

Instrumenting and Analyzing Fabrication Activities, Users, and Expertise

Jun Gong^{*†}, Fraser Anderson^{*}, George Fitzmaurice^{*}, Tovi Grossman^{*‡}

^{*}Autodesk Research
Toronto, ON, Canada
{first.last}@autodesk.com

[†]Dartmouth College
Hanover, NH, USA
jun.gong.gr@dartmouth.edu

[‡]University of Toronto
Toronto, ON, Canada
tovi@dgp.toronto.edu




Figure 1: Overview of the identification of users, recognition of activity and inference of expertise from contextual data.

ABSTRACT

The recent proliferation of fabrication and making activities has introduced a large number of users to a variety of tools and equipment. Monitored, reactive and adaptive fabrication spaces are needed to provide personalized information, feedback and assistance to users. This paper explores the sensorization of making and fabrication activities, where the environment, tools, and users were considered to be separate entities that could be instrumented for data collection. From this exploration, we present the design of a modular system that can capture data from the varied sensors and infer contextual information. Using this system, we collected data from fourteen participants with varying levels of expertise as they performed seven representative making tasks. From the collected data, we predict which activities are being performed, which users are performing the activities, and what expertise the users have. We present several use cases of this contextual information for future interactive fabrication spaces.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300554>

CCS CONCEPTS

Human-centered computing → Human computer interaction (HCI)

KEYWORDS

Fabrication; activity recognition; user identification; skills evaluation;

1 INTRODUCTION

Fabrication and making activities have increased in popularity with the general population as well as in research [6, 12, 45, 48]. These activities have become more accessible to a wider range of individuals with varying abilities, requirements and skills, but the tools and environments used in the process do not sense or adapt to the fabrication context.

Prior work has shown the benefits of detecting the workflow or expertise of a user to create adaptive interfaces that respond to a user's context and situation [32, 40]. These interfaces can adapt to the needs and goals of the user [26] or provide information, feedback, or guidance that is more appropriate to the user's expertise and tasks. Likewise, recent work has also demonstrated the need for identifying users, their skill sets, and tasks in fabrication spaces through interviews and observation of approximately 85 people across 11 local workshops [10]. It envisions a "hybrid workshop" that can exploit contextual information to "create a personalized

space”, “tailor assistance based on users’ current skill sets” and “guide users towards equipment, practices, and workflows”. We believe that having more contextual information will allow for novel applications not yet explored, for example, by monitoring the tools and activities that are occurring within a workshop, a manager of the space could have a better awareness of what is happening in the shop at any given time, and use the historic information for decision support and space planning.

There is a large amount of work within the ubiquitous and pervasive computing literature that examines how to sense contextual information from various data sources [18, 19, 55]. Often this involves automatically detecting activities, occupancy, or usage patterns by instrumenting environments, objects and people with sensors [27, 35]. However, unlike work on physical activity recognition, within the domain of fabrication [14, 20, 21, 57], we explore sensorization by instrumenting the environment, tools, and users involved in fabrication to detect combination of activities, users and their expertise.

This paper makes three contributions: (1) an exploration of how the activities, users and expertise can be detected by using a Random Forest classifier to evaluate a combination of 29 independent data streams from sensors located in the environment, fixed to the tools and equipment, and worn on the person performing the activities, (2) a study which evaluates the feasibility of this approach across seven representative fabrication tasks, and fourteen users with varying expertise levels, (3) a discussion of the relative merit of environment, tools and user-based sensors. Using all sensors available, the system predicted user identity with 93.8% accuracy, predicted the performed activity with 94.2% accuracy, and could predict the expertise with an average RMSE of 1.3 on a 7 point scale, which was similar to the human experts’ accuracy. Building on our study results, we discuss design implications and provide a set of sample applications which could make use of the contextual information inferred from this newly available data.

2 RELATED WORK

This work builds and extends upon prior work in a number of domains, including context and activity recognition, the identification of skills in physical tasks and user identification.

2.1 Context and Activity Recognition

Building from Weiser’s initial paradigm [67] of a ubiquitous computing environment that responds to user’s activities and situations, there has been a multitude of research investigating sensing and context recognition. A full review of all approaches work in the area is beyond the scope of this paper, but summaries can be found in several papers [1, 11, 16]; here we outline the more relevant prior work.

Several prior research projects have examined sensing and instrumentation of an environment to infer context. A number of projects have examined being able to sense activities of daily living through cameras, wearable sensors, and sensors embedded in objects and environments [22, 36, 46, 49]. These approaches typically collect multi-dimensional data streams which are used to train a classifier that can later identify the context or activity being performed. Closer to our work, recent approaches have presented novel sensing mechanisms that can be worn on the body to detect when different appliances or machines or tools are being used, such as Ward et al. [66] and Lukowicz et al. [41] who used microphones and accelerometers to recognize different activities within a workshop context (e.g., sawing, drilling, etc.) This approach was extended into a car manufacturing context, where the wearable sensors were used to track subsequences in an assembly process [61]. More recently, generalized sensor packages, such as SyntheticSensors by Laput et al. have been used to detect several environmental states, including when some workshop equipment was powered on through a combination of sensing modalities [35]. This paper extends these prior approaches by considering the context of fabrication and analyzing how the user, tools and environment can be instrumented to provide different types of contextual information. This paper goes beyond identifying *what* activity is being carried out and attempts to identify *how* that activity is being performed and by *who*.

2.2 Identification of Skills in Physical Tasks

Within many domains outside of HCI (e.g., medical, sports, dance), there are a number of approaches used to measure the expertise of the individual, many of which are tailored to the particular domain. To evaluate a surgeon’s performance on a laparoscopic cholecystectomy, Rosen et al. recorded the force and torque applied to the laparoscopic instruments in 3D space and used these features to train Hidden Markov

Models on novice and expert performance [54]. This approach is representative of many of the prior approaches within the medical domain where the tools or objects are instrumented with spatial or force sensors, and motion analysis is used to predict expertise [7, 29, 42]. Other approaches have instrumented the individual and tracked their bodily movements through motion capture of the upper body, an approach that's common in sports and dance [3, 4] in addition to medicine [39, 53]. We build on this work by integrating sensors throughout the environment, in addition to on the tools and individuals, and extend it to a novel domain.

More directly related to the current work is recent work on evaluating the skill of tradespeople [20, 21]. Enokibori and Mase used wearable accelerometers and gyroscopes [20] to measure an individual's skill by analyzing movement during a metal filing task, finding a visible difference between coaches and learners in three measures computed from the sensor data. Erden and Tomiyama used motion analysis to find differences between novices and experts in the amount of speed and position deviation [21]. We extend these works into broader and more generalized fabrication tasks with a much richer set of data and more diverse expertise. We also evaluate the utility of physiological data for expertise inference.

2.3 User Identification

User identification technologies span a variety of techniques and technologies. Vision-based methods have already proved robust and accurate in user identification and have been applied in many authentication applications [50, 65], and other biometric techniques have drawn attention, such as the identification of users from body-worn accelerometers or mobile phones [25, 59, 64]. More recently, user identification was achieved by sensing the body's electrical properties through electrical frequency response sensing [56]. We build on these works and try to identify users from their unique characteristics – in particular how they use tools, and their natural actions as they perform fabrication activities.

3 INSTRUMENTATION FOR FABRICATION

Prior work has explored instrumenting spaces, people and objects in a variety of use cases, including activity recognition, person identification, and identification of other contextual information [11]. Fabrication

activities often occur in spaces that are transitory, dynamic, noisy, dangerous, and require manual dexterity [10, 31]. Given these considerations, identifying appropriate instrumentation to sense the context of these activities remains challenging.

3.1 Sensor Types

A wide range of sensors are available which may be able to provide contextual information in the space of fabrication and may have unique considerations within this domain.

Vision

Cameras provide rich information about a scene, which can be obtained through machine learning and computer vision algorithms to yield sensor-like feeds. There is a large body of work in video-based sensing [24, 44] to determine states, activities, and changes in an environment. While these computer vision based methods are powerful, we deliberately excluded cameras in this work since the spaces that fabrication can take place in can be large, transitory and dynamic, the issue of occlusion can cause loss of information from camera streams.

Audio

Audio data has proved useful in different context-aware applications [35, 41] for sensing the state of the environment and identifying what machinery or appliances are being used. While the fabrication space can be quite noisy (e.g., a running air compressor), we still chose to exploit audio data for capturing contextual information due to the promising results found in the prior work [35].

Motion

Inertial measurement units (IMUs) include an accelerometer, gyrometer and magnetometer, which can be used to detect orientation, movement, and even tiny vibrations of the instrumented objects. Recent work has demonstrated the unique sensing capabilities of accelerometers [33] that are sampled at a high frequency, and even at lower frequencies accelerometers can capture unique signatures of movement [51]. We used an accelerometer sampled at a high frequency to capture vibration signatures while different users were performing different tasks. We also exploited a lower sampling rate IMU to track the tool's orientation and vibration.

Ambient Environment Sensors

Temperature, humidity, and light sensors are quite common in IoT devices, which can be used to monitor

the ambient environmental conditions. As machinery is used, it may give off some heat or changes in how the tools are held may cause changes in the amount of ambient light that is detected. While the sensors may not be a rich source of information, they may provide some usable information at a very low cost given their integration into many existing electronics.

Biometric Sensors

Given that fabrication tasks inherently involve an individual, biometric sensor (e.g., heart rate sensor, eyetracker) may provide unique information as the user performs the task. These data streams may capture a user's confidence, their familiarity with the task, or other inherent data that may be impossible to detect with more indirect methods [23, 28]. Given that most biometric data requires direct instrumentation of the user, safety issues must be considered, and the sensors must not impede the user's natural dexterity and operations of the hands. Thus, we used the off-the-shelf products (e.g., Apple Watch, Pupil-labs Eye Tracker) to obtain biometric data.

3.2 Sensor Placement

As fabrication activities are conducted by an individual, using a tool, within an environment, it is worth examining the value of placing sensors at or on each of these entities. Depending on the context, it may be possible to instrument all three of the entities, or some subset depending on the technology available, the environmental context, and other factors.

Environment

In several fabrication contexts (within a workshop, machine shop, makerspace, or on a construction site), the environment where fabrication takes place is relatively fixed. Several tools may be used within the space, and many different individuals may be present in the space performing the activities. In many instances, it may be possible and useful to place fixed sensors within a space to record what is happening inside that space (e.g., who is in the space, what they are doing, and what skill level they have). However, fixed environmental sensors may not have the fidelity necessary to capture the desired information, as sensors may be limited in bandwidth, or come with privacy concerns.

Tools

As nearly all fabrication activities require tools (manual, power, digital, etc.), sensorizing these tools may provide valuable input. By placing sensors on the

tools, the system may be able to understand how the tool is being operated, what task is being completed, and how well the operator is performing the task. Instrumenting the tool is relatively unobtrusive with small sensors and electronics having minimal impact on the natural operation of the tools, and minimal impacts to privacy. However, instrumenting tools can be costly, with many tools being used throughout a fabrication process. Tools may also offer limited information, only able to sense how they are being manipulated and used, and not necessarily the environment or task they are being used in, or the intent or goals of the operator.

Users

As the fabrication activities of interest are carried out by users, it is beneficial to consider what data can be collected directly from the person performing the action. With this approach, the sensors stay with the person regardless of where they move throughout the environment or what tools they are using. This type of instrumentation has the potential to be the richest source of information, as it is closest to the user and can detect their focus, physiological measures, and movement. However, instrumenting the user requires them to perform additional steps prior to fabrication (e.g., equipping the sensor array and performing necessary calibrations) and has the highest impact on the user's actions and dexterity. Additionally, requiring users to equip physiological sensors comes with additional privacy concerns due to the ability to monitor their activities outside the scope of the fabrication activities.

4 SYSTEM IMPLEMENTATION

A modular system was developed to record all sensor data in a synchronous manner.

4.1 Sensor Selection

A wide array of sensors is available on the market to measure a plethora of physical phenomena. While some niche sensors, such as radar or electromagnetic sensors [34, 37] have been used for different sensing and interaction techniques, we chose to focus on commercially available sensors that could be readily deployed in a relatively unobtrusive manner that would likely be able to detect the actions and activities of fabrication.

Environment

Building on prior work on environmental sensing and activity recognition [35], an accelerometer (MPU9250)

sampled at 4kHz, and an acoustic microphone (ADMP401) sampled at 17kHz were used. These sensors are small, inexpensive and can be placed discreetly in the environment. The sampling rate on these sensors is sufficiently high to capture many types of activities. While other environmental sensors may provide useful information, prior work has demonstrated that many activities can be readily recognized using only these two modalities [35].

Tools

All tools were instrumented with a TI SensorTag [68]. This small, wireless sensor contains a 9-axis IMU, as well as ambient light, humidity, temperature, barometric pressure, and magnetic sensor. We use the SensorTag as it contains a wide variety of sensors and it has a compact size, and the wireless nature would have minimal effect on user's operation. For our purposes, we only used the 9-axis IMU and ambient light sensor. The firmware was modified such that the 9-axis IMU was sampled at 80-100Hz while the ambient light sensor was sampled at 10Hz. The variable sampling rate for 9-axis IMU was due to changes in battery voltage, and samples were interpolated during post-processing to achieve a constant 100Hz sampling rate.

Users

For instrumenting users, the primary consideration is unobtrusiveness and having minimal effect on user's operation. Therefore, we use off-the-shelf wearable devices to capture users' data. Each person was instrumented with an eye tracker from Pupil-labs [69], which recorded the eye position at 120Hz relative to a head-worn camera which captured video of the scene the person was looking at. Each person was also outfitted with two Apple watches [70] – one worn on each wrist. The watches recorded the movement of each hand via accelerometers and gyrometers sampled at 100Hz, and recorded the heart rate, sampled at 0.2Hz.

4.2 System Architecture

To record and synchronize the data from the disparate sensor sources, a custom system was developed. Software written in Java ran on a laptop running Mac OSX to collect data from the smartwatches, environmental and tool sensors. Recording software for the eyetracker also ran on the same laptop capturing and processing the video and eye position data from the device.

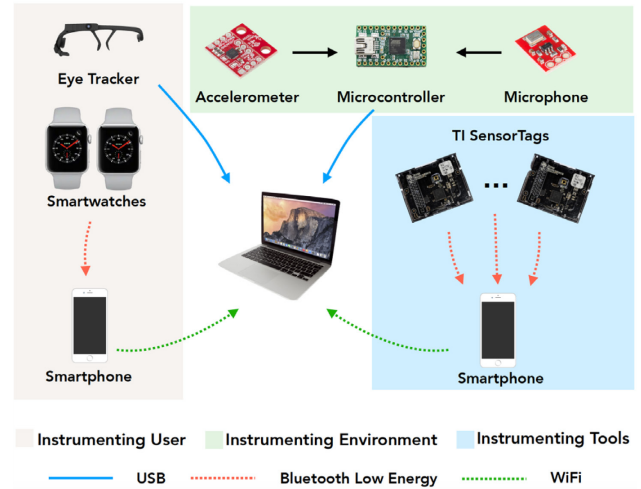


Figure 2: Overview of the data capture system, with one central computer integrating data from a wide range of sensors.

Figure 2 shows the overview of our data capture system. The environmental sensors (microphone and accelerometer) are soldered to a custom PCB, where they were sampled by a Teensy 3.2 [71], which communicated with the laptop via USB. Each SensorTag instrumented on the tools communicated via BLE to an Apple iPhone, which transmitted the data back to the laptop via WiFi. The smartwatches were each connected to Apple iPhone, where a custom application which read the data from the sensor and relayed that to the laptop via WiFi. The eyetracker transmitted video and eye position data via USB.

5 DATA COLLECTION EXPERIMENT

A study was conducted to collect data for representative fabrication tasks.

5.1 Participants

We solicited participation from a broad audience within our institution and asked them to complete a screening questionnaire. This questionnaire was used to collect demographic information, handedness, vision, overall expertise with fabrication, and expertise on the seven specific fabrication tasks that were being tested. From the 25 responses to this questionnaire, participants were selected based on their self-assessed expertise to maximize the range of skills that participants had. Finally, 14 participants were selected who were right handed (13) or ambidextrous (1), had normal (9), or corrected to normal vision using contact lenses (5). All respondents to the questionnaire were entered for a draw to receive a \$25 gift card.

All participants (12 male, 2 females; 20-35 years old) had on-site workshop safety training, as required by our institution, were aware of the risks of using power tools, and provided informed consent to participate in the study. All 14 participants received a \$25 gift card for completing the study.

5.2 Apparatus

The experiment was conducted within an active workshop, using the system described previously. The workshop is approximately 1000 sqft, and contains a number of 3D printers (FDM, SLA, SLS), a CNC machine, electronics equipment, and a wide variety of hand tools. While data was being collected, the rest of the workshop was open to its users and other tools and equipment were in use.

Participants were outfitted with the wearable sensors described previously – a mobile eye tracker, and Apple watches worn on both wrists.

Six different tools were instrumented and used in the study. Sensors were placed on each tool in a location that minimized their interference with the natural operation of the tool. Different locations on the tool may influence the signals, though we empirically tested several locations and found little difference in readings, likely as a result of the tools having rigid structures which transfer vibrations and movements without dampening. A DeWalt 18V cordless drill/driver was used for the drilling and screwing tasks, with the SensorTag fixed to the base of the tool. A 15" Stanley handsaw was used for the sawing task, with the SensorTag fixed to the side of the blade, near the handle. A YIHUA 907F soldering iron was used for the soldering task, with the SensorTag fixed to the iron, near the cable. A swivel-head deburring tool was used for the deburring task, with the SensorTag attached to the top of the handle. A ball-peen hammer was used for the nailing task, with the SensorTag attached to the hammer halfway up the shaft. A Delta drill press with a sanding drum spinning at 1150 RPM was used for the sanding task, with the SensorTag fixed to the front of the drill press. Figure 3 shows six instrumented tools and sensor locations.

The environment was outfitted with the microphone and accelerometer module described previously. For the sanding task, the module was mounted to the work table of the drill press. For all other tasks, the sensor was mounted on the workbench where the activities were taking place.




Figure 3: Instrumented tools used in the data collection study.

Two video cameras captured video and audio of the tasks for later review but were not used as part of the data capture process.

5.3 Design

Each participant performed 7 unique tasks: deburring, nailing, screwing, sawing, drilling, soldering, and sanding. For each task, participants performed two blocks in succession, each of four trials, resulting in 7 x 4 x 2 total trials per participant. The presentation order of the tasks was counterbalanced across the 14 participants. The study took approximately 60-80 minutes to complete.

Between the two blocks, participants watched an instructional video which contained a tutorial on how to perform that task. This experimental design was intended to allow the capture of a wider range of behavior, from novice to more expert.

5.4 Procedure

All participants had received workshop safety training prior to participating in the study. This training covered general hazards, safety, and protocols in the workshop, but did not include any explicit instruction relevant to the tasks being studied.

Prior to performing the tasks, participants were outfitted with the two smartwatches and the eye tracker. The eye tracker was then calibrated for the participant by adjusting the camera to capture the eye location and performing a routine to calibrate the eye-camera coordinates to the coordinates of the head-worn video camera.

For each task, participants were given a piece of paper that had an image of the desired result for the task (e.g., a photo of the circuit board with the pin soldered into place). No further task guidance, except in cases where the experimental proctor intervened for safety reasons. Only two interventions took place, one where the participant attempted to set the screw in the wood

using the battery pack of the drill, and another when the same participant was about to sand the wood without placing it on the sanding table.

For each trial of the deburring task, participants were asked to use the deburring tool to clean the inside edge of a roughly-cut aluminum pipe. For the nailing task, participants were asked to hammer a nail into a 2" by 2" piece of $\frac{3}{4}$ " thick MDF at a marked location near one of the corners. Similarly, for screwing, participants were asked to use the cordless drill to screw a Robertson-head $\frac{3}{4}$ " screw into a 2" by 2" MDF wood. For sawing, participants were asked to use the hand saw to make a straight and clean cut at a marked location on a 2" x 1" piece of pine. For drilling, participants were asked to use the cordless drill and drill a 5mm hole through a piece of acrylic at a marked location near the corner. For soldering, participants were asked to solder separate male headers on a general-purpose PCB. For sanding, participants were asked to sand a 2" by 2" $\frac{1}{4}$ " thick birch plywood to the marked contour with the spindle-sander and drill press.




Figure 4: Seven physical tasks in the user study.

For all tasks, the tools were setup by the experimenter (e.g., installed a proper drill bit and set correct speed and clutch) and materials were secured in place (e.g., clamp the material) before the participant performed the tasks. In this work, we only focus on the task operation, but it would be interesting to explore how the preparation process can be used to infer contextual information (e.g., detect expertise prior to the task beginning). The participant then performed the first block of trials, repeating the task four times, verbally indicating the beginning and end of each task. After each block of trials, the participant recorded a self-assessment of their performance on each of the four trials.

6 DATA ANALYSIS

From the collected data, we explore three elements of contextual information that would help fabrication tools and environments adapt to the user: (1) identification of which user is performing the task, (2) which task the user is performing, (3) what expertise the user has. We also aim to build an understanding of how different sensors in different locations can contribute to identifying this information.

We use machine learning to identify users, recognize physical tasks and infer user expertise. In particular, we use Random Forest in our current implementation. Random Forest has previously been found to be accurate, robust, scalable, and efficient in many different applications [15, 37].

6.1 Feature Extraction

Extracting relevant and meaningful features is critical to obtain the contextual information. Prior work has demonstrated successful identification of different physical activities from high frequency sensor data [35], which we adapt to our system's equivalent sensors and derive similar features for the unique sensors in our system.

For each trial, data from different types of sensors are divided into a set of instances via a 0.5 second sliding window (10% overlapping between each window). Within each data instance, we calculate spectral features (FFT) and seven statistical features (mean, STD, min, max, range, sum, energy) for each sensor data except for heart rate, gaze position and light intensity. As the heart rate is only sampled at 0.2Hz, we compute the mean of the heart rate readings on both watches that are closest to the time window. For gaze position data, we are interested in whether the user is looking at a fixed point while performing the task. Thus, we calculate the number of saccades (operationalized as changes in position at greater than approximately 5° per second). Due to the low sampling rate of the light sensor, we only compute statistical features. Note that we use 1024 FFT points for high-sample-rate sensors (e.g., environmental microphone and accelerometer) and 32 FFT points for low-sample-rate sensors (e.g., IMU sensors on Apple Watch and SensorTag). The phase information from the FFT is not used. The feature data for every sensor is then concatenated into a single feature vector and used to train the machine learning model.

6.2 Results

We test the prediction capabilities of the trained model for user identification, physical activity recognition and expertise inference.

User Identification

Classification accuracy. In order to overcome the possible strong correlation in data samples due to the high data frame rates of the sensor and the time series nature of data, we compute the accuracy using a leave-one-block-out approach for each task by training the model using the data from one block and testing it using the data from the other block. The model is trained on half the data, then tested on the other half, this process is done twice (with the two data sets swapped), and the overall accuracy is computed as the average prediction accuracy between those two tests. While we are intentionally trying to change the user's behavior between the two behaviors between the two blocks, we suspect there is unique characteristics for each user that persist even when their skill improves. We also tested with a leave-four-trials-out approach with slightly better results.

The overall average accuracy (Figure 5) across all seven tasks with all sensor data was 93.8% (SD = 3.2%). And the average accuracy dropped to 87.0% (SD = 5.3%), 79.5% (SD = 7.7%) and 59.8% (SD = 13.0%) when only using the sensor data from environment, tools and user respectively. From the results (Figure 5), several interesting patterns can be seen. First, it was difficult to identify different participants using the data from the user-worn sensors on the deburring task. We believe this is because of the complex nature of the hand movements produced in the deburring task – some participants would run the tool around the inside repeatedly, then use a bimanual approach where they rotated the pipe and tool simultaneously; very few participants exhibited consistent movement patterns within a single trial. Second, nailing achieved the most balanced performance using these three types of data, which meant that we could capture meaningful information from all three data sources in order to identify which user is hammering across a number of instrumentation scenarios. Third, sanding achieved the lowest accuracy with all sensor data available, which may be due to similar sanding operation behaviors among all participants. Overall, the environmental and tool-mounted sensors performed better than the user-worn sensors in identifying which

participant was performing the tasks. This is likely due to the tool-mounted sensors capturing the motion of the tool and the environmental sensors capturing the tool's interaction with the material. Unique individual movements of the wrists are not captured by the smartwatches, and those movements may contain inherent information in how the task was being performed.




Figure 5: Leave-One-Block-Out classification results for User Identification.

Sensor importance. Beyond the classification accuracy, we were also interested in which sensor played more important roles in user identification. To understand relative feature importance, a weighted breakdown of merit was calculated by Random Forest weights when all sensor features were supplied to the classifier (Figure 6, percentages are calculated from normalized Random Forest weights). As expected, the microphone and accelerometer from the environment were weighted quite high compared to all sensors. Interestingly, the magnetometer on the tools was weighted highest. We believe this is due to the unique way in how many participants held the tools (i.e., tool orientation).




Figure 6: Sensor feature merit for user identification, activity recognition and expertise inference.

We also examined the top ten features computed from the sensors. It turned out that statistical features (e.g., mean, sum, energy) played a more important role than spectral features among most sensors except for the microphone. This may be because the unique signatures that the user is producing are a result of gross motor movements and posture and positioning, not in the higher-frequency, cyclic signals resulting from things like the drill or sander spinning.

Activity Recognition

Classification accuracy. Similarly, we tested whether the machine learning model worked well across different users. To calculate the accuracy, we used the data from thirteen participants for training and the remaining participant for testing. The overall accuracy was then calculated by averaging the results from all fourteen combinations of training and test data. The overall average accuracy with all sensor data (Figure 7, left) was 94.2% (SD = 7.7%). The average accuracy dropped to 78.3% (SD = 14.9%), 89.3% (SD = 13.5%) and 80.5% (SD = 11.5%) when only using the sensor data from environment, tool and user respectively.




Figure 7: Activity classification accuracy for leave-one-subject-out (left); Confusion Matrix when using all sensors data (right).

The confusion matrix (Figure 7, right) illustrates that it was difficult to clearly distinguish between screwing and drilling since these two tasks used the same tool and had similar procedure to operate, though accuracy was still relatively high for those two tasks (around 80%) with all data. More exploration is needed to be done to further investigate how to reliably distinguish between activities that utilize an equivalent tool set with similar processes, perhaps new sensors (e.g., force sensors) could be used, or capturing the activities before and after the task performance (e.g., loading a screw) may provide more insight.

Sensor importance. We also calculated the sensor feature merit for activity recognition (Figure 6, orange bars). Unlike user identification, accelerometers and

gyrometers on the tools contributed more to the results, with a strong contribution from the environmental accelerometer as well. We believed this is due to inherent similarities in how the tools are handled, and the unique signatures that they cause during operation (e.g., Drilling vs. Sanding have unique frequency responses due to the different rotation speeds among other differences). And it was also clear that accelerometer and gyrometer in the Apple Watch worn on the right hand contributed more than the same sensors from left hand – this is representative of the fact that all actions were performed primarily with the dominant (right) hand, with the left hand used for support and largely idle or supportive in several tasks.

Expertise Inference

Expertise is a relatively complex and subjective measurement, and different individuals will judge expertise by differing criteria. In this work, we try to provide initial insights into how expertise can be inferred from data collected during fabrication tasks.

Ground truth. Two experts in fabrication (one author) watched the videos clips of participants performing the tasks. Each expert gave a rating from one to seven for each trial based on their opinion and expertise. A Spearman's rank correlation shows a highly significant correlation between the ratings ($p < 0.001$; $rs(784) = 0.7$). The mean of these two ratings was used as the ground truth.

Feature extraction. While a 0.5 second sliding window is long enough for user identification and physical activity recognition as those are relatively continuous events, we believe it is too short for expertise inference. Expertise is evident in how the tools and material are handled, and manifested in different ways throughout the task (e.g., starting the saw cut with slow pull-strokes before making more aggressive oscillating cuts). Thus, instead of using a sliding window and creating several data instances for each trial, we compile a single feature vector for each trial. We use the same features as in user identification and activity recognition, but they are used differently. For statistical features, they are calculated using the data from a whole trial, while for spectral features, we use a sliding window of 1024 points, but then collapse the windows into a single window of length 1024 by computing the mean for each point in the window.

Measurement. Since we want to obtain a continuous evaluation score of user's expertise, we use Random Forest Regressor instead as the model for expertise

inference. And we evaluate the results using root-mean-square error (RMSE).

Regression Accuracy. For each task, we used the data from thirteen participants for training and inferred the expertise on the remaining participant. The overall RMSE was then calculated by averaging the RMSEs from all fourteen combinations of training and test data. The overall average RMSE with all sensor data was 1.27 (SD = 0.45). The average RMSE increased when only using the sensor data from environment (1.35 (SD = 0.46)), tool (1.42 (SD=0.43)) and user (1.42 (SD = 0.43)). While an RMSE of 1.3 does not allow the system to reliably predict the expertise of a user to within a single point on the seven point scale, it is more than enough to distinguish novices from experts, and it reflects some gradation in skill. Additionally, the RMSE between the two expert ratings across all trials of all tasks was 1.37, indicating that expertise evaluation is somewhat subjective, and the automated approach performs very similarly to the human evaluator.




Figure 8: Leave-one-subject out results for expertise inference.

Among the seven different tasks, soldering skill appeared to be the most distinguishable, possibly due to the fact that the experts held the soldering iron with more stability than the novice, which could be easily captured by the accelerometers. Screwing and drilling were also well-recognized, possibly due to the differences in tool orientation between novices and experts.

Sensor Importance. Environmental sensors (e.g., accelerometer and microphone) were still the most useful sensors in expertise inference (Figure 6). We believe this is because different operations from experts and novice users can lead to different patterns of sound and vibration. For example, experts tend to start the saw cut with slow pull-strokes before making

more aggressive oscillating cuts. Compared with user identification and activity recognition, sensors on the users, especially the accelerometer on the user's right hand, played more important role in expertise inference. One possible reason was that stable hands during operation were strong indicators for expertise, which was echoed by relative good regression performance. Similar effects may be seen in drilling and screwing as well, as novices tended to have a more wobbly approach to drilling and a less consistent application of force. It was interesting to find that we captured useful information from the accelerometer and gyrometer on user's left hand for expertise. We believe that while the non-dominant hand does not play a large role in task identification, its use may be an indicator of expertise as experts used it for support, stability, or guidance while novices had a much less consistent approach to the actions of their left hand.

Summary

By instrumenting the environment, tools and users with different types of sensors, we can achieve promising results of user identification, activity recognition and expertise inference. Among the different kinds of sensors, the accelerometer on the environment has shown significant usefulness in identifying these three types of contextual information. Instrumenting the environment and tools provides more information than only instrumenting users, however, sensors on the users appear to have unique benefit for expertise inference.

Interestingly, we did not find that biometric data including gaze position and heart rate conveyed much information. Heart rate helps to identify different users but does not contribute much to activity recognition and expertise inference. This is likely due to the little heart rate variation across tasks (e.g., most tasks were of similar difficulty and would not cause participants to worry or feel more confident). The gaze data that was used did not seem to be of much value in providing contextual information. One possible reason is that when users are performing the physical tasks, the camera and tracker would occasionally shift on the face, causing errors in alignment. Additionally, a very coarse measure was used from the eye-tracker (number of saccades), as it was a simple measure that could be computed independent of the tasks and scene. However, more complex eye-tracking or scanpath comparison methods may lead to more useful information from eye gaze [9].

7 SAMPLE USE CASES

To illustrate the potential value in being able to identify expertise, users and activities during making, we outline several sample use cases.

7.1 Adaptive Fabrication Environment

By leveraging the newly available contextual information, and an active workspace [10, 31], the fabrication space could respond intelligently to a user's needs. As the system detects that the user is drilling holes in wood, it could highlight the locations of the bolts, screws, or other drill bits via projection. It could also project the drilling state in the environment, or render it on a public display to provide other users with an awareness of what is happening in the space. An adaptive environment could also enforce safety measures if it detects that a user has insufficient skill with a particular tool. For instance, if the system detects the user has limited expertise in drilling through plastic, the system may lockout certain drawers, such as those containing metal drilling bits which are more difficult and potentially more dangerous to use.

7.2 Workshop Management

By tracking the users, activities and skills within a workshop or makerspace, the managers or coordinators of the space could have better tools and information at their disposal as they oversee the operation of their space. A dashboard that visualizes the types of activities currently happening in the space could help the managers understand the frequent and typical activities that occur within the space to help them plan for the future usage of the space. By observing the expertise of the users in the space, the manager could understand how effective their training programs are for different aspects, or could use such a view to help facilitate peer learning by identifying users who could learn from each other's expertise.

7.3 Adaptive Tutorials and Dynamic Learning Content

If the fabrication environment or system can detect a user's expertise using a given tool or performing a task, it could provide more appropriate guidance or instruction to the user, helping to realize the idea of digital apprenticeship [14]. If a user is sawing wood with low expertise, the system could detect this and suggest video tutorials or other media that would offer training. Additionally, by being able to recognize which user is performing the action, the system can track which videos were shown to that user and show

a variety of different videos over time to help users achieve greater levels of expertise. Additionally, if relatively high levels of expertise were detected, the system could provide tips and hints from expert craftsmen, or recommend alternative and more advanced workflows or tools in a non-intrusive manner similar to Ambient Help [43].

7.4 Automated Workflow Capture

If activities and expertise can be captured and identified, then the system can begin to build models of common workflows used in fabrication contexts. By recognizing the workflows used by experts, the system can begin to build an understanding of efficient fabrication processes, and use that knowledge to assist novices in navigating their tasks. For instance, a system may suggest to novices to perform sanding after drilling, as that ordering is more common amongst experts than novices as it avoids potentially having to repeat a sanding step.

8 DISCUSSION

Overall, our results suggest that adding instrumentation can aid in the detection of contextual information for fabrication tasks, but there are important considerations when determining how that instrumentation should take place.

8.1 Instrumenting for Context

Depending on the use of the contextual information, issues of privacy and cost, different factors may impact what sensors are deployed and used to gather the data. For simple user identification tasks, no advanced sensor processing may be needed at all – badge-ins to access the space, or simply syncing a personal smartwatch to a central server may provide all the identification that is necessary to achieve the desired levels of user identification. However, by instrumenting the environment and tools, the space can track who is using what equipment and have an idea of how the space was used, or who contributed to certain projects without any additional information from the user.

Similarly, activity recognition could largely be performed by instrumenting each tool, as in the Smart Makerspace [31]. However, this approach could be costly and cumbersome to instrument each sensor. This work shows that by instrumenting the user and environment (which can be done with minimal expenditure), high rates of activity recognition can be

achieved across several tasks. Additionally, as the data from the tool sensors was not tagged with the tool ID, a single sensor that is transferred between the tool can be used to provide information across the tasks (e.g., users could move a single Sensor Tag between the tools to allow the system to recognize user, activity and expertise).

8.2 Limitations and Future Work

While the results from the study are promising, there are several assumptions made in the study design, and several limitations that need to be addressed when transferring these results to more ecologically valid situations.

Expertise. We attempted to recruit participants in a wide range of expertise in different tasks but the expertise across different tasks is not completely balanced, and also does not capture the entire range of expertise in the general population (e.g., we did not include very young children or veteran tradespeople). Additionally, the nature of expertise is quite subjective, as evidenced by the variance between the expert raters. While taking the average of the two expert raters reduces the subjectivity somewhat, more work is needed to determine more objective measures of expertise that can be used as a baseline.

Task complexity. While the studied tasks did have some complexity and varied stages, the tasks were still simplified and controlled. While we expect many of the results will transfer to the real world, more work is needed to ensure that tasks, activities and expertise can be identified as the environment, orientation, or materials used in the task are changed. Furthermore, we are also interested in detecting multiple activities and users simultaneously, which needs a thorough evaluation.

Long term robustness. Users would change their behaviors while gaining more experience and expertise, which may affect the user identification. However, with consistent data collection, we may be able to model these changes over time but a thorough long-term study is necessary for future work.

Machine Learning. This work did not explore different machine learning models and all possible features that could be derived from those sensors data. Across all three of these aspects, future work could explore how to improve the results by adding and changing the various aspects used to process and classify the data.

9 CONCLUSION

The addition of contextual data such as user identification, activity recognition, and expertise inference have the potential to improve a users' experience with interactive fabrication tasks. However, to achieve these visions of a responsive workspace, the system needs to have a better understanding of the context of the space. This paper shows that through instrumentation of the environment, tools, and users, a system can recognize the activities and users with a high degree of accuracy and can infer the expertise of the user at a similar level as human raters.

ACKNOWLEDGMENTS

We thank Chris Chekan for providing expert feedback and video tutorials and Michelle Annett for reviewing and rating participant performance.

REFERENCES

- [1] Gregory D. Abowd and Elizabeth D. Mynatt. 2000. Charting past, present, and future research in ubiquitous computing. *ACM Transactions on Computer-Human Interaction (TOCHI)*. 7, 1: 29–58.
DOI=<https://doi.org/10.1145/344949.344988>
- [2] Amin Ahmadi, Edmond Mitchell, Francois Destelle, et al. 2014. Automatic activity classification and movement assessment during a sports training session using wearable inertial sensors. *Wearable and Implantable Body Sensor Networks (BSN)*, 2014 11th International Conference on, IEEE, 98–103.
DOI=<https://doi.org/10.1109/BSN.2014.29>
- [3] Amin Ahmadi, David Rowlands, and Daniel Arthur James. 2010. Towards a wearable device for skill assessment and skill acquisition of a tennis player during the first serve. *Sports Technology* 2, 3–4: 129–136.
DOI=<http://doi.org/10.1002/jst.112>
- [4] Dimitrios S. Alexiadis and Petros Daras. 2014. Quaternionic Signal Processing Techniques for Automatic Evaluation of Dance Performances From MoCap Data. *IEEE Trans. Multimedia* 16, 5: 1391–1406.
DOI=<http://doi.org/10.1109/TMM.2014.2317311>
- [5] Dimitrios S. Alexiadis, Philip Kelly, Petros Daras, Noel E. O'Connor, Tamy Boubekeur, and Maher Ben Moussa. 2011. Evaluating a dancer's performance using kinect-based skeleton tracking. In *Proceedings of the 19th ACM international conference on Multimedia*, 659–662.
DOI=<http://dx.doi.org/10.1145/2072298.2072412>
- [6] Meryl Alper. 2013. Making space in the makerspace: Building a mixed-ability maker culture. *Humanities, Arts, Science, and Technology Alliance and Collaboratory*.
- [7] Fraser Anderson, Daniel W. Birch, Pierre Boulanger, and Walter F. Bischof. 2012. Sensor fusion for laparoscopic surgery skill acquisition. *Computer Aided Surgery* 17, 6: 269–283.
DOI=<http://dx.doi.org/10.3109/10929088.2012.727641>
- [8] Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2013. YouMove: Enhancing Movement Training with an Augmented Reality Mirror. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*, 311–320.
DOI=<http://doi.org/10.1145/2501988.2502045>

- [9] Nicola C. Anderson, Fraser Anderson, Alan Kingstone, and Walter F. Bischof. 2015. A comparison of scanpath comparison methods. *Behavior research methods* 47, 4: 1377–1392. DOI=<https://doi.org/10.3758/s13428-014-0550-3>
- [10] Michelle Annett, Tovi Grossman, Daniel Wigdor, George Fitzmaurice. 2018. Workshops for the 21st Century: Exploring the Role of Environments in Personal Fabrication Processes. *ACM Transactions Computer-Human Interaction* (TOCHI).
- [11] Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. 2007. A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing* 2, 4: 263–277. DOI=<http://dx.doi.org/10.1504/IJAHUC.2007.014070>
- [12] Patrick Baudisch and Stefanie Mueller. 2017. Personal fabrication. *Foundations and Trends® in Human-Computer Interaction* 10, 3–4: 165–293.
- [13] Michael Boyle and Saul Greenberg. 2005. The language of privacy: Learning from video media space analysis and design. *ACM Transactions on Computer-Human Interaction* (TOCHI) 12, 2: 328–370.
- [14] Tim Campbell, Jonathan Harper, Bjorn Hartmann, and Eric Paulos. Towards Digital Apprenticeship: Wearable Activity Recognition in the Workshop Setting. 7.
- [15] Liwei Chan, Yi-Ling Chen, Chi-Hao Hsieh, Rong-Hao Liang, and Bing-Yu Chen. 2015. Cyclopsring: Enabling whole-hand and context-aware interactions through a fisheye ring. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST'15), 549–556. DOI=<https://doi.org/10.1145/2807442.2807450>
- [16] Guanling Chen and David Kotz. 2000. A survey of context-aware mobile computing research. Technical Report TR2000-381, Dept. of Computer Science, Dartmouth College.
- [17] J. Chen, M. Yeasin, and R. Sharma. 2003. Visual modelling and evaluation of surgical skill. *Pattern Analysis & Applications* 6, 1: 1–11. DOI=<http://doi.org/10.1007/s10044-002-0165-7>
- [18] Anind K. Dey. 2001. Understanding and Using Context. *Personal Ubiquitous Comput.* 5, 1: 4–7. DOI=<http://doi.org/10.1007/s007790170019>
- [19] Anind K. Dey, Gregory D. Abowd, and Daniel Salber. 2001. A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-aware Applications. *Hum.-Comput. Interact.* 16, 2: 97–166. DOI=http://doi.org/10.1207/S15327051HCI16234_02
- [20] Yu Enokibori and Kenji Mase. 2012. A method to evaluate metal filing skill level with wearable hybrid sensor. In *Proceedings of the 3rd Augmented Human International Conference* (AH'12), 1–8. DOI=<http://doi.org/10.1145/2160125.2160131>
- [21] M. S. Erden and T. Tomiyama. 2009. Identifying welding skills for training and assistance with robot. *Science and Technology of Welding and Joining* 14, 6: 523–532. DOI=<http://doi.org/10.1179/136217109X437150>
- [22] Miikka Ermes, Juha Pärkkä, Jani Mäntyjärvi, and Ilkka Korhonen. 2008. Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions. *IEEE transactions on information technology in biomedicine* 12, 1: 20–26. DOI=<http://doi.org/10.1109/TITB.2007.899496>
- [23] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2009. OpenEAR—introducing the Munich open-source emotion and affect recognition toolkit. In *3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*, 1–6. DOI=<http://doi.org/10.1109/ACII.2009.5349350>
- [24] Jerry Fails and Dan Olsen. 2003. A design tool for camera-based interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (CHI'03), 449–456. DOI=<http://dx.doi.org/10.1145/642611.642690>
- [25] Davrondzhon Gafurov, Einar Snekkenes, and Patrick Bours. 2007. Gait authentication and identification using wearable accelerometer sensor. *IEEE Workshop on Automatic Identification Advanced Technologies*, 220–225. DOI=<https://doi.org/10.1109/AUTOID.2007.380623>
- [26] Krzysztof Z. Gajos, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the working conference on Advanced visual interfaces* (AVI'06), 201–208. DOI=<https://doi.org/10.1145/1133265.1133306>
- [27] Andy Harter, Andy Hopper, Pete Steggles, Andy Ward, and Paul Webster. 2002. The anatomy of a context-aware application. *Wireless Networks* 8, 2/3: 187–197. DOI=<https://doi.org/10.1023/A:1013767926256>
- [28] James E. Hoffman and Baskaran Subramaniam. 1995. The role of visual attention in saccadic eye movements. *Perception & psychophysics* 57, 6: 787–795. DOI=<https://doi.org/10.3758/BF03206794>
- [29] F. C. Huang, C. M. Pugh, J. L. Patton, and F. A. Mussa-Ivaldi. 2010. Learning kinematic mappings in laparoscopic surgery. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (EMBC'10), 2097–2102. DOI=<https://doi.org/10.1109/IEMBS.2010.5626188>
- [30] Scott E. Hudson and Ian Smith. 1996. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (CSCW'96), 248–257. DOI=<http://dx.doi.org/10.1145/240080.240295>
- [31] Jarrod Knibbe, Tovi Grossman, and George Fitzmaurice. 2015. Smart Makerspace: An Immersive Instructional Space for Physical Tasks. In *Proceedings of the 2015 International Conference on Interactive Tabletops and Surfaces* (ITS '15), 83–92. DOI=<http://doi.org/10.1145/2817721.2817741>
- [32] Pat Langley. 1999. User modeling in adaptive interface. In *UM99 User Modeling*. Springer, 357–370.
- [33] Gierad Laput, Robert Xiao, and Chris Harrison. 2016. Viband: High-fidelity bio-acoustic sensing using commodity smartwatch accelerometers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST'16), 321–333. DOI=<https://doi.org/10.1145/2984511.2984582>
- [34] Gierad Laput, Chouchang Yang, Robert Xiao, Alanson Sample, and Chris Harrison. 2015. Em-sense: Touch recognition of uninstrumented, electrical and electromechanical objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST'15), 157–166. DOI=<https://doi.org/10.1145/2807442.2807481>
- [35] Gierad Laput, Yang Zhang, and Chris Harrison. 2017. Synthetic Sensors: Towards General-Purpose Sensing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (CHI'17), 3986–3999. DOI=<http://doi.org/10.1145/3025453.3025773>
- [36] Matthew L. Lee and Anind K. Dey. 2015. Sensor-based observations of daily living for aging in place. *Personal and Ubiquitous Computing* 19, 1: 27–43. DOI=<http://doi.org/10.1007/s00779-014-0810-3>
- [37] Jaime Lien, Nicholas Gillian, M. Emre Karagozler, et al. 2016. Soli: Ubiquitous gesture sensing with millimeter wave radar. *ACM Transactions on Graphics* (TOG) 35, 4: 142. DOI=<https://doi.org/10.1145/2897824.2925953>
- [38] Henry C. Lin, Izhak Shafran, David Yuh, and Gregory D. Hager. 2006. Towards automatic skill evaluation: Detection and segmentation of robot-assisted surgical motions. *Computer Aided Surgery* 11, 5: 220–230. DOI=<https://doi.org/10.3109/10929080600989189>
- [39] Zhuohua Lin, Munenori Uemura, Massimiliano Zecca, et al. 2013. Objective skill evaluation for laparoscopic training based on motion analysis. *IEEE Transactions on Biomedical Engineering* 60, 4: 977–985. DOI=<https://doi.org/10.1109/TBME.2012.2230260>

- [40] Jiming Liu, Chi Kuen Wong, and Ka Keung Hui. 2003. An adaptive user interface based on personalized learning. *IEEE Intelligent Systems* 18, 2: 52–57.
DOI=<http://dx.doi.org/10.1109/MIS.2003.1193657>
- [41] Paul Lukowicz, Jamie A Ward, Holger Junker, et al. 2004. Recognizing Workshop Activity Using Body Worn Microphones and Accelerometers. In *Pervasive Computing*, 18–32.
DOI=http://doi.org/10.1007/978-3-540-24646-6_2
- [42] John D. Mason, James Ansell, Neil Warren, and Jared Torkington. 2013. Is motion analysis a valid tool for assessing laparoscopic skill? *Surgical endoscopy* 27, 5: 1468–1477.
DOI=<https://doi.org/10.1007/s00464-012-2631-7>
- [43] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2011. Ambient help. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'11)*, 2751–2760.
DOI=<https://doi.org/10.1145/1978942.1979349>
- [44] Dan Maynes-Aminzade, Terry Winograd, and Takeo Igarashi. 2007. Eyepatch: prototyping camera-based interaction through examples. In *Proceedings of the 20th annual ACM symposium on User interface software and technology (UIST'07)*, 33–42.
DOI=<https://doi.org/10.1145/1294211.1294219>
- [45] David A. Mellis, Leah Buechley, Mitchel Resnick, and Björn Hartmann. 2016. Engaging Amateurs in the Design, Fabrication, and Assembly of Electronic Devices. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*.
DOI=<http://doi.org/10.1145/2901790.2901833>
- [46] Emmanuel Munguia Tapia, Stephen S. Intille, and Kent Larson. 2007. Portable Wireless Sensors for Object Usage Sensing in the Home: Challenges and Practicalities. In *Proceedings of the 2007 European conference on Ambient intelligence (Aml'07)*, 19–37.
DOI=http://doi.org/10.1007/978-3-540-76652-0_2
- [47] Y. Munz, B. D. Kumar, K. Moorthy, S. Bann, and A. Darzi. 2004. Laparoscopic virtual reality and box trainers: is one superior to the other? *Surgical Endoscopy and Other Interventional Techniques* 18, 3: 485–494.
DOI=<https://doi.org/10.1007/s00464-003-9043-7>
- [48] Brian O'Connell. 2015. Going from Curious to Maker: New User Experiences in a University Makerspace. VentureWell. Proceedings of Open, the Annual Conference, National Collegiate Inventors & Innovators Alliance, 1.
- [49] H. Pirsiavash and D. Ramanan. 2012. Detecting activities of daily living in first-person camera views. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2847–2854.
DOI=<http://doi.org/10.1109/CVPR.2012.6248010>
- [50] Nalini K. Ratha, Jonathan H. Connell, and Ruud M. Bolle. 2001. Enhancing security and privacy in biometrics-based authentication systems. *IBM systems Journal* 40, 3: 614–634.
DOI=<http://doi.org/10.1147/sj.403.0614>
- [51] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L. Littman. 2005. Activity recognition from accelerometer data. In *Proceedings of the 17th conference on Innovative applications of artificial intelligence (IAAI'05)*, 1541–1546.
- [52] Carol E. Reiley, Henry C. Lin, David D. Yuh, and Gregory D. Hager. 2011. Review of methods for objective surgical skill evaluation. *Surgical Endoscopy* 25, 2: 356–366.
DOI=<http://doi.org/10.1007/s00464-010-1190-z>
- [53] Neil Rittenhouse, Bharat Sharma, Ranil Sonnadara, Alex Mihailidis, and Teodor Grantcharov. 2014. Design and validation of an assessment tool for open surgical procedures. *Surgical endoscopy* 28, 3: 918–924.
DOI=<https://doi.org/10.1007/s00464-013-3247-2>
- [54] Jacob Rosen and Blake Hannaford. 2000. Objective Evaluation Of Laparoscopic Surgical Skills Using Hidden Markov Models Based On Haptic Information And Tool/tissue Interactions 13.
- [55] Daniel Salber, Anind K. Dey, and Gregory D. Abowd. 1999. The Context Toolkit: Aiding the Development of Context-enabled Applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, 434–441.
DOI=<http://doi.org/10.1145/302979.303126>
- [56] Munechiko Sato, Rohan S. Puri, Alex Olwal, et al. 2017. Zensei: Embedded, multi-electrode bioimpedance sensing for implicit, ubiquitous user recognition. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI'17)*.
DOI=<https://doi.org/10.1145/3025453.3025536>
- [57] Eldon Schoop, Michelle Nguyen, Daniel Lim, Valkyrie Savage, Sean Follmer, and Björn Hartmann. 2016. Drill Sergeant: Supporting Physical Construction Projects through an Ecosystem of Augmented Tools. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA'16)*, 1607–1614.
DOI=<http://doi.org/10.1145/2851581.2892429>
- [58] Makoto Segawa, Masaya Okada, Kazumi Renge, Masahiro Tada, Haruo Noma, and Akira Utsumi. 2014. Elderly driver retraining using automatic evaluation system of safe driving skill. *IET Intelligent Transport Systems* 8, 3: 266–272.
DOI=<http://doi.org/10.1049/iet-its.2013.0027>
- [59] Weidong Shi, Jun Yang, Yifei Jiang, Feng Yang, and Yingen Xiong. 2011. Senguard: Passive user identification on smartphones using multiple sensors. *IEEE 7th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob'11)*, 141–148.
DOI=<http://doi.org/10.1109/WiMOB.2011.6085412>
- [60] Dimitrios Stefanidis, James R. Korndorffer Jr, B. Todd Heniford, and Daniel J. Scott. 2007. Limited feedback and video tutorials optimize learning and resource utilization during laparoscopic simulator training. *Surgery* 142, 2: 202–206.
DOI=<http://doi.org/10.1016/j.surg.2007.03.009>
- [61] Thomas Stiefmeier, Daniel Roggen, Georg Ogris, Paul Lukowicz, and Gerhard Tröster. 2008. Wearable activity tracking in car manufacturing. *IEEE Pervasive Computing* 7, 2.
DOI=<http://doi.org/10.1109/MPRV.2008.40>
- [62] Jeanett Strandbygaard, Flemming Bjerrum, Mathilde Maagaard, et al. 2013. Instructor feedback versus no instructor feedback on performance in a laparoscopic virtual reality simulator: a randomized trial. *Annals of surgery* 257, 5: 839–844.
DOI=<http://doi.org/10.1186/1472-6920-12-7>
- [63] Yasutaka Tashiro, Hiromasa Miura, Yoshitaka Nakanishi, Ken Okazaki, and Yukihide Iwamoto. 2009. Evaluation of skills in arthroscopic training based on trajectory and force data. *Clinical orthopaedics and related research* 467, 2: 546–552.
DOI=<http://doi.org/10.1007/s11999-008-0497-8>
- [64] Hoang Minh Thang, Vo Quang Viet, Nguyen Dinh Thuc, and Deokjai Choi. 2012. Gait identification using accelerometer on mobile phone. *IEEE International Conference on Control, Automation and Information Sciences (ICCAIS'12)*, 344–348.
DOI=<http://doi.org/10.1109/ICCAIS.2012.6466615>
- [65] Massimo Tistarelli and Enrico Grosso. 2000. Active vision-based face authentication. *Image and Vision Computing*.
DOI=[http://doi.org/10.1016/S0262-8856\(99\)00059-1](http://doi.org/10.1016/S0262-8856(99)00059-1)
- [66] Jamie A Ward and Gerhard Tro ster. 2006. Activity Recognition of Assembly Tasks Using Body-Worn Microphones and Accelerometers. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 10: 15.
DOI=<http://doi.org/10.1109/TPAMI.2006.197>
- [67] M. Weiser. 1995. The computer for the 21st century. *Scientific American* 272, 3: 78–89.
- [68] Simplelink SensorTag - TI.com. Retrieved April 11, 2016 from http://www.ti.com/ww/en/wireless_connectivity/sensortag2015/?INTC=SensorTag&HQS=sensortag
- [69] Pupil Labs - Pupil. Retrieved September 21, 2018 from <https://pupil-labs.com/pupil/>
- [70] Watch. Apple. Retrieved September 21, 2018 from <https://www.apple.com/watch/>
- [71] PJRC Store. Retrieved September 21, 2018 from <https://www.pjrc.com/store/teensy32.html>