate multiple active and passive data streams. Adopting a modular

system architecture is critical in order to address the system complexity associated with heterogeneous and variable data-stream inputs. A modular architecture will allow these systems to be adapted as new apps and device capabilities appear, to be tailored to particular needs and populations, and to serve as a building block for other health related applications. Moreover, as particular signal processing methods improve, they can more readily be evaluated and integrated into modular systems. We provide two examples of modular system architectures in this section, one optimized for mobile device processing, and the second for server side processing. It is worth noting here that though we focus here on modular systems, another interesting yet under-investigated line of research in *mHealth* is the use of advances in data compression. Since power consumption is a limiting factor where local processing is often required, efficiency is desirable both in terms of transmission energy conservation [90] and the codification of the various aforementioned signals to reduce data size.

1) SAINT: Scalable Sensing and Inference Toolkit: Each passive sensor data type generates a large volume of raw data whose transmission can tax bandwidth and battery resources. Moreover, these data may require local processing to address privacy concerns, e.g., detecting the presence of a human voice [91]. However, it is challenging to run the sensing and signal processing in real-time on resource constrained mobile phones. A poorly written signal processing algorithm may consume large amounts of CPU cycles and be a significant drain on the battery. Similarly if intermediate storages during raw sensor recording or signal processing are not cleaned then the system can be out of memory easily for high data rate sensor streams (e.g., audio or video).

To this end, the open-source sensing and inference toolkit SAINT was developed [92] that can assist sensor data collection and signal processing inside the phone. In SAINT APIs to collect 30 different raw sensor traces and processed data-streams are provided. SAINT tackles these challenges with time- and memory-efficient implementations of its sensing and processing libraries. Fast signal processing algorithms were implemented in native layers (i.e., not on a virtual machine) with low time-complexity [93]. The implementations are also tested over multiple deployments and were found to be stable and battery efficient. In addition to fast and memory efficient implementations, SAINT also provides admission control and duty-cycling for demanding signal processing tasks to further reduce battery usage. Admission control allows for on-demand triggering of signal processing, e.g. speaker or stress recognitions are only triggered when there is human voice in audio data. Duty-cycling on the other hand saves battery by periodically sampling data streams to approximate a human behavior signal. For example, running activity recognition at 10 second intervals can approximate the subject's total level of physical activity fairly accurately.

The modular design enables reuse and extension of SAINT's built-in sensing and inference capabilities. At its core, SAINT uses a publish-subscribe [94] or *pub-sub* pattern. In a *pub-sub* pattern, publishers produce data for which subscribers listen. An intermediate *event bus* or *broker* manages subscription and data transfer among publisher and subscribers (Fig. 5). Subscribers register for specific publisher data to the *broker*; publishers then send data directly to the *broker*, and the *broker* relays the information to registered subscribers. Inside

SAINT, any existing raw sensor or derived stream is a publisher. A new derived stream can be created in three steps: (1) registering as subscribers to necessary publishers, (2) performing necessary computation on the publisher data to create the new derived stream, and (3) making the derived stream available as a publisher for later reuse. For instance, a new sleep detector could use the presence of movement, location, human voice, and phone charging status from SAINT to detect if a person is sleeping, then make the output of this detector available for other components to use. Furthermore, registering to existing publishers and making the new derived stream data available can be easily implemented.

Another benefit of the pub-sub structure is the centralization of sensing and processing that reduces redundancy in computation. In SAINT, if multiple subscribers want the same publisher's data, then the publisher just produces the data once and the broker relays the information to each subscriber, reducing redundancy. For example, if a sleep and stress detector both run on the phone and want to know if human voice is detected, SAINT detects human voice from audio data once and relays the information to the sleep and stress detectors.

2) Lifestreams: The ultimate goal of many novel data collection techniques described previously is to enable robust and actionable behavioral indicators that can be used to characterize a patient's baseline, and then to identify significant variations, trends and shifts in specific behaviors or symptoms that are relevant to an individual's health. However, making sense and acting on these multi-dimensional, heterogeneous data streams requires iterative and intensive exploration of the data sets, and development of customized analysis techniques that are appropriate to the health domain of interest. Lifestreams is a modular and extensible open-source data analysis stack that runs on the server and is designed to accelerate the development and refinement of sense making techniques [34].

Lifestreams runs on top of mobile data collection systems, such as Ohmage or SAINT [92], [30]. It processes the raw or intermediate data provided by these tools with a multi-layer analysis stack, depicted in Fig. 6, consisting of the following: 1) feature extraction, 2) feature selection, 3) inference, and 4) visualization. Each layer consists of modules that process the data provided by the lower layer, and then send the result to the next layer up. From the input data streams, the feature extraction layer computes various statistics, such as daily min/max, average, or the weights of the principal components derived from a set of measurements. The feature compression and selection layer leverages techniques, such as cross-entropy or correlation analysis, and the domain experts' input to select or compose the features that are most relevant to the analysis goal. Then, the inference layer uses the selected features to detect patterns and trends with correlation estimation, and change detection algorithms [95]. Finally, the visualization layer uses different visualization methods to make the inferred patterns or trends actionable for the end user. Lifestreams was found to be useful as a tool for the research coordinators to quickly navigate through the data and provide visual aids to guide the discussion with participants during interview sessions. Its extensible design allows it to be easily extended to support different studies exploring various new data streams and transfer to other domains

III. Open Challenges in Inferencing and Feedback

So far we have highlighted, through examples, state of the art methods and tools to achieve the *measurement-inference-management* loop. However, in this section we discuss further factors that are important but either underutilized or remain unresolved. These challenges are most conspicuous when considering *inference* and *management*. In this section we discuss these factors, as well as advances from related studies that have a bearing on future directions.

A. Accounting for Temporal Patterns

The presence of a circadian rhythm in the experience and perception of pain is well understood. This relates to a variety of chronobiological factors such as daily variations in endorphin and enkephalin concentrations in the pain processing parts of the brain. The review by Junker & Wirz summarizes [96] several studies that include both naturally occurring as well as experimenter induced pain. The studies within showed specific and consistent times of the day when pain perception tended to peak depending on the cause. Investigations in using smart-phone technology as a way to monitor daily rhythms have been recently investigated for sleep applications. Most studies have focused on sleep pattern inference by utilizing app usage, charging, screen unlocking, along with an ambient light detector and audio streams to detect sleep durations and quality [97]–[99]. Recently, Abdullah *et al.* [100] demonstrated that by using periods of smartphone non-use, sleep patterns can be inferred using a rule-based algorithm which accurately matched with ground truth from a sleep journal. With the combination of such sleep measures, along with the domain knowledge of the effects of circadian rhythm on specific types of pain, there is great potential for more closely personalized and more timely issuances of interventions for pain management in a daily cycle.

As we have discussed pain conditions with high levels of chronicity require long term management and therapy. Also over enough time, persistent negative psychological states can emerge. It has been shown that, given enough data from long periods, such mental states can be detected. With data spanning 2 weeks Saeb *et al.* [101] demonstrated the use of movement, activity and phone use derived descriptors to infer depression severity levels by way of *supervised learning*, resulting in an accuracy 86.5%. Over a 10 week period, Wang *et al.* [57] investigated correlations between mental well-being measures: depression, stress, flourishing and loneliness scores with tracked information relating to conversation, activity and sleep. Such results demonstrate progress toward inferences on low frequency changes in mental state or even *traits* which are ultimately necessary for CBT related management. Longer period factors such as climate and seasonality have yet to be leveraged in this regard but are readily measurable if studies of this time length are done.

B. Advances in Pattern Recognition

In mobile sensing, a common paradigm to generate predictive models is *supervised learning* [27]. Although a widely used and powerful approach, there remains two critical dependencies. The first is the generation of good features from raw sensor data and the second is the truthfulness of labelled states or behavioral categorizations. Given that there is

no standard method to optimally engineer features, a handcrafting approach is often done using expert domain knowledge, intuition, or other prior experimentation. For lower level inferences good descriptors can be readily determined. However, it is often the case that abstract higher level inferences (such as ‘anxiety’) are desirable, but such tasks would require complex feature engineering and are difficult to map, often leading to poorly performing models. Moreover, such abstract labels have loose definitions or are ambiguous to people and thus difficult to assign correctly when creating a training set.

In light of this, we point to the increasingly popular paradigm of *deep learning* [102], [103], which has the capacity to learn feature representations and map low level data to highly abstract categories in a less handcrafted way. However, to date, this has mostly been demonstrated with large unimodal datasets in other application domains. That said, Martinez & Yannakakis [104] proposed a convolutional neural network framework designed to fuse combinations of continuous and discrete 1 dimensional physiological signals with differing sample rates to successfully infer 6 abstract affective states (anxiety, frustration, fun, relaxation, challenge and excitement). Such a method naturally lends itself to the high level behavioral inferences needed for pain management and from the assortment of 1 dimensional signals acquired from mobile sensing. Notably, in this case from the use of wearable devices that can extract such physiological signals. The drawback is the high computational expense in the training and use of highly parametric models, which may not be able to deliver timely output for intervention.

A further significant challenge when dealing with human mobile sensing is the high degree of idiosyncrasy within the data streams. In recent years generalizable methods to capture latent structures in behavioral routine from mobile data have been proposed; such structures could be used as personal baselines from which anomalous behavior could be inferred. Early work by Eagle & Pentland [105] proposed a method to establish eigen-representations of behavior from mobile data to model structures in day to day routine. This method was effective in modelling the idiosyncratic patterns in a personal routine as well as generating a way to find similarities between groups of people. However, this approach relies on rich, densely-sampled high quality data which cannot be realistically acquired in real world sensing. To this end, Zheng *et al.* [106] proposed a collaborative filtering model to overcome sparsity in a single user’s data based on the intuition that many users follow similar behavior patterns; of course this is dependent on the availability of multiple users’ data.

Another paradigm which can address the problem of gaps in data and/or idiosyncratic variability is *multi task learning* (MTL) [107]. Simultaneous *supervised learning* is applied to a set of different but related classification or regression tasks. For example, data from a user from two different weeks can be used to infer a behavior for each week where some data maybe missing one of the weeks. Learned parameters common to both tasks (weeks) can be leveraged to bolster the accuracy of the partial week when compared to training two models separately. Similarly this can be done in terms of different people as tasks thus accounting for commonalties or differences between people. Romera-Paredes *et al.* [108] recently showed this capability with facial expressions of pain as well as with electromyographic signals by exploiting the *transfer learning* property in MTL. In a similar study, these authors also show an improvement in the recognition of pain face expression

using OrthoMTL [109] which assumes features that describe identity are unrelated (orthogonal) to features that describe pain expression. In essence this method calibrates to each person by simultaneously learning on identity features and pain expressive features and separates them. In principle this can be applied to data from any modality or combination of modalities as long as this is consistent between tasks. Moreover, advances in MTL show the learning process is equivalent to a single convex optimization process [110], as such there is potentially fewer process iterations when compared to repeatedly training multiple models.

Finally, it is worth noting here that a major barrier to the advancement in predictive modelling is the availability of data or lack thereof. However, in terms of *spatiotemporally small scale behaviors* in pain expression a new extensive multi-modal dataset called *EmoPain* was recently released [14]. It explicitly addresses this requirement gap for chronic back pain. Contained within are: high resolution face imagery, surface electromyography, acoustics and whole body motion capture. Two sets of pain labels are included, based on facial expression and known body motion based behaviors [7], the authors set baseline modelling frameworks for each label set using detected facial points with SVM classifiers for pain face recognition and speed and posture based features with Random Forest models for the body behaviors. Public datasets for *large scale* behaviors are still very rare due to the challenges of de-identification, especially for a dataset with patient information.

C. Feedback to Users

Interfacing with the user once inputs, measurements, and inferences have been made is dependent on the type and nature of the required management. Currently, there are numerous commercial pain management mobile apps which are principally designed for specific functions, including self-monitoring, pain intensity tracking, information provision, medication management, and relaxation training. However, there remains a lack of traction with users. Recent reviews [111], [112] and commentaries [113], [114] unanimously suggest a lack of clinical engagement and regulation in the development of such software, and efficacious claims are underpinned by little evidence. Moreover, the apps relating to CP self-management are still relatively simplistic in how they account for cognitive behavioral factors which are essential for persuasive and effective interfacing.

Some commercial products such as Habit Changer is described as utilizing cognitive behavioral components within the management strategy but details are not provided. Some newer products such as the WebMD PainCoach enables the user to monitor pain, set and track activity goals and generates related messages. These new systems are promising in terms of personal monitoring, but future versions call for further functionality in using this information to effectively respond to episodes of de-motivation which are ultimately at the root of poor progress [115].

In terms of general applications, mobile systems can be categorized in three groups in terms of feedback strategy: (1) aggregation of data into summary statistics, often augmented with attractive visualizations. For instance, Ubifit [116] or BeWell [117] which uses background wallpaper to show overall physical activity, social interaction and sleep. This is often purposed for goal setting or gentle priming towards the goal [118], (2) visualization of data which relies on the users to self-explore and reflect on the data [119], [120] and (3) the

delivery of generic recommendations that are globally applicable or are tailored to specific subgroups based on demographics, culture or lifestyle [121], [122]. However, these strategies do not make use of any in-depth analysis and may not only miss potential opportunities to effectively intervene, but may also fail to deliver effective feedback that is actionable and relevant to the individual.

That said, increased persuasiveness in a real world deployment has been demonstrated. Rabbi *et al.* [123] proposed a context driven food and exercise suggestion engine *MyBehavior*. In this system text based suggestions which are perceived to be low effort and familiar to each user are generated based on both passive and actively sensed data. Such perceptions of familiarity and low effort increases the likelihood of actualization according to behavioral theory [124]. Results over 10-week user trials show significant increases in adoption when compared to randomly issued suggestions. Similar persuasive feedback would be of great benefit for musculoskeletal pain for example where avoidance and demotivation to do regular exercise is a prevalent problem.

An interesting direction in *spatiotemporally small scale* feedback is the sonification of patient's movements during physiotherapy [115]. Common motion related traits are self-guarding movement and a reduction in proprioception. Sound feedback exploits the tight links between the auditory and motor parts the brain. Singh *et al.* [125] proposes a design framework: *Go-with-the-flow* for sonified wearable systems for CP exercising. Initial results indicate an impact on body awareness, increased motivation to reach exercise targets, and positive gains in compensatory motion and self-efficacy.

Ultimately, all patient sensing has the purpose of informing decisions and actions. The focus of this review is on the capture and processing of relevant data sources to inform user feedback but in the end effectively closing critical feedback loops for health will depend on effective ways of presenting to and interacting with the end user/patient.

IV. Conclusion

Patient centric data streams will increasingly inform disease management and diagnosis. They will be used to contextualize and personalize patient response to treatment and to build applications that support patients in their own self-management. In this review we focused on the use case of pain measurement and management because it has broad applicability, significant implications for both clinical outcomes and patient quality of life, and because mobile data streams can fill a significant gap in current approaches to measurement. While today's technology is ready to be put to use to improve pain measurement and management, there remain challenges and opportunities for further research.

The greatest short term challenges are to develop the evidence base for these approaches and to address usability and relevance for both patient and clinician. In addition, a modular data architecture is critical to promote the inclusion of new data sources (applications and devices) and the adaptation to new and specialized users (conditions and demographics) over time.

Biographies

Min S. Hane Aung received the Ph.D. degree in applied machine learning from the University of Liverpool, Merseyside, U.K. He is currently a Visiting Lecturer at the UCL Interaction Centre, University College London and also a Research Associate within the People-Aware Computing group at Cornell University. He is a Member of the Institution of Engineering and Technology. His main research interests include the development of recognition systems for human behavior.

Longqi Yang (M'15) was born in China in 1992. He received the B.Eng. degree from Zhejiang University, Zhejiang, China, in 2014. He is currently working toward the Ph.D. degree and AOL Ph.D. Fellow in the Computer Science Department, Cornell University, Ithaca, NY, USA, and Cornell Tech. He is currently advised by Prof. D. Estrin. He is a member of the small data lab and the connected experiences lab at Cornell Tech. His main area of research interests are personalization and mobile health.

Deborah Estrin (F'04) received the B.S. degree from University of California Berkeley, Berkeley, USA, in 1980, and the M.S. and Ph.D. degrees from Massachusetts Institute of Technology, Cambridge, USA, in 1982 and 1985, respectively. She was a Faculty Member of Computer Science at Computer Science Department, University of Southern California, from 1986 to 2000 and University of California Los Angeles. She is currently a Professor of Computer Science at Cornell Tech, New York City and founded the Health Tech Hub of the Jacobs Institute. Her primary areas of research are mobile health and the analysis and use of small data. She is an elected member of the American Academy of Arts and Sciences and the National Academy of Engineering.

Tanzeem Choudhury received the Ph.D. degree from MIT Media Lab, Cambridge, MA, USA. She is currently an Associate Professor in the Information Science Department and directs the People-Aware Computing Group, Cornell University, Ithaca, NY, USA. Her primary research interests include mobile sensing of health and ubiquitous computing. She is a Member of the ACM.

REFERENCES

- [1]. Estrin D and Sim I, "Open mHealth architecture: An engine for health care innovation," *Science*, vol. 330, no. 6005, pp. 759–760, 2010. [PubMed: 21051617]
- [2]. Nilsen W, Kumar S, Shar A, Varoquiers C, Wiley T, Riley WT, and Atienza AA, "Advancing the science of mHealth," *J. Health Commun*, vol. 17, no. sup1, pp. 5–10, 2012.
- [3]. Eccleston C, "Role of psychology in pain management," *Brit. J. Anaesth*, vol. 87, no. 1, pp. 144–152, 2001. [PubMed: 11460803]
- [4]. Merskey H and Bogduk N, *Classification of Chronic Pain*. Seattle, WA, USA: International Association for the Study of Pain Press, vol. 210, 1994.
- [5]. Merskey H, "The taxonomy of pain," *Med. Clin. North Amer*, vol. 91, no. 1, pp. 13–20, 2007. [PubMed: 17164101]
- [6]. Loeser JD and Melzack R, "Pain: An overview," *The Lancet*, vol. 353, no. 9164, pp. 1607–1609, 1999.
- [7]. Keefe FJ and Block AR, "Development of an observation method for assessing pain behaviour in chronic low back pain patients," *Behav. Ther*, vol. 13, no. 4, pp. 363–375, 1982.

- [8]. Fordyce W, Behavioural Method for Chronic Pain and Illness. St. Louis, MO, USA: Mosby, 1976.
- [9]. Sullivan MJL, Thibault P, Savard A, Catchlove R, Kozey J, and Stanish WD, "The influence of communication goals and physical demands on different dimensions of pain behavior," *Pain*, vol. 125, pp. 270–277, 2006. [PubMed: 16860479]
- [10]. Narayanan S and Georgiou PG, "Behavioral signal processing: Deriving human behavioral informatics from speech and language," *Proc. IEEE*, vol. 101, no. 5, pp. 1203–1233, 5 2013.
- [11]. Pantic M, Pentland A, Nijholt A, and Huang TS, "Human computing and machine understanding of human behavior: A survey," in *Artificial Intelligence for Human Computing*. Berlin, Germany: Springer, pp. 47–71, 2007.
- [12]. Zeng Z, Pantic M, Roisman GI, and Huang TS, "A survey of affect recognition methods: Audio, visual, and spontaneous expressions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 1, pp. 39–58, Jan. 2009. [PubMed: 19029545]
- [13]. Vyzas E and Picard RW, "Affective pattern classification," in *Proc. Assoc. Adv. Artif. Intell. Fall Symp. Ser.: Emotional Intelligent: Tangled Knot Cognit*, pp. 176–182, 1998.
- [14]. Aung MSH, Kaltwang S, Romera-Paredes B, Martinez B, Singh A, Cella M, Valstar M, Meng H, Kemp A, Shafizadeh M, Elkins AC, Kanakam N, de Rothschild A, Tyler N, Watson PJ, Williams A. C. de C., Panitc M, Bianchi-Berthouze N et al., "The automatic detection of chronic pain-related expression: requirements, challenges and the multimodal *EmoPain* dataset," in *IEEE Transactions on Affective Computing*, to appear.
- [15]. Vlaeyen JWS and Linton SJ, "Fear-avoidance and its consequences in chronic musculoskeletal pain: A state of the art," *Pain*, vol. 85, no. 3, pp. 317–332, 2000. [PubMed: 10781906]
- [16]. Felipe S, Singh A, Bradley C, Williams A, and Bianchi-Berthouze N, "Roles for personal informatics in chronic pain," in *Proc. Int. Conf. Pervasive Comput. Technol. Healthcare*, vol. 15, 2015, pp. 161–168.
- [17]. Jensen MP, Turner JA, and Romano JM, "Changes after multidisciplinary pain treatment in patient pain beliefs and coping are associated with concurrent changes in patient functioning," *Pain*, vol. 131, no. 1/2, pp. 38–47, Sep. 2007. [PubMed: 17250963]
- [18]. Rainville P, Bao QVH, and Chrétien P, "Pain-related emotions modulate experimental pain perception and autonomic responses," *Pain*, vol. 118, no. 3, pp. 306–318, 2005. [PubMed: 16289802]
- [19]. King WC and S Bond D. "The importance of pre and postoperative physical activity counseling in bariatric surgery," *Exercise Sport Sci. Rev* vol. 41, no. 1, pp. 26–35, 2013.
- [20]. Bravata DM, Smith-Spangler C, Sandaram V, Gienger AL, Lin N, Lewis R, Stave CD, Olkin I, and Sirard JR, "Using pedometers to increase physical activity and improve health: A systematic review," *JAMA*, vol. 298, no. 19, pp. 2296–2304, 2007. [PubMed: 18029834]
- [21]. Greaves CJ, Sheppard KE, Abraham C, Hardeman W, Roden M, Evans PH, and Schwarz P, "Systematic review of reviews of intervention components associated with increased effectiveness in dietary and physical activity interventions," *BMC Public Health*, vol. 11, no. 1, 2011.
- [22]. Patel MS, Asch DA, and Volpp KG. "Wearable devices as facilitators, not drivers, of health behavior change," *JAMA*, vol. 313, no. 5, pp. 459–460, 2015. [PubMed: 25569175]
- [23]. Richardson CR, Newton TL, Abraham JJ, Sen A, Jimbo M, M Swartz A, "A meta-analysis of pedometer-based walking interventions and weight loss," *Ann. Family Med* vol. 6, no. 1, pp. 69–77, 2008.
- [24]. Dodge HH, Mattek NC, Austin D, Hayes TL, and Kaye JA, "In-home walking speeds and variability trajectories associated with mild cognitive impairment," *Neurology*, vol. 78, no. 24, pp. 1946–1952, 2012. [PubMed: 22689734]
- [25]. Bort-Roig J, Gilson ND, Puig-Ribera A, Contreras RS, and Trost SG, "Measuring and influencing physical activity with smartphone technology: A systematic review," *Sports Med*, vol. 44, no. 5, pp. 671–686, 2014. [PubMed: 24497157]
- [26]. Case MA, A Burwick H, Volpp KG, and Patel MS, "Accuracy of smartphone applications and wearable devices for tracking physical activity data," *JAMA*, vol. 313, no. 6, pp. 625–626, 2015. [PubMed: 25668268]

- [27]. Lane ND, Miluzzo E, Lu H, Peebles D, Choudhury T, and Campbell AT, "A survey of mobile phone sensing," *IEEE Commun. Mag.*, vol. 48, no. 9, pp. 140–150, Sep. 2010.
- [28]. Sohn T, Varshavsky A, LaMarca A, Chen MY, Choudhury T, Smith I, and De Lara E, "Mobility detection using everyday GSM traces," in *UbiComp 2006: Ubiquitous Computing*. Berlin, Germany: Springer, pp. 212–224, 2006.
- [29]. Say P, Stein DM, Ancker JS, Hsieh A, Pollak JP, and Estrin D, "Smartphone data in rheumatoid arthritis—What do rheumatologists want?" in *Proc. AMIA Annu. Symp.*, Nov. 2015, pp. 1130–1139.
- [30]. Tangmunarunkit H, Hsieh CK, Longstaff B, Nolen S, Jenkins J, Ketcham C, Selsky J, Alquaddoomi F, George D, Kang J, Khalapyan Z, Ooms J, Ramanathan N, and Estrin D, "Ohmage: A general and extensible end-to-end participatory sensing platform," *ACM Trans. Intell. Syst. Tech.*, vol. 6, no. 3, article 38, 2015.
- [31]. Longstaff B, Reddy S, and Estrin D, "Improving activity classification for health applications on mobile devices using active and semi-supervised learning," in *Proc. Pervasive Comput. Technol. Healthcare*, 2010, pp. 1–7.
- [32]. Welch LR, "Hidden Markov models and the Baum-Welch algorithm," *IEEE Inf. Theory Soc. Newslett.*, vol. 53, no. 4, pp. 10–13, Dec. 2003.
- [33]. Liao L, Fox D, and Kautz H, "Extracting places and activities from gps traces using hierarchical conditional random fields," *Int. J. Robot. Res.*, vol. 26, no. 1, pp. 119–134, 2007.
- [34]. Cheng YC, Chawathe Y, LaMarca A, and Krumm J, "Accuracy characterization for metropolitan-scale Wi-Fi localization," in *Proc. 3rd Int. Conf. Mobile Syst., Appl. Services*, 2005, pp. 233–245.
- [35]. Lou Y, Zhang C, Zheng Y, Xie X, Wang W, and Huang Y, "Map-matching for low-sampling-rate GPS trajectories," in *Proc. 17th ACM SIGSPATIAL Int. Conf. Adv. Geographic Inform. Syst.*, 2009, pp. 352–361.
- [36]. Hsieh C, Tangmunarunkit H, Alquaddoomi F, Jenkins J, Kang J, Ketcham C, Longstaff B, Selsky J, Dawson B, Swendeman D, Estrin D, and Ramanathan N, "Lifestreams: A modular sense-making toolset for identifying important patterns from everyday life," in *Proc. 11th ACM Conf. Embedded Netw. Sensor Syst*, 2013, vol. 11, no. 5.
- [37]. Weninger F, Wollmer M, and B. Schuller, "Emotion recognition in naturalistic speech and language—A survey," *Emotion Recog., A Pattern Anal. Approach*, pp. 237–267, 2014.
- [38]. Rockwell P, "Lower, slower, louder: Vocal cues of sarcasm," *J. Psycholinguistic Res.*, vol. 29, no. 5, pp. 483–495, 2000.
- [39]. Hansen JH, Kim W, Raurkar M, Ruzanski E, and Meyer-hoff J, "Robust emotional stressed speech detection using weighted frequency subbands," *EURASIP J. Adv. Signal Process.*, vol. 1, pp. 1–10, 2011.
- [40]. Lai M, Chen Y, Chu M, Zhao Y, and Hu F, "A hierarchical approach to automatic stress detection in English sentences," in *Proc. Acoust. Speech, Signal Process*, 2006, vol. 1, pp. I-753–I-756.
- [41]. Grimm M, Kroschel K, Mower E, and Narayanan S, "Primitives-based evaluation and estimation of emotions in speech," *Speech Commun.*, vol. 49, no. 10, pp. 787–800, 2007.
- [42]. Yang Y, Fairbairn C, and Cohn JF, "Detecting depression severity from vocal prosody," *IEEE Trans. Affective Comput.*, vol. 4, no. 2, pp. 142–150, Apr-Jun 2013.
- [43]. Rabbi M, Ali S, Choudhury T, and Berke E, "Passive and in-situ assessment of mental and physical well-being using mobile sensors," in *Proc. 13th Int. Conf. Ubiquitous Comput.*, 2011, pp. 385–394.
- [44]. Arnow BA, Hunkeler EM, Blasey CM, Lee J, Constantino MJ, Fireman B, Kraemer HC, Dea R, Robinson R, and Hayward C, "Comorbid depression, chronic pain, and disability in primary care," *Psychosomatic Med.*, vol. 68, no. 2, pp. 262–268, 2006.
- [45]. Ehde DM, Dillworth TM, and Turner JA, "Cognitive-behavioral therapy for individuals with chronic pain: Efficacy, innovations, and directions for research," *Amer. Psychologist*, vol. 69, no. 2, pp. 153–166, 2014.
- [46]. Turner JA, Holtzman S, and Mancl L, "Mediators, moderators, and predictors of therapeutic change in cognitive-behavioral therapy for chronic pain," *Pain*, vol. 127, no. 3, pp. 276–286, 2007. [PubMed: 17071000]

- [47]. Saon G, Thomas S, Soltau H, Ganapathy S, and Kingsbury B, "The IBM speech activity detection system for the DARPA RATS program," in Proc. INTERSPEECH, 2013, pp. 3497–3501.
- [48]. Thomas S, Saon G, Van Segbroeck M, and Narayanan SS, "Improvements to the IBM speech activity detection system for the DARPA RATS program," in Proc. IEEE Int. Conf. Acoust., Speech Signal Process., 2015, pp. 4500–4504.
- [49]. Lu H, Hong, Yang J, Liu Z, Lane ND, Choudhury T, and Campbell AT, "The Jigsaw continuous sensing engine for mobile phone applications," in Proc. 8th ACM Conf. Embedded Netw. Sensor Syst., 2010, pp. 71–84.
- [50]. Basu S, "Conversational scene analysis," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, MA, USA, 2002.
- [51]. Wyatt MD, "Measuring and modeling networks of human social behavior," Ph.D. dissertation, Univ. Washington, Washington, DC, USA, 2010.
- [52]. Choudhury T et al., "The mobile sensing platform: An embedded activity recognition system," IEEE Pervasive Comput, vol. 7, no. 2, pp. 32–41, Apr-Jun 2008.
- [53]. Choudhury TK, "Sensing and modeling human networks," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, MA, USA, 2003.
- [54]. Wyatt D, Choudhury T, Bilmes J, and Kitts JA, "Inferring colocation and conversation networks from privacy-sensitive audio with implications for computational social science," ACM Trans. Intell. Syst. Technol, vol. 2, no. 1, 2011, Art. no. 7.
- [55]. Lu H, Brush AJ, Priyantha B, Karlson AK, and Liu J, "Speak-erSense: Energy efficient unobtrusive speaker identification on mobile phones," in Pervasive Computing. Berlin, Germany: Springer, 2011, pp. 188–205.
- [56]. Lu H, Frauendorfer D, Rabbi M, Schmid Mast M, Chittaranjan GT, Campbell AT, Gatica-Perez D, and Choudhury T, "StressSense: Detecting stress in unconstrained acoustic environments using smart-phones," in Proc. ACM Conf. Ubiquitous Comput., 2012, pp. 351–360.
- [57]. Wang R, Chen F, Chen Z, Li T, Harari G, Tignor S, Zhou X, Ben-Zeev D, and Campbell AT. "StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones," in Proc. ACM Conf. Ubiquitous Comput., 2014, pp. 3–14.
- [58]. Ben-Zeev D, Wang R, Abdullah S, Brian R, Scherer EA, Mistler LA, Hauser M, Kane JM, Campbell A, and Choudhury T, "Mobile behavioral sensing for outpatients and inpatients with schizophrenia," Psychiatric Services, vol. 67, no. 5, pp. 558–561, 2016. [PubMed: 26695497]
- [59]. Choudhury T, "Tracking behavioral symptoms of bipolar disorder using automated sensing and delivering personalized interventions using smartphones," in Proc. 17th Annu. Conf. Int. Soc. Bipolar Disorders, 2015, vol. 17, p. 10.
- [60]. Prkachin KM, "The consistency of facial expressions of pain—A comparison across modalities," Pain, vol. 51, no. 3, pp. 297–306, 1992. [PubMed: 1491857]
- [61]. Kappesser J and Williams A. C. de C., "Pain and negative emotions in the face: Judgements by health care professionals," Pain, vol. 99, nos. 1/2, pp. 197–206, 2002. [PubMed: 12237197]
- [62]. Monwar M and Rezaei S, "Pain recognition using artificial neural network," in Proc. IEEE Int. Symp. Signal Process. Inform. Technol, 2006, pp. 28–33.
- [63]. Lucey P, Cohn JF, Prkachin KM, Solomon PE, and Matthews I, "Painful data: The UNBC-McMaster shoulder pain expression archive database," in Automatic Face & Gesture Recognition and Workshops (FG 2011), pp. 57–64, 2011.
- [64]. Prkachin K and Solomon P, "The structure, reliability and validity of pain expression: The evidence from patients with shoulder pain," Pain, vol. 139, pp. 267–274, 2008. [PubMed: 18502049]
- [65]. Ashraf AB, Lucey S, Cohn JF, Chen T, Ambadar Z, Prkachin KM, and Solomon PE, "The painful face—Pain expression recognition using active appearance models," Image Vis. Comput, vol. 27, pp. 1788–1796, 2009. [PubMed: 22837587]
- [66]. Hammal Z and Cohn JF, "Automatic detection of pain intensity," in Proc. 14th ACM Int. Conf. Multimodal Interaction, 2012, pp. 47–52.

- [67]. Kaltwang S, Rudovic O, and Pantic M, "Continuous pain intensity estimation from facial expressions," *Advances in Visual Computing*. Berlin, Germany: Springer, 2012, vol. 7432, pp. 368–377.
- [68]. Martinez B, Valstar M, Binefa X, and Pantic M, "Local evidence aggregation for regression-based facial point detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 5, pp. 1149–1163, 5 2013.
- [69]. Ahonen T, Hadid A, and Pietikainen M, "Face description with local binary patterns: Application to face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 12, pp. 2037–2041, Dec. 2006. [PubMed: 17108377]
- [70]. Sikka K, Dhall A, and Bartlett M, "Weakly supervised pain localization using multiple instance learning," in *Proc. 10th IEEE Int. Conf. Workshops Autom. Face Gesture Recognit.*, 2013, pp. 1–8.
- [71]. Sikka K, Ahmed AA, Diaz D, Goodwin MS, Craig KD, Bartlett MS, and Huang JS, "Automated assessment of children's postoperative pain using computer vision," *Pediatrics*, vol. 136, no. 1, pp. e124–e131, 2015. [PubMed: 26034245]
- [72]. Littlewort GC, Barlett MS, and Lee K, "Automatic coding of facial expressions displayed during posed and genuine pain," *Image Vis. Comput.*, vol. 27, no. 12, pp. 1797–1803, 2009.
- [73]. Bartlett MS, Littlewort GC, Frank MG, and Lee K, "Automatic decoding of facial movements reveals deceptive pain expressions," *Current Biol.*, vol. 24, no. 7, pp. 738–743, 2014.
- [74]. Eisenberger NI, "The pain of social disconnection: examining the shared neural underpinnings of physical and social pain," *Nature Rev. Neurosci.*, vol. 13, no. 6, pp. 421–434, 2012. [PubMed: 22551663]
- [75]. Martel MO, Thibault P, and Sullivan MJL, "The persistence of pain behaviours in patients with chronic back pain is independent of pain and psychological factors," *Pain*, vol. 151, pp. 330–336, 2010. [PubMed: 20813461]
- [76]. Wu J and Rehg JM, "Centrist: A visual descriptor for scene categorization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 8, pp. 1489–1501, Aug. 2011. [PubMed: 21173449]
- [77]. Dhall A, Joshi J, Sikka K, Goecke R, and Sebe N, "The more the merrier: Analysing the affect of a group of people in images," in *Proc. IEEE Int. Conf. Automat. Face Gesture Recog.*, 2015, vol. 12, pp. 1–8.
- [78]. Suk M and Prabhakaran B, "Real-time mobile facial expression recognition system—A case study," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog. Workshops*, 2014, pp. 132–137.
- [79]. McDuff D, El Kaliouby R, Senechal T, Demirdjian D, and Picard R, "Automatic measurement of ad preferences from facial responses gathered over the internet," *Image Vis. Comput.*, vol. 32, no. 10, pp. 630–640, 2014.
- [80]. Hicks CL, von Baeyer CL, Spafford PA, van Korlaar I, and Goodenough B, "The faces pain scale—revised: Toward a common metric in pediatric pain measurement," *Pain*, vol. 93, no. 2, pp. 173–183, 2001. [PubMed: 11427329]
- [81]. Adams P, Rabbi M, Rahman T, Matthews M, Volda A, Gay G, and Volda S, "Towards personal stress informatics: Comparing minimally invasive techniques for measuring daily stress in the wild," in *Proc. 8th Int. Conf. Pervasive Comput. Technol. Healthcare*, 2014, pp. 72–79.
- [82]. Schiavenato M, von Baeyer CL, and Craig KD, "Self-report is a primary source of information about pain, but it is not infallible a comment on "Response to voepel-lewis's letter to the editor," bridging the gap between pain assessment and treatment: Time for a new theoretical approach?," *Western J. Nursing Res.*, vol. 35, no. 3, pp. 384–387, 2013.
- [83]. George JM, "Trait and state" affect," in *Individual Differences and Behavior in Organizations*, Murphy KR, Ed. Hoboken, NJ, USA: Wiley, 1996, pp. 145–171.
- [84]. Stone AA and Shiffman S, "Capturing momentary, self-report data: A proposal for reporting guidelines," *Ann. Behavioral Med.*, vol. 24, no. 3, pp. 236–43, 2002.
- [85]. Moskowitz DS and Young SN, "Ecological momentary assessment: What it is and why it is a method of the future in clinical psychopharmacology," *J. Psychiatry Neurosci.*, vol. 31, no. 1, pp. 13–20, 2006. [PubMed: 16496031]

- [86]. Rahman T, Zhang M, Volda S, and Choudhury T, "Towards accurate non-intrusive recollection of stress levels using mobile sensing and contextual recall," in Proc. 8th Int. Conf. Pervasive Comput. Technol. Healthcare, 2014, pp. 166–169.
- [87]. Pollak JP, Adams P, and Gay G, "PAM: A photographic affect meter for frequent, in situ measurement of affect," in Proc. SIGCHI Conf. Human Factors Comput. Syst., 2011, pp. 725–734.
- [88]. Yang L, Freed D, Wu A, Wu J, Pollak JP, and Estrin D, "Your activities of daily living (YADL): An image-based survey technique for patients with arthritis," ArXiv:1601.03278, 2016.
- [89]. Adams P, Adams A, Gay G, and Choudhury T, "Keppi: A tangible user interface for the self-report of scalar values," [Online]. Available at: <http://pac.cs.cornell.edu/>
- [90]. Abbas Z and Yoon W, "A survey on energy conserving mechanisms for the internet of things: Wireless networking aspects," Sensors, vol. 15, no. 10, pp. 24818–24847, 2015. [PubMed: 26404275]
- [91]. Klasnja P, Consolvo S, Choudhury T, Beckwith R, and Hightower J, "Exploring privacy concerns about personal sensing," in Pervasive Computing, Berlin, Germany: Springer, 2009, pp. 176–183.
- [92]. Rabbi M, Caetano T, Costa J, Abdullah S, Zhang M, and Choudhury T, "SAINT: A scalable sensing and inference toolkit," in Proc. Hotmobile, 2015.
- [93]. Press WH, Teukolsky SA, Vetterling WT, Flannery BP, Numerical Recipes in C: The Art of Scientific Computing, 2nd Ed., New York, NY, USA: Cambridge Univ. Press, 1992.
- [94]. Eugster PT, Felber PA, Guerraoui R, and Kermarrec AM, "The many faces of publish/subscribe," ACM Comput. Surveys, vol. 35, no. 2, pp. 114–131, 2003.
- [95]. Ross GJ, "Parametric and nonparametric sequential change detection in R: The CPM package," J. Statist. Softw, vol. 78, 2013.
- [96]. Junker U and Wirz S, "Review article: Chronobiology: influence of circadian rhythms on the therapy of severe pain," J. Oncol. Pharmacy Practice, vol. 16, no. 2, pp. 81–87, 2010.
- [97]. Chen Z, Lin M, Chen F, Lane ND, Cardone G, Wang R, Li T, Chen Y, Choudhury T, and Campbell AT, "Unobtrusive sleep monitoring using smartphones," in Proc. 7th Int. Conf. IEEE Pervasive Comput. Technol. Healthcare, 2013, pp. 145–152.
- [98]. Min JK, Doryab A, Wiese J, Amini S, Zimmerman J, and Hong JI, "Toss 'N' turn: Smartphone as sleep and sleep quality detector," in Proc. SIGCHI Conf. Human Factors Comput. Syst., 2014, pp. 477–486.
- [99]. Bai Y, Xu B, Ma Y, Sun G, and Zhao Y, "Will you have a good sleep tonight?: Sleep quality prediction with mobile phone," in Proc. 7th Int. Conf. Body Area Netw., 2012, pp. 124–130.
- [100]. Abdullah S, Matthews M, Murnane EL, Gay G, and Choudhury T, "Towards circadian computing: 'Early to bed and early to rise', makes some of us unhealthy and sleep deprived," in Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput., 2014, pp. 673–84.
- [101]. Saeb S, Zhang M, Karr CJ, Schueller SM, Corden ME, Kording KP, Mohr DC, "Mobile phone sensor correlates of depressive symptom severity in daily-life behavior: An exploratory study," J. Med. Internet Res, vol. 17, no. 7, 2015, Art. no. e175.
- [102]. Bengio Y, Courville A, and Vincent P, "Representation learning: A review and new perspectives," IEEE Trans. Pattern Anal. Mach. Intell, vol. 35, no. 8, pp. 1798–1828, Aug. 2013. [PubMed: 23787338]
- [103]. Schuller B, "Deep learning our everyday emotions," in Advances in Neural Networks: Computational and Theoretical Issues. Berlin, Germany: Springer, 2015, pp. 339–346.
- [104]. Martinez HP and Yannakakis GN, "Deep multimodal fusion: combining discrete events and continuous signals," in Proc. Int. Conf. Multi-modal Interaction, 2014, pp. 34–41.
- [105]. Eagle N and Pentland S, "Eigenbehaviors: Identifying structure in routine," Behav. Ecol. Sociobiol, vol. 63, pp. 1057–66, 2009.
- [106]. Zheng J, Liu S, Li LM, "Effective routine behavior pattern discovery from sparse mobile phone data via collaborative filtering," in Proc. IEEE Int. Conf. Pervasive Comput. Commun., 2013, pp. 29–37.
- [107]. Caruana R, "Multitask learning: A knowledge-based source of inductive bias," in Proc. 10th Int. Conf. Mach. Learning, 1993, pp. 41–48.

- [108]. Romera-Paredes B, Aung MSH, Pontil M, Williams A. C de C., Bianchi-Berthouze N, and Watson P, "Transfer learning to account for idiosyncrasy in face and body expressions," in Proc. IEEE 10th Int. Conf. Workshops Autom. Face Gesture Recog., 2013, pp. 1–6.
- [109]. Romera-Paredes B, Argyriou A, Berthouze N, and Pontil M, "Exploiting unrelated tasks in multi-task learning," in Proc. Int. Conf. Artif. Intell. Statist., 2012, pp. 951–959.
- [110]. Argyriou A, Evgeniou T, and Pontil M, "Convex multi-task feature learning," Mach. Learning, vol. 73, no. 3, pp. 243–272, 2008.
- [111]. Rosser BA and Eccleston C, "Smartphone applications for pain management," J. Telemed. Telecare, vol. 17, no. 6, pp. 308–312, 2011. [PubMed: 21844177]
- [112]. Lalloo C, Jibb LA, Rivera J, Agarwal A, and N Stinson J, "There's a pain app for that: A review of patient targeted smartphone applications for pain management," Clin. J. Pain, vol. 31, no. 6, pp. 557–563, 2015. [PubMed: 25370138]
- [113]. Richardson JE and Reid MC, "The promise and pitfalls of leveraging mobile health technology for pain care," Pain Med, vol. 14, pp. 1621–1626, 2013. [PubMed: 23865541]
- [114]. Reynoldson C, "Assessing the quality and usability of smartphone apps for pain management," Pain Med, vol. 15, pp. 898–909, 2014. [PubMed: 24422990]
- [115]. Singh A, Klapper A, Jia J, Fidalgo A, Tajadura-Jimenez A, Kanakam N, Bianchi-Berthouze N, Williams A, "Motivating people with chronic pain to do physical activity: Opportunities for technology design," in Proc. SIGCHI Conf. Human Factors Comput. Syst., 2014, pp. 2803–2812.
- [116]. Consolvo S, McDonald DW, Toscos T, Chen MY, Froehlich J, Harrison B, Klasnja P, LaMarca A, LeGrand L, Libby R, and Smith I, "Activity sensing in the wild: A field trial of ubifit garden," in Proc. SIGCHI Conf. Human Factors Comput. Syst., 2008, pp. 1797–1806.
- [117]. Lane ND, Mohammad M, Lin M, Yang X, Lu H, Ali S, Doryab A, Berke E, Choudhury T, and Campbell A, "Bewell: A smartphone application to monitor, model and promote wellbeing," in Proc. 5th Int. ICST Conf. Pervasive Comput. Technol. Healthcare, 2011, pp. 23–26.
- [118]. Consolvo S, McDonald DW, and Landay JA, "Theory-driven design strategies for technologies that support behavior change in everyday life," in Proc. SIGCHI Conf. Human Factors Comput. Syst., 2009, pp. 405–414.
- [119]. Kay M, Choe EK, Shepherd J, Greenstein B, Watson N, Consolvo S, and Kientz JA, "Lullaby: A capture & access system for understanding the sleep environment," in Proc. ACM Conf. Ubiquitous Comput., 2012, pp. 226–234.
- [120]. Choe EK, Lee B, Kay M, Pratt W, and Kientz JA, "SleepTight: Low-burden, self-monitoring technology for capturing and reflecting on sleep behaviors," in Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput., 2015, pp. 121–132.
- [121]. Pellegrini CA, Hoffman SA, Collins LM, and Spring B, "Optimization of remotely delivered intensive lifestyle treatment for obesity using the multiphase optimization strategy: Opt-IN study protocol," Contemporary Clin. Trials, vol. 38, no. 2, pp. 251–259, 2014.
- [122]. Kukafka R, "Tailored health communication," Consum. Health Inform.: Inform. Consum. Improving Health Care, pp. 22–33, 2005.
- [123]. Rabbi M, Aung MSH, Zhang M, Choudhury T, "MyBehavior: Automated personalized health feedback from user behavior and preference using smartphones," in Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput., 2015, pp. 707–718.
- [124]. Fogg BJ, "A behavior model for persuasive design," in Proc. 4th Int. Conf. Persuasive Technol, 2009, vol. 40, pp. 1–7.
- [125]. Singh A, Piana S, Pollarolo D, Gualtierio V, Varni G, Tajadura-Jimenez A, Williams A, Camurri A, and Bianchi-Berthouze N, "Go-with-the-flow: Tracking, analysis and sonification of movement and breathing to build confidence in activity despite chronic pain," Human Comput. Interaction, vol. 31, no. 3/4, pp. 1–40, 2016.

Fig. 1.

Calendar view with the daily color coded mobility index for RA management. As described above the clinical validity of this measures is undergoing evaluation within RA and other orthopedic contexts.

Fig. 2. (a) Amplitude values for 5 second recording of audio. (b) Spectrogram of the recording with blue lines showing inferred segments of human voice.

Fig. 3. Example screenshots from PAM (left) prompting the self-report of emotional state using representative images and YADL (right) prompting selection of difficult daily activities as a proxy to pain level.

Fig. 4.
Prototype of *Keppi*—tangible and portable self-reporting device for pain.

Fig. 5.

The architecture of SAINT sensing and inference framework. SAINT provides a unified bus interface to share data across sensing and inference modules. Client applications can connect with SAINT to receive sensed and inferred data.

Fig. 6.

Lifestreams stack, data enters from the lowest level and is successfully filtered, aggregated and analyzed to produce actionable visualizations for patients and clinician.