# SmokeMon: Unobtrusive Extraction of Smoking Topography Using Wearable Energy-Efficient Thermal

ANONYMOUS AUTHOR(S)

Smoking is the leading cause of preventable death worldwide. Cigarette smoke includes thousands of chemicals that are harmful and cause tobacco-related diseases. To date, the causality between human exposure to specific compounds and the harmful effects is unknown. A first step in closing the gap in knowledge has been measuring smoking topography, or how the smoker smokes the cigarette (puffs, puff volume, and duration). However, current gold standard approaches to smoking topography involve expensive, bulky, and obtrusive sensor devices, creating unnatural smoking behavior, and preventing their potential for real-time interventions in the wild. Although motion-based wearable sensors and their corresponding machine-learned models have shown promise in unobtrusively tracking smoking gestures, they are notorious for confounding smoking with other similar hand-to-mouth gestures such as eating and drinking. In this paper, we present SmokeMon, a chest-worn thermal-sensing wearable system that can capture spatial, temporal, and thermal information around the wearer and cigarette all-day to unobtrusively and passively detect smoking events and estimates of topography. We also developed a deep learning-based framework to extract puffs and smoking topography. We evaluate SmokeMon in both controlled and free-living experiments with a total of 19 participants, more than 110 hours of data, and 115 smoking sessions achieving an F1-score of 0.9 for puff detection in lab and 0.8 in wild. By providing SmokeMon as an open-platform, we enable measuring smoking topography in free-living settings to enable testing of smoking topography in the real world, with potential to facilitate timely smoking cessation interventions.

CCS Concepts: • **Human-centered computing** → *Ubiquitous and mobile devices*; *Ubiquitous and mobile computing design and evaluation methods*; • **Applied computing** → *Law, social and behavioral sciences.*

Additional Key Words and Phrases: Wearable computing

**ACM Reference Format:**
Anonymous Author(s). 2022. SmokeMon: Unobtrusive Extraction of Smoking Topography Using Wearable Energy-Efficient Thermal. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 1 (August 2022), 24 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

## 1 INTRODUCTION

Smoking is the leading cause of preventable death worldwide [1]. Globally, over 8 million deaths are attributed to smoking each year [3]. In the United States, smoking related illness is responsible for roughly one in five deaths annually [1]. The staggering toll of smoking on human life has spurred large-scale public health efforts to raise awareness of the health risks of smoking and reduce smoking prevalence, including the US Department of Health and Human Services' 'Health People 2020' initiative[1]. While smoking is known to be harmful, the majority of what is known regarding the relationship between human exposure and the harmful effects is predominantly based on self-report, controlled lab environments, environmental sensors, and blood, saliva, or urine tests. These

[1] https://www.healthypeople.gov/2020/


2474-9567/2022/8-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

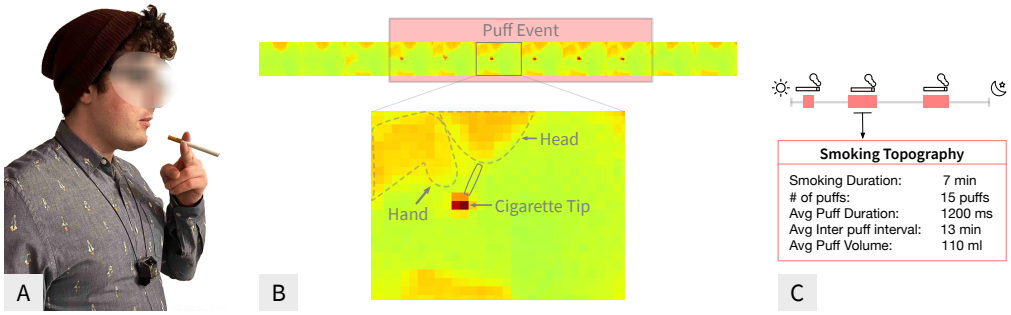


Fig. 1. SmokeMon is a chest-worn thermal-sensing wearable system with the lens pointed at the wearer [A] that can capture spatial, temporal, and thermal information around the wearer and cigarette [B] all-day to unobtrusively and passively detect smoking events and estimates of topography [C].

methods are either unreliable, expensive, or do not capture active smoking behaviors. Tobacco smoke exposure still lacks an ideal method of measurement in natural settings [15].

Recent nicotine and tobacco research studies demonstrate the importance of measuring smoking exposure through *smoking topography*, or fine-grained measures of smoking patterns. Benefits of studying smoking topography include: (1) measurement and assessment of harmful carbon monoxide exposure among smokers [12], (2) understanding the relationship between chemical exposure and tobacco related diseases [2], and even (3) providing measures to predict abstinence post treatment [44]. Smoking topography is measured during a smoking bout, which provides a summary of fine-grained measures of exposure, including timing of a puff, number of puffs, puff duration, puff volume, inter puff interval, and smoking duration. Collectively, these topography measures provide a valid and reliable index of conventional smoking exposure [25].

Gold standard methods in measuring smoking topography require smokers to attach the cigarette to one end of the device and smoke through a mouthpiece on the other end (e.g., CReSSPocket). The reduced realism of these smoking events alters the way participants smoke, impeding the ability of these devices to capture natural habituated smoking behavior [11]. To address this, researchers have investigated non-obtrusive ways to measure smoking behavior, including the use of Inertial Measurement Unit sensors (IMUs) in smartwatches [39]. However, such approaches are often confounded by non-smoking hand to mouth gestures and consequently return many false positives, suggesting motion-based sensing is inadequate for detecting smoking events and topography accurately when used alone. Wearable cameras, on the other hand, can capture temporal and spatial color information with which smoking can be distinguished from its confounding gestures [19]. However, using video data introduces privacy and stigma concerns, limiting the applicability of camera-based approaches in natural settings [9].

In this paper, we introduce SmokeMon (Fig. 1), an all day wearable device that utilizes low-resolution thermal imaging to capture temporal, spatial, and thermal features of wearer smoking activity, including puffing behavior. The *key observation* used in designing SmokeMon, is that smoking episodes have visibly distinctive thermal signatures (i.e., the temperature of the tip of a lit cigarette), as well as spatially and temporally distinctive features of hand to mouth smoking gestures. Previous work utilized multiple low-thermal sensors to classify human activity showing the promise of thermal in classifying smoking activity in a controlled environment [8]. SmokeMon extends current smoking behavior research by using low-resolution not only to classify smoking sessions but to extract further information about the smoking activity (i.e., smoking topography) using machine learning in both controlled and natural settings.

The following are the **key contributions of this paper**:

(1) We present the design and implementation of an all-day wearable system, SmokeMon, along with the data processing pipeline that enables detection of smoking activity and topography automatically and unobtrusively — that is, without reliance on user input and without physically interfering with the cigarette or smoking gesture.
(2) We evaluate SmokeMon approach to detecting smoking events and extracting smoking topography in both controlled and natural settings.
(3) We introduce and share the first dataset where smoking topography estimation using thermal imaging is attempted. We release two annotated datasets (in-lab and in-wild) to the public, enabling others to expand on the detection of smoking and estimation of smoking topography using thermal image sensing.
(4) We release and open-source the SmokeMon platform to enable further research that aims to understand or measure smoking behavior.
(5) We collect user feedback to highlight the advantages and the disadvantages of SmokeMon to inform future designs.

## 2 BACKGROUND AND RELATED WORK

### 2.1 Smoking Topography

Smoking topography provides comprehensive quantification of the manner in which someone smokes through variables described in Table 1. The ability to quantify smoking behavior at such a level of detail has allowed researchers to further understand factors that influence or maintain nicotine use. For example, smoking topography metrics can be used to predict abstinence following treatment with nicotine replacement therapy, where inter-puff interval significantly predicted abstinence at the end of the trial [44]. Due to these significant predictors, researchers are now paying closer attention to the effect of smoking cessation pharmacotherapy on smoking topography [33]. Smoking topography has also helped explain why cigarettes labeled as "light," "low," or "mild" (i.e., low-yield cigarettes) have failed to minimize the harms of smoking. Herning et al. [17] show that when people smoke low-yield cigarettes, they often take puffs with greater volume and duration than puffs taken while

Table 1. Smoking topography measurements and their definitions

| **Measurement** | **Definition** |
|---|---|
| A Puff | Starts when the person puts the cigarette in their mouth and begins inhaling; ends when the person stops inhaling the smoke. The action of exhaling the smoke is not considered part of the puff. |
| Number of puffs | Total number of puffs in a smoking session while smoking a single cigarette. |
| Puff duration | During a single puff, the time (in milliseconds) between the start and end time of the puff. |
| Puff volume | During a single puff, the volume (in milliliters) of the smoke inhaled |
| Inter puff interval (IPI) | The inter puff interval is the duration (in seconds) between the end of one puff and the start of the next |
| Smoking duration | The time (in seconds) between the start time of the first puff and the end time of the last puff within a session. |

smoking high-yield cigarettes. This is to say that smokers often compensate for the reduced nicotine content of low-yield cigarettes by changing their puffing behavior.

Devices measuring smoking topography have improved significantly throughout their relatively short history. They began as in-lab devices (e.g., Clinical Research Support System - CReSS) with which users smoke a cigarette through a mouthpiece connected to a pressure transducer to measure airflow, from which smoking topography is then calculated [30]. To facilitate smoking topography measurement in natural settings, portable devices have been introduced. CReSS pocket [4] is a portable version of CReSS. Users insert a cigarette into one end and inhale from the other end. Other smoking topography devices have been introduced as well, such as Smoking Puff Analyzer Mobile (SPA-M)[SODIM SAS, Fleury-les-Aubrais, France], CReSSmicro[Plowshare Technologies, Inc, Baltimore, MD, United States], and Wireless Personal Use Monitor (wPUM) [32]. Despite their utility, these devices are bulky and obtrusive, requiring the smoker to inhale through a mouthpiece and altering the subjective experience of smoking [30, 34]. Moreover, sharing a smoke-through device between participants requires thoroughly opening and cleaning the device between uses and constitutes a safety risk made especially salient by the ongoing Covid-19 pandemic. These challenges have inspired many researchers to propose alternative approaches to smoking topography, as we discuss in the next section.

## 2.2 Approaches to Detect Smoking Behavior Using Wearable Sensors

Tracking smoking events and topography in natural settings remains challenging. Imtiaz et al. [22] conducted a systematic review of wearables used to monitor smoking behavior in free-living environments, along with the behavioral and physiological manifestations of smoking behavior that drive choices of sensing modalities and approaches in smoking detection. For example, hand-to-mouth gestures are a behavioral manifestation of the puffing event, prompting researchers to use motion-based sensors in wrist-worn devices to track hand-to-mouth gestures as a proxy for smoking behavior. Below we summarize additional behavioral manifestations of smoking events identified by Imtiaz et al. and expand their list to include the thermal manifestation of smoking.

**Hand-to-mouth smoking gestures:** Motion-based sensors such as inertial measurement units (IMUs) have been extensively used to track hand-to-mouth gestures as a proxy to smoking behavior. Researchers have used inertial sensors (accelerometer, gyroscope, magnetometer) both individually and in combination to detect these gestures. Parate et al. demonstrate the possibility of detecting the smoking gesture using quaternion information from inertial sensors [31], while Skinner et al. [43], and Shoaib et al. [42] used data from the smartwatch's accelerometer and gyroscope to detect smoking activity. Senyurek et al. utilized only one axis from a wrist-worn accelerometer to determine the periodic repetition of smoking gestures [41]. While hand-to-mouth gestures are strong features of smoking behavior and can be used to extract puffing behavior, approaches that rely solely on hand-to-mouth gestures are often confounded by other common daily behaviors that require hand-to-mouth gestures (e.g., eating, drinking, touching the face).

**Respiratory pattern of smoking:** Researchers have shown a difference between the pattern of inhalation/exhalation when smoking and the pattern of inhalation/exhalation when breathing regularly. Specifically, Imtiaz et al. [22] review paper summarized the respiratory pattern of smoking or puffs into 4 phases: (1) a short period of breath-holding when the cigarette is in the mouth, (2) a sharp increase of airflow when smoke is inhaled, (3) a brief period to hold the smoke in the lungs, and finally (4) smoke exhalation through the nose or mouth. To measure this pattern, researchers have used wearable Respiratory Inductance Plethysmography (RIP) sensor. RIP sensors consist of inductive threads usually embedded in belts or sewn into clothing. During the inhale and exhale process, the lungs contract and expand, causing a change in the inductance of the belt[22]. While such an approach has shown promise in controlled lab experiments, in free-living contexts, it requires the use of other sensors that capture hand-to-mouth gestures to increase confidence in the detection and eliminate confounding breaths[36].

RIP signals are also known to be similar across multiple activities: standing, walking, resting idly in a chair, eating, and smoking cigarettes [22].

**Cigarette ignition event** One way to track the frequency of smoking is by tracking the number of times a cigarette is ignited. Ubi-Lighter [37] and PACT lighter [21] are examples of smart lighters that can timestamp the cigarette ignition event as a proxy to the start time of a smoking event. While such an approach does not provide further smoking topography, it is often used with other methods to reduce false-positive events. For example, Senyurek et al. [40] and Imtiaz et al. [21] used a smart lighter in combination with a wrist-based sensor to detect smoking more accurately than could be done with either device alone.

**Thermal signature of lit cigarettes:** HeatSight [8] explored using multiple low resolutions of 8×8 thermal sensors to detect human activity in a controlled environment. Classification of smoking sessions achieved a very high f-score (83.66%) which shows the promise of using thermal sensors for smoking session classification. In this paper, we extend previous work and utilize the thermal signature of lit cigarettes to classify smoking sessions and further detect and classify puffs within the smoking session to extract smoking topography. We also collect thermal data from both controlled and natural living to provide insights on the feasibility of using thermal as a smoking topography device in the wild. Finally, we provide participants feedback on their experience in wearing SmokeMon in the wild, including feedback in collecting thermal data, which can inform future studies that use low-resolution thermal sensors in the wild.

Sensing modalities or systems that can capture multiple manifestations or proxies of smoking behaviors have reported higher scores in detecting both smoking events and smoking topography. For example, PACT2 [21] employs a smart lighter in addition to IMU and RIP to further increase confidence in smoking detection by capturing cigarette lighting events. While multi-sensing approaches improved the scores of smoking detection, they impose an extra burden on users as they require charging and wearing multiple devices. Ego-centric RGB cameras have shown significant improvement and potential in measuring many behavioral manifestations of smoking [20]. However, they create privacy concerns, discomfort, or stigma for the wearer and the people around them [10]. *SmokeMon uses a low-resolution thermal camera that can sense many smoking behavior proxies such as the thermal signature of a lit cigarette, hand-to-mouth gestures, some respiratory patterns (inhalation of smoke causes the temperature tip of the cigarette to increase), and cigarette lighting events; all without capturing unnecessary information in the environment that can trigger privacy concerns (as is the case with ego-centric RGB cameras).*

### 2.3 Thermal Imaging in Human Activity Recognition and Context-Aware Computing

Thermal cameras capture infrared radiation emitted by objects in the environment. Thermal images comprise pixels representing the thermal radiation value of the object in the field of view, which can be used to calculate the temperature of the object. Humans, as well as many objects (including a lit cigarette), naturally emit thermal radiation that can be captured with the thermal sensor unobtrusively (i.e., without the need for contact). Over the years, researchers have used thermal sensing as a *non-wearable* sensing modality for in-room activity monitoring such as walking [23], or as an over-head sensor to detect hand interaction with surfaces [24], mid-air gestures [5], face-fronting to measure cognitive load [6] or stress [7], as a hand-held sensor to recognize materials [13], or perform thermographic based energy auditing [27–29]. As a *wearable* sensing modality, thermal sensors have been used primarily to extract information about the human body. Glimpse [16] used it in a wearable camera to discard frames that have bystanders by detecting humans in the scene. FingerTrak [18] used multiple thermal cameras in a bracelet configuration to predict the wearer's hand poses. HeatSight [8] has utilized multiple 8x8 thermal sensors to classify human activity in a controlled setting, including smoking events. However, they do not classify or detect puffing events to extract smoking topography. SmokeMon extends the work of HeatSight by exploring the feasibility of a low-resolution thermal sensor to extract smoking topography. Moreover, we do not

use the same thermal sensors used in HeatSight (i.e., Grid-EYE) as it fails to work in sub-zero Celsius climates, ruling out the potential of using the sensor in many parts of the world.

# 3 SMOKEMON: SYSTEM DESIGN, IMPLEMENTATION AND ENERGY CONSUMPTION

In this section, we describe SmokeMon's design goals along with explorations performed to determine the best design choices, followed by the final system designed to measure smoking topography. We conclude this section by discussing SmokeMon energy consumption and battery life.

## 3.1 Design Goals

**Goal 1: Reliable sensing of smoking behavior.** SmokeMon must provide rich information that can capture multiple proxies of smoking behaviors to be more reliable in extracting smoking topography among known confounding activities.

**Goal 2: Capture data that humans can interpret.** Data obtained from SmokeMon should be easily verifiable and interpretable by a human. One of the main challenges in the area of human activity detection is the lack of fine-grained labeled datasets because it is difficult to obtain ground truth in free-living settings. SmokeMon data should avoid this challenge by capturing data that is intelligible to the human and the machine.

**Goal 3: Unobtrusive low-burden sensing.** SmokeMon should minimally interfere with everyday activity, especially smoking. Wearable devices can introduce stigma and privacy concerns. The design, sensor position, and data collected from SmokeMon should not introduce additional burden on the wearer, affecting them physically or psychologically.

**Goal 4: Low-power sensing to support all-day wear.** Smoking can happen at any time of the day. Therefore SmokeMon should last for at least 16 hours on a single charge. This number is calculated based on waking hours, assuming that a person needs 8 hours of sleep a day. It is also essential that SmokeMon incorporates energy-efficient approaches and components to reduce the weight and size of the battery while still providing all-day sensing.

**Goal 5: Open-platform to support smoking cessation.** Smoking topography research has shown great promise in understanding smoking and nicotine intake habits and in assessing the impact of smoking cessation interventions. SmokeMon should provide an open platform to the research community to continue to advance unobtrusive capture of longitudinal information about smoking behavior.

## 3.2 Design Exploration: Motivations and Challenges in Using Thermal For Smoking Detection

To attain the aforementioned design goals, we explored multiple approaches. The thermal-based sensing approach provided the greatest promise in attaining the stated goals. We describe the advantages of using a thermal-based approach in detecting smoking behavior.

**Motivation 1: Low-resolution thermal camera for comfortable all-day sensing.** Low-resolution thermal cameras are light weight and energy efficient, making it possible to wear them continuously without burdening the participant. Because the thermal camera is worn on body, the person will be in close proximity to the camera, making it sufficient to use a low-resolution thermal camera. However, it is necessary to ensure that the field of view (FoV) of the thermal camera is capable of capturing the wearer's hand and head movement. Table 2 provides metrics and comparisons between thermal cameras available off-the-shelf. We specifically compare three different cameras that have different characteristics: (1) FLIR [2] is a high resolution thermal camera, (2) MLX [3] is a low

[2]FLIR 3.5 https://www.flir.com/products/lepton/?model=500-0771-01&vertical=lwir&segment=oem (Accessed June 5, 2022)
[3]MLX 90640 https://www.melexis.com/en/documents/documentation/datasheets/datasheet-mlx90640 (Accessed June 5, 2022)

resolution thermal camera, and (3) Grid-EYE [4] is an ultra low resolution thermal camera. As evident from the sample images presented in Table 2, each device is capable of capturing a silhouette of the human head when the camera is pointed upward from the chest, regardless of the device's resolution. However, we note that the FLIR has an extremely high current draw as compared to the other devices, which makes it unsuitable for longitudinal data collection (affecting our previously stated Goal 4). We thus eliminated the FLIR from our consideration. Initially, we were interested in using the Grid-EYE sensor (similar to the one used in HeatSight [8]) as it has a low current draw. To ensure that the Grid-EYE addressed our needs, we conducted a small experiment with the Grid-EYE device where we asked people to smoke while wearing the device. The Grid-EYE fails to work in contexts with an ambient temperature below $0°C$ and therefore restricts the use of the Grid-EYE sensor to a few months and limited regions of the world (contradict with Goal 1 and Goal 5). We also noticed that when the participants smoked sideways, the cigarette would go out of the FoV. Therefore, we opted for the MLX sensor due to its larger FoV, allowing for a more robust capture of cigarette motion under many ambient temperature settings. Additionally, the larger FoV made it less sensitive to wearing position.

One major advantage of thermal cameras is that they can capture thermal images even when they are covered with opaque materials. The ability to conceal the thermal camera increases the participant's comfort in wearing them. Researchers have shown that people grow uncomfortable when they see a camera lens, which in turn makes the wearer feel stigmatized, resulting in social avoidance, even if the camera does not record the bystander [10]. Therefore, to increase comfort in wearing SmokeMon, we investigated different materials to cover the MLX lens. We tried a thermal passing acrylic sheet and a thermal passing black thin plastic sheet (similar to a garbage bag). We chose the thin plastic sheet as it was easier to wrap around the opening of the camera.

Table 2. A comparison between thermal cameras available off the shelf.

| | **FLIR** | **MLX** | **Grid-EYE** |
|---|---|---|---|
| **Resolution** $(p)$ | $160 \times 120$ | $32 \times 24$ | $8 \times 8$ |
| **FOV** $(°)$ | 57 | $110 \times 75$ | $60 \times 60$ |
| **Temperature Range** $(°C)$ | -1 to 400 | -40 to 300 | 0 to 80 |
| **Sensitivity** $(°C)$ | 0.05 | 0.1 | 0.16 |
| **Current** $(mA)$ | 45.5 | 23 | 4.5 |
| **Per Unit Cost** $(USD)$ | 233.99 | 74.95 | 45.95 |
| **Sample Image** | 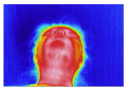 | 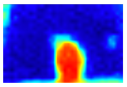 |  |

**Motivation 2: Thermal provide rich information to model smoking topography.** It is well understood that the more data we can capture from diverse smoking instances, the better SmokeMon's ability to predict and correctly model smoking behavior. Several observations confirmed that the thermal sensor provides rich smoking-related information. SmokeMon takes advantage of spatio-temporal thermal data to capture the moment when someone lights up the cigarette. Also, the temperature of the cigarette during a smoking session provides valuable information about the puff. Figure 2 shows a time-series signal of the max pixel temperature during a smoking session. We also highlight the ground truth of a puff in the same figure. One must note that thermal information alone can not predict puffs accurately. We can see this in Fig. 2[c], where the cigarette is in the FoV

[4]Grid-EYE https://industry.panasonic.eu/products/components/sensors/ir-thermophile-array-sensor-grid-eye (Accessed June 5, 2022)

but not in the mouth. We see a peak in the max pixel temperature when the cigarette is in the FoV. However, the spatio-temporal thermal information clearly shows that the cigarette is not near the mouth. Moreover, during the puff, we can see an increase in the cigarette temperature caused by the combustion process when the participant inhales air though the cigarette. This information can be used to approximate puff volume. Fig. 3 shows the correlation (Pearson's 0.77) between the normalized ground truth puff volume obtained from CReSS Pocket and the one calculated from the thermal data by integrating over the temperature rate of change during the puff.

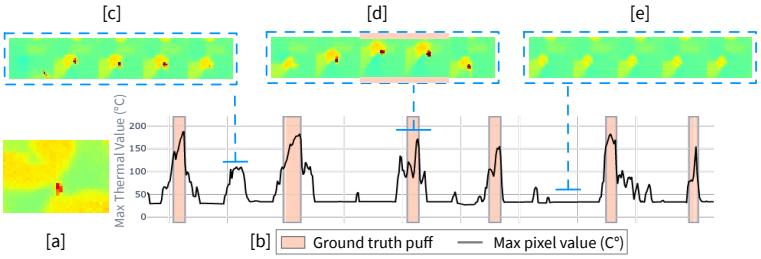


Fig. 2. Thermal captures fine-grained information about the smoking session including: [a] cigarette ignition [b] cigarette tip during a smoking session. Thermal, spatial and temporal information can help in distinguishing between a puff (when the cigarette is near the mouth [d]) and a non-puff (when the cigarette is in the field of view but not in the mouth [c] or when there is no cigarette [e])

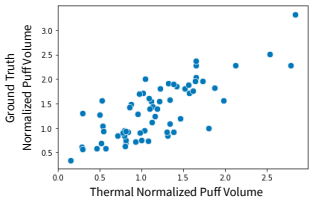


Fig. 3. the correlation (person-r 0.77) between the normalized ground truth puff volume obtained from CReSS Pocket and the one calculated from the thermal data.

**Motivation 3: Thermo-spatial data are useful in distinguishing confounding hand-to-mouth gestures.** Hand-to-mouth gestures are often confounding; research has shown how motion-based sensors like IMU often confound similar hand-to-mouth gestures like eating, drinking, and smoking [38]. Thermal images provide useful information to the human and the machine, which helps in distinguishing between similar and often confounding hand-to-mouth gestures. Fig. 4 shows examples of different hand-to-mouth gestures with and without objects in hand, such as smoking, eating, drinking, and touching the head or mouth. We use a color map that maximizes the viewability of the image to a human reviewer. From the images, we can observe that the smoking gesture has a clear signature that is easily observable by a human reviewer. As we will show in subsequent sections, a machine can also learn to capture this signature.

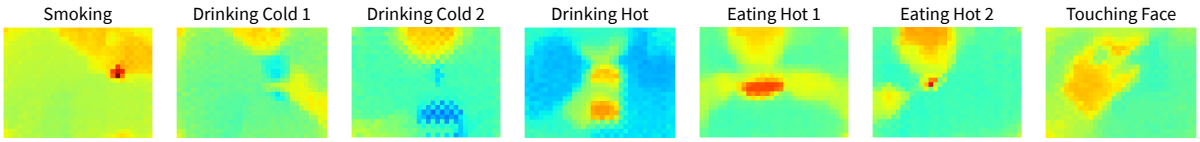


Fig. 4. Thermal capture rich information that can help in distinguish between many hand-to-mouth gestures.

## 3.3 SmokeMon Hardware, Firmware, and Encapsulation

**Hardware:** For the development of the SmokeMon prototype, we used a development board [5] equipped with a Cortex-M4-based Apollo 3 microcontroller and a BLE 5.0 radio. In addition, the board also contains an MCP73831 single cell LiPo charger, a Real Time Clock (RTC), and a micro-SD card socket, which are essential components for

[5]Sparkfun OpenLog Artemis https://www.sparkfun.com/products/16832 (Accessed August 15, 2022)

in-wild deployment. BLE and the micro-SD card slot are required for communication and data storage. SmokeMon obtains the thermal images from a breakout board that contains MLX90604 thermal sensor array [6], and connects to the main board via an I2C-based connector. The size of each captured thermal image is $32 \times 24$, where each pixel represents a temperature reading ranging from -40° C to 300° C.

**Encapsulation and attachments:** It is important to encapsulate SmokeMon in a wearable casing that maximizes comfort and allows reliable data collection. The thermal camera in SmokeMon should be positioned in a way that allows capturing of the head and as much of the hand trajectory as possible. This will allow confirmation of the cigarette going towards the mouth. Therefore, we printed multiple mounts with different angles (centering the camera towards the chin, at 45° to 60° from the chest) to hold the thermal sensor based on the wearer's chest size (see Fig. 5). For example, a person with a flat chest will have a mount with a 55° or more, while a person with a protruding chest will require a smaller angle. We also fixed the position of the device on the chest using a necklace chain of fixed length. We determine the best angle and chain length for the participants during the lab visit; we confirm the angle and length by visualizing the data obtained from SmokeMon device. To prevent SmokeMon from flipping or moving out of place, we provide a magnet attachment mechanism – participants place a magnet behind their cloth to connect with the magnets with opposing poles embedded behind SmokeMon's case.

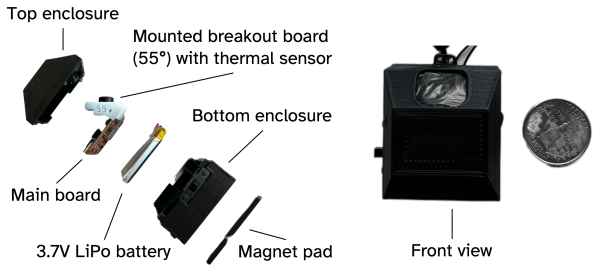


Fig. 5. Enclosure and attachment mechanisms in SmokeMon to maximize comfort while enabling reliable data collection.

**Firmware:** SmokeMon firmware samples a timestamped thermal frame at 4 Hz (equivalent to 4 fps). Depending on the goal of using SmokeMon, the firmware can be set to save all timestamp data to an external SD Card (*Storage State*) or can be set to stream data to another device using BLE (*BLE State*). The *Storage State* is useful when continuous and reliable data is favored over real-time streaming or process of the data. For example, in our case, we would rather save all collected data on an SD card to increase the reliability and validity of our experiment. Similarly, smoking topography devices (e.g., CReSS Pocket) perform analysis offline at the end of the data collection process. On the other hand, a smoking cessation application might utilize *BLE State* to send notifications to the user as a form of intervention or might need to offload data for further real-time processing on smartphones. SmokeMon facilitates both cases by enabling the researchers or users of SmokeMon to choose the firmware state based on their goals.

## 3.4 Energy Consumption and Battery Life

**Energy Consumption:** Table 3 presents the average current draw of each state that SmokeMon supports, including the *Storage State* and the *BLE State* we explained above. In addition to these two states, SmokeMon also supports the *Sensing only* state. We use the INA219 current sensor [7] to measure the power consumption of

[6] Adafruit MLX90640 https://www.adafruit.com/product/4469 (Accessed August 15, 2022)
[7] INA219 https://www.adafruit.com/product/904 (Accessed June 5, 2022)

SmokeMon in each firmware configuration: *Sensing Only*, *Storage State* and *BLE State*. To measure the current, we connected the current sensor between the power source (500 mAh LiPo battery) and SmokeMon GND and VIN contacts on the board. We then connected an Arduino board to the same current sensor via I2C to read the power profile of the system at a sampling frequency of 50 Hz. For each state, we collected the power profile data for 1 hour and calculate the median current consumption for each state, as reported in Table 3. We used these numbers to estimate the battery life of SmokeMon.

Table 3. Current draw and power consumption of SmokeMon firmware under different settings. Sensing only state collects and processes the thermal data. Storage state collects, process, and store the data in an SD Card. BLE state collects, process, and stream the thermal data over BLE.

| | **SmokeMon Firmware** | | |
|---|---|---|---|
| | *Sensing Only* | *Storage State* | *BLE State* |
| **Current Draw (mA)** | 24.2 | 26.1 | 33.6 |
| **Power (mW)** | 91.9 | 104.4 | 132.7 |

**Battery Life:** For our data collection, we used SmokeMon *Storage State* firmware. Based on the estimated current consumption for SmokeMon *Storage State* firmware shown in Table 3 and based on the fact that we want to system to last for a full waking day (at least 16 hours, assuming 8 hours of sleep), we chose to equip SmokeMon with a 500mAh battery, resulting in a battery lifetime of 19 hours. To ensure that our estimation is correct, we fully charged SmokeMon and asked a volunteer to wear the device until the battery drained out completely. We instructed the volunteer not to switch off the system but rather take it off and place it on a table at the end of the study. When the battery is fully depleted, we calculate the battery lifetime by calculating the difference between the first timestamp and the last time stamp of readings stored on the SD card.

Values in Table 3 can also be used to estimate battery lifetime when multiple states are used. For example, instead of continuously saving or streaming data, a more energy-efficient approach will only save or send the data when there is a high probability of a smoking event (i.e., when a very hot pixel is detected). In Section 6.3, we perform such analysis based on the data that we collected to demonstrate how an opportunistic data collection approach can further reduce the power consumption of the system to enable a longer battery life or to facilitate real-time processing via BLE.

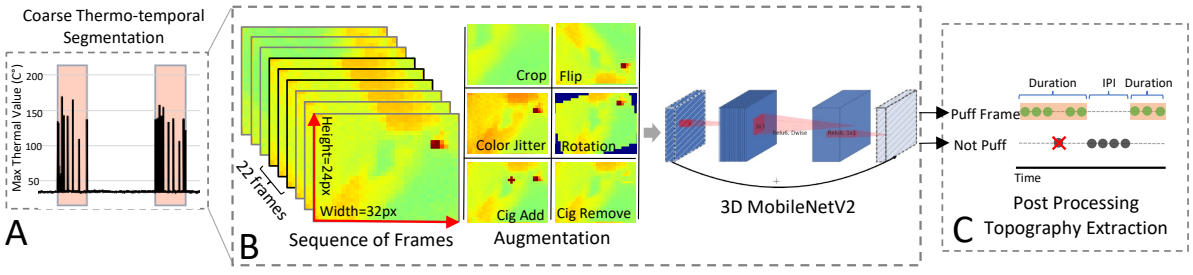


Fig. 6. SmokeMon Framework to Extract Smoking Topography: [A] we first extract possible smoking sessions based on thermal information. Then we [B] run a puff detection model on each frame in the extracted session. Finally, [C] we perform post processing to extract puffing events and the smoking topography of the session.

## 4 SMOKEMON FRAMEWORK TO EXTRACT SMOKING TOPOGRAPHY

Figure 6 illustrates the SmokeMon framework for extracting smoking topography. SmokeMon first performs a coarse thermo-temporal segmentation to automatically discard time segments with a low probability of smoking. Within each remaining time segment, SmokeMon runs a puff classification model on each frame and then constructs a puff event by post-processing the frame-level puff prediction. Finally, SmokeMon extracts the smoking topography based on the constructed puff events.

### 4.1 Step A: Coarse Thermo-Temporal Segmentation

Passive wearable systems deployed in free-living settings collect massive amounts of data about both relevant (i.e., smoking) and other irrelevant activities. Running a machine learning model on each time segment is inefficient as it unnecessarily increases computational load and energy consumption. As we are interested in smoking segments of the data, we perform a coarse thermo-temporal segmentation to discard irrelevant time segments and keep time segments with a high probability of being a smoking session (i.e., segments that have at least one pixel with a high temperature value that may belong to the tip of a burning cigarette). To reduce the runtime and energy consumption of the coarse thermo-temporal segmentation step, we designed a heuristic-based algorithm (See Algorithm 1) that can efficiently run either on-device on the SmokeMon MCU or on a offline CPU after extracting the data from the SD card. The algorithm loops over the frames and checks the maximum pixel value of each frame $F_i$ and creates a possible smoking session if the max value is $> 70^\circ$ (determined empirically based on the low bound range of the cigarette tips in our dataset). For each $F_i$ frame that is flagged as containing potential smoking activity, the algorithm also flags previous and future frames within the range of $[F_{i-start\_offset}, F_{i+end\_offset}]$ to include likely puff events without cigarette visibility and obtain a coarse thermo-temporal segment. The offset is a variable in the algorithm that can be adjusted based on the interpuff interval duration. The resulting segments include both smoking and other confounding contexts or activities that can have a high thermal pixel (e.g., hot drinks and spotlight lighting). In the next step, we use a deep learning approach to classify puff frames and further eliminate confounding and non-smoking frames.

**Algorithm 1** Coarse Thermo-Temporal Segmentation

**Parameters:** start_offset = 1320; end_offset = 2640; temperature_threshold = 70

```
historical_buffer = [ ]                          ▷ The buffer to hold the latest frames (size=start_offset)
segment_end = 0
smoking_sessions = { }
while True do
    time, data ← getFrame()
    historical_buffer.update((time,data))
    if time ≤ segment_end then                                   ▷ session segmentation continues
        smoking_sessions[time] ←data                             ▷ flag the time instance as smoking
    else if MaxTemp(data) ≥ temperature_threshold then           ▷ check if cigarette exist
        segment_end ← time + end_offset                          ▷ initiate a segment
        for time, data in historical_buffer do                   ▷ flag previous time instances as smoking
            smoking_sessions[time] ←data
```

### 4.2 Step B: Puff Frame Detection

We train a deep neural net model to determine whether a participant was performing a cigarette puff in frame $F_i$. For each frame $F_i$, frames $F_{i-11}$ to $F_{i+11}$ were stacked to create a 3D input data of size 23×32×24, which was fed

into the neural network. We use the MobileNetV2 architecture [35] as it is a light model that can run efficiently on edge devices and enable real-time inference if needed. We replaced the 2D convolution layers with their 3D counterparts to capture temporal information related to the puff gesture. Instead of randomly initializing the weights in our neural network, we utilize transfer learning and initialize our network with weights obtained from a similar model pre-trained on a gesture recognition dataset [26]

To overcome the challenge of between-subjects and within-subjects smoking gesture variability, we augmented our training set by using the following transformations (See Fig. 6 for visuals): (1) We randomly remove cigarettes in the frame to force the neural network to learn puffing gestures with and without a visible cigarette. We accomplish this by first locating the cigarette in the frame and then using an in-painting method to remove the cigarette. (2) We also randomly add cigarette confounding blobs in the image to ensure that the network learns the difference between a cigarette in hand and a random hot blob that looks like a cigarette. (3) Since temperature readings get affected by environmental factors such as cold weather, we arbitrarily applied brightness jitter augmentation. (4) We applied other common augmentation methods such as random crop, random flip, and random rotation to account for device position variability. To further increase the resilience of our models, we collect data from two participants performing intentionally confounding gestures such as drinking water of varying temperature and touching their face (one hour each person) to augment the training set.

We used binary cross entropy loss function with Adam optimizer and chose a learning rate scheduler to run gradient descent until convergence. The model's performance was measured with a cross-validation method (leave one participant out). Currently, we do not perform any on-device machine learning prediction, however we made efficient choices about the model architecture to enable future on-device processing.

### 4.3 Step C: Post-processing and Smoking Topography Extraction

After detecting the smoking frames, we cluster the puff frames based on temporal proximity (max is 12 seconds) to create a puffing event. Clusters with less than 2 positive puffing frames are discarded. After obtaining the puff events, we calculate the following smoking topography: number of puffs, puff duration, and inter-puff interval. For the in-lab data only, we also calculate the puff volume and use a smoking topography device (CReSS Pocket) for ground truth. Puff volume is calculated from thermal data by first normalizing the data using the mean and standard deviation of each participant. Then for each puff event, we extract a time series of the max value of each frame in the puff event to create a max temperature curve. The change of max temperature within a puff is highly correlated with airflow going through the cigarette and measured by CReSS Pocket. To measure volume, CReSS Pocket generates a flow curve (a time series measurement of flow generated using a pressure transducer) for each puff, and integrates over the flow curve. In SmokeMon, we integrate over the max temperature curve within the puff to estimate volume. To obtain a mapping between our estimation and ground truth, we train a linear regression model using leave-one-participant-out cross-validation to estimate the puff volume.

## 5 EXPERIMENTS AND DATASETS

To evaluate SmokeMon, we conducted an in-lab and in-wild experiment that we further explain below.

### 5.1 In-Lab Experiment

We conducted a controlled in-lab experiment to compare the performance of SmokeMon against the research-grade smoking topography measurement device (CReSS Pocket). The elements of smoking topography we are interested in for the in-lab evaluation are the number of puffs, puffs duration, puff volume and the duration of the inter puff interval.

*5.1.1 Smoking Topography Ground Truth.* To extract ground truth smoking topography, we use CReSS Pocket, a popular portable device used in clinical settings. CReSS Pocket is a battery-powered device that automatically

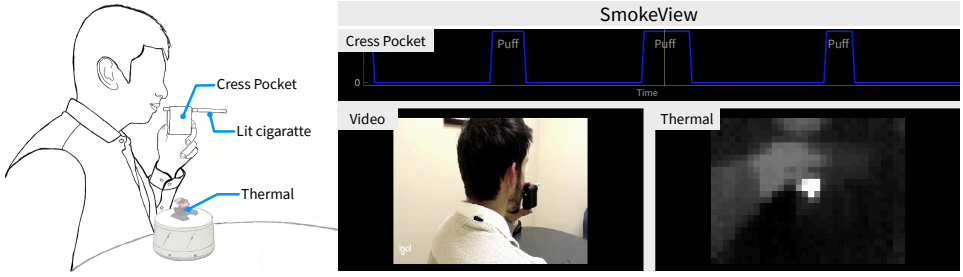


Fig. 7. In lab experimental setup.

measures smoking behavior parameters such as date, time, start and end of the smoking episode, puffs per cigarette, puff volume, and puff duration. In our evaluation, we compare the smoking topography extracted from our pipeline with the smoking topography obtained from CReSS Pocket.

*5.1.2 Protocol.* Participants were asked to smoke a cigarette using CReSS Pocket and SmokeMon together to provide smoking topography measurements from both devices for each smoking episode. To use CReSS Pocket, participants insert a lit cigarette into one end of the device and inhale through the other end using a disposable (See Fig. 7). During the pilot study, we noticed that the unnatural position of the cigarette was caused by using CReSS Pocket obstructed SmokeMon view of the cigarette. In light of this, the experiment was revised so that SmokeMon was not worn but rather placed on the table in front of the participant to compensate for the cigarette offset created by the CReSS Pocket (See Fig. 7). Because the timekeeping mechanisms of the two systems differ, all sessions were videotaped to facilitate synchronization of SmokeMon and CReSS Pocket data. After teaching participants how to smoke using CReSS Pocket, we leave the participant alone in the room and ask them to smoke a cigarette. Participants smoked one cigarette of their choice (provided by the study team) and were compensated $10 for their time. This experiment was conducted before the COVID-19 pandemic.

*5.1.3 Dataset.* We collected data from 8 participants (1 female). All participants (except one) were students at [anonymized for review] university. More details about the collected smoking topography using CReSS Pocket can be found in Table 4. To synchronize data collected from SmokeMon, we developed SmokeView to visualize the data we collected during the experiment: video, SmokeMon thermal data, and CReSS Pocket puff time series (see Fig. 7). We mark the lighting event, the first puff in the video, and the first puff observed in SmokeMon to calculate the offset between the video and SmokeMon. We also mark the first puff using CReSS Pocket to synchronize CReSS Pocket data to the video. After syncing the data, we remove the puffs that were not captured by CReSS Pocket, which were puffs used to light the cigarette. The annotated thermal data along with CReSS Pocket data will be released to the public.

## 5.2 In-Wild Experiment

The objective of the in-wild experiment is to test SmokeMon under realistic everyday conditions. We ask participants to wear the device in a natural setting for at least 10 hours (section 5.2.1). Participants were asked to keep a log of their smoking session, which we then used to obtain a fine-grained annotation of the smoking session and its puffs (section 5.2.3). Table 5 provides summary statistics of the data collecting in the wild which we discuss in 5.2.4

*5.2.1 In-Wild Protocol.* We recruited participants using ResearchMatch.org, Craigslist.com, and flyers distributed around the campus located in a metropolitan city [anonymized for review]. Participants were instructed on how

Table 4. Puffing statistics obtained from CReSS Pocket

| p | Average Puff Duration (s) | Average IPI (s) | Average Puff Volume (ml) | Smoking Duration (m) | Number of Puffs |
|---|---|---|---|---|---|
| P1 | 2.62 | 16.41 | 170.06 | 4.00 | 15 |
| P2 | 2.36 | 17.46 | 207.70 | 6.55 | 21 |
| P3 | 1.63 | 11.80 | 91.02 | 5.71 | 22 |
| P4 | 2.53 | 13.63 | 102.92 | 6.65 | 26 |
| P5 | 3.90 | 17.77 | 274.89 | 8.27 | 23 |
| P6 | 4.45 | 11.57 | 233.48 | 4.59 | 38 |
| P7 | 2.91 | 8.20 | 124.22 | 4.20 | 22 |
| P8 | 1.91 | 9.72 | 129.05 | 2.84 | 11 |

to wear and operate the device. They were asked to wear the device in the lab and smoke and drink a beverage of their choice in the lab to ensure that they understood the instruction and answer any further questions they might have. Participants also had a chance to see a sample of the thermal images collected in the lab to demonstrate the type of information that we are collecting as part of the informed consent process. We asked the participants to wear SmokeMon for at least 10 hours in the wild and requested that it include smoking and non-smoking events (i.e., eating, drinking, etc.). Participants were also asked to log the start time of their smoking event by noting it down on an electronic or paper log. All participants were compensated with 60 USD, and they were given a chance to review and delete data if requested (no one deleted any data). At the end of the experiment, we asked participants to fill out a post-experiment questionnaire.

*5.2.2 Post-Experiment Questionnaire.* At the end of the study, participants engaged in a post-experiment questionnaire. We asked them about their general impression of their experience and whether they faced any discomfort while wearing the device or if it interfered with their daily lives. We asked participants to report bystander reactions and any discomfort communicated to them by the bystanders. We also ask participants questions about their thermal data and if they have any concerns about sharing it.

**General feedback questions.** To gain a general impression of SmokeMon, we asked participants the following two open-ended questions: [Q1] *"What did you think about wearing the device to collect data about your activities?"* [Q2] *"What can we do to improve the experience of wearing the device?"*.

**Participants comfort questions.** To understand the user concerns that SmokeMon can illustrate, we asked participants the following open-ended question: [Q3] *"Were there any situations when wearing the device that made you feel uncomfortable?"*, and [Q4] *"Did you ever take off or cover the device? If yes, when and why?"*. We also asked the participants if the device intervened with everyday activity as a proxy to measure possible user burdens and concerns. We specifically asked: [Q5] *"Did wearing the device change how you normally go about your daily activities?... how you smoked? (if yes, how?)"*.

**Bystander reaction questions.** Imaging sensors often cause bystander concerns. However, it is unknown how bystanders will react to the thermal sensor in SmokeMon. To capture bystander reaction, we asked participants: [Q6] *"Did anyone ask you about the devices? If yes, what did you them?"* and [Q7] *Did anyone ask you to take the device off or turn it off/around?.*

**Attitude toward data sharing questions.** To understand potential data sharing concerns, we asked participants to review part of their data and delete any segments they did not want to share with the researchers. We then ask them [Q8] *"Have you deleted any part of the data you collected? Why?"*. We also asked them to report their

concern level if we released their data to multiple stakeholders to understand the potential privacy concerns of thermal data produced by SmokeMon. Specifically, we asked them [Q9] *"To what extent would it concern you if we released the data you just viewed to each of the following: other researchers, members of your family, your friends, your doctor, your dietitian, and finally the general public"*.

*5.2.3 In-Wild Smoking Topography Ground Truth.* We label the start and end of each puff gesture in a smoking event. We use participants' self annotation of their smoking to locate the smoking events. We also view the entire dataset in the case a participant reported that they forgot to log a smoking event. We then calculate the following topography from our ground truth labels for each smoking event: 1) start and end of smoking event, (2) puffs per cigarette and (3) puff duration, and (4) inter puff interval. Puff volume can not be extracted from the in-wild data as we do not have ground truth for it. For the in-lab experiment, we used CReSS Pocket for ground truth. However, we can not use CReSS Pocket as a ground truth device in the wild as it interferes with smoking behavior. Moreover, sharing an inhale through a device during the recent and ongoing COVID-19 pandemic might impose unnecessary risk on the participants.

*5.2.4 In-Wild Dataset.* We collected a total of 110 hours of data containing 115 smoking sessions obtained from 11 participants (different participants that in lab). Table 5 shows the number of smoking and puff events statistics per participant. The annotated dataset of participants who consented to release their data will be shared publicly.

Table 5. In-wild participants self-reported demographics and number of cigarettes smoked per day along with the In-wild participants smoking data and topography summary.

| P# | Age | Gender | Cigarette per day | Number of sessions | Average Smoking Duration (m) | Number of Puffs (s) | Average Puff Duration (s) | Average IPI (s) |
|---|---|---|---|---|---|---|---|---|
| P9 | 23 | Non-Binary | 2 | 1 | 4.22 | 12 | 1.64 | 21.25 |
| P10 | 56 | Woman | 12 | 12 | 6.28 | 198 | 4.11 | 27.88 |
| P11 | 75 | Woman | 20 | 14 | 5.11 | 185 | 2.75 | 22.09 |
| P12 | 24 | Man | 3 | 11 | 6.52 | 205 | 2.48 | 22.54 |
| P13 | 32 | Woman | 2 | 6 | 3.95 | 48 | 1.58 | 32.03 |
| P14 | 39 | Woman | 25 | 19 | 5.94 | 172 | 1.46 | 45.78 |
| P15 | 45 | Woman | 10 | 10 | 2.26 | 33 | 1.71 | 33.20 |
| P16 | 42 | Woman | 10 | 15 | 4.18 | 143 | 3.56 | 29.85 |
| P17 | 39 | Man | 3 | 11 | 3.94 | 81 | 2.57 | 39.82 |
| P18 | 33 | Woman | 13 | 15 | 7.41 | 146 | 3.43 | 54.68 |
| P19 | 24 | Man | 2 | 2 | 5.70 | 20 | 1.56 | 36.48 |

## 5.3 Evaluation Criteria and Metrics

We evaluate our smoking topography extraction approach using a leave one participant out method in which we test the models on one participant and train the models with the remaining participants. We report our performance in two tasks: puff detection and smoking topography extraction. Our model predicts puffing frames which are followed by a post-processing method that creates events and is evaluated at an event level. We report metrics such as the positive precision, recall,F1 score (positive), and intersection over union (IoU) between ground truth and predicted puff events for each participant individually. Using the true positive puff events, we report the error in our smoking topography measurements. We only use the true positive puffs to prevent carry-over errors from the detection task, which we report independently. To analyze the agreement between SmokeMon measurement and ground truth, we use Bland-Altman plots - a method often used to compare between gold standards and new instruments in medical and clinical research

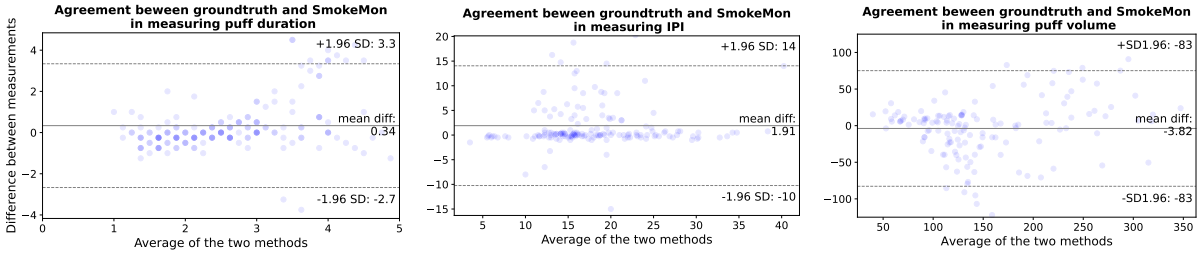


(a) (b) (c)

Fig. 8. In-lab Bland–Altman plots comparing between ground truth and SmokeMon smoking topography measurements.

Table 6. In-Lab puff detection and puff count per participant

| P# | Puff Events Detection | | | | Puff Count | |
|---|---|---|---|---|---|---|
| | Recall | Precision | F1 | mIoU | Ground truth | Predicted |
| P1 | 0.93 | 0.82 | 0.87 | 0.69 | 15 | 17 |
| P2 | 0.95 | 1.00 | 0.98 | 0.78 | 21 | 20 |
| P3 | 0.91 | 0.87 | 0.89 | 0.76 | 22 | 23 |
| P4 | 0.96 | 0.60 | 0.74 | 0.71 | 26 | 42 |
| P5 | 1.00 | 0.92 | 0.96 | 0.80 | 23 | 25 |
| P6 | 1.00 | 0.83 | 0.90 | 0.55 | 38 | 46 |
| P7 | 1.00 | 0.96 | 0.98 | 0.79 | 22 | 23 |
| P8 | 0.82 | 0.90 | 0.86 | 0.72 | 11 | 10 |
| **mean** | 0.95 | 0.86 | 0.90 | 0.73 | - | - |

# 6 RESULTS

## 6.1 In lab: Puff Detection and Smoking Topography

Table 6 shows the result of puff event detection from our in-lab experiment using SmokeMon. The average positive F1 score for puff detection is high (0.9). The per-person F1 score is also high (above 0.85) for most of the participants. We inspected low puff detection results to understand the source of the error and the limits of our current model. P4 recall is very low compared to other participants. Upon inspecting P4's false positives, it appears that P4 keeps the cigarette near the face outside of puff events, and our model predicted this action as a puff. Augmenting the training dataset with such cases can further improve the model prediction. Overall, SmokeMon achieves 0.95 recall, 0.86 precision, and 0.90 F1 score in the in-lab experiment with more than 50% mean intersection over union.

Fig. 8 shows the Bland–Altman plots that compare between smoking topography obtained from ground truth and the ones obtained from SmokeMon. The mean difference is minimal in all of the plots (0.34 seconds for puff duration, 1.91 seconds for IPI, and -3.82 for puff volume). For puff duration and IPI, most points are close to the zero line (the closer the points to the zero-line, the more agreements between ground truth and SmokeMon measurements). The puff volume plot shows greater variability, particularly as the volume average increases, indicating a systematic bias in our measurement or the calibration process of our ground truth device CReSS Pocket. However, most points are between the ±1.96 standard deviation (SD), which is acceptable. We investigated the outlier points (i.e., located outside of the ±1.96 SD limit) for each plot and found that all puffs which lie above

the +1.96 SD line belong to P6 with long duration puffs (mean = 6.1 sec). These long duration puffs occur when the participant keeps the cigarette in their mouth while taking multiple puffs.

## 6.2 In Wild: Puff Detection and Smoking Topography

Table 7 highlights the results of puff event and smoking session detection from the in-wild experiment using SmokeMon. We extracted all the smoking sessions (recall 1 as shown in Table 7). The model average F1 score is 0.8 across all participants and above 0.75 for most of the participants. We investigated P15's false positive and false negative cases as the recall and precision were very low. It appears that the position of the device for P15 was quite different than other participants (i.e., head not fully visible in the FoV), which accounts for low recall. Also, P15 kept the cigarette near their mouth and visible in the FoV between puffs, which confused the model with false positives. Overall, SmokeMon achieves 0.82 recall, 0.78 precision, and 0.8 F1 scores for in-wild experiment with more than %50 mean intersection over the union.

Fig. 9 shows the in-wild Bland–Altman plots that compare smoking topography obtained from ground truth with the ones obtained from SmokeMon. The mean difference is minimal in all of the plots (1.92 seconds for puff duration and -3.33 seconds for IPI). In both plots, we can see that most points cluster around the zero-line, showing high similarity between the ground truth measure and SmokeMon. The puff duration and the IPI error increase and fluctuate more frequently as the average increases. However, similar to the in-lab results, most points are between the ±1.96 standard deviation (SD), which is acceptable. We investigate the outlier points (i.e., located outside of the ±1.96 SD limit) for each plot and found that all puffs that lie above the +1.96 SD line belong to a few sessions from P10 with long duration puffs (mean = 5.4 sec).

Table 7. In-Wild average puff prediction per participant, puff count and smoking session detection. P10 smoking session detection analysis is removed because the ratio of smoking events to nonsmoking events is not similar to other participants as the participants not following the study protocol and turned the device on during smoking sessions only.

| **P#** | **Puff Events Detection** | | | | **Puff Count** | | **Smoking Session Detection** | | |
|---|---|---|---|---|---|---|---|---|---|
| | Recall | Precision | F1 | mIoU | Ground truth | Predicted | Recall | Precision | F1 |
| P9 | 0.75 | 0.82 | 0.78 | 0.28 | 12 | 11 | 1.0 | 1.00 | 1.00 |
| P10 | 0.94 | 1.00 | 0.97 | 0.57 | 198 | 194 | - | - | - |
| P11 | 0.88 | 0.75 | 0.81 | 0.70 | 185 | 282 | 1.0 | 0.93 | 0.97 |
| P12 | 0.84 | 0.83 | 0.84 | 0.65 | 205 | 204 | 1.0 | 1.00 | 1.00 |
| P13 | 0.79 | 0.63 | 0.70 | 0.70 | 48 | 54 | 1.0 | 1.00 | 1.00 |
| P14 | 0.86 | 0.83 | 0.85 | 0.72 | 172 | 158 | 1.0 | 1.00 | 1.00 |
| P15 | 0.67 | 0.62 | 0.65 | 0.56 | 33 | 48 | 1.0 | 0.62 | 0.77 |
| P16 | 0.75 | 0.75 | 0.75 | 0.44 | 143 | 178 | 1.0 | 1.00 | 1.00 |
| P17 | 0.97 | 0.77 | 0.86 | 0.59 | 81 | 101 | 1.0 | 0.92 | 0.96 |
| P18 | 0.78 | 0.78 | 0.78 | 0.63 | 146 | 135 | 1.0 | 1.00 | 1.00 |
| P19 | 0.78 | 0.78 | 0.78 | 0.65 | 20 | 21 | 1.0 | 0.67 | 0.80 |
| **mean** | 0.82 | 0.78 | 0.80 | 0.59 | - | - | 1.00 | 0.92 | 0.95 |

## 6.3 Energy Consumption Analysis

In section 3.4, we profiled each state supported by the SmokeMon firmware, including sensing only state, storage state, and BLE state. For our data collection, we set the system to operate in the storage state to increase our experiment's reliability and validity. Using the storage state continuously still produced an energy-efficient system as we achieved 19 hours of battery life. In this section, we further analyze SmokeMon energy consumption

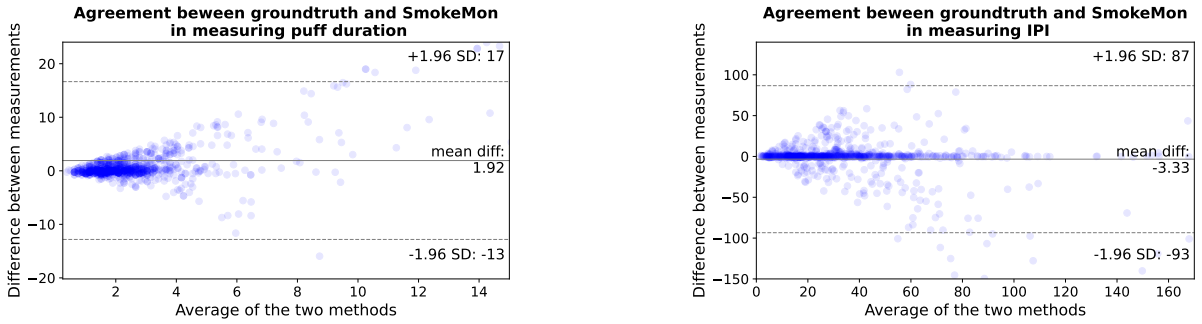


(a) (b)

Fig. 9. In-wild Bland–Altman plots comparing between ground truth and SmokeMon smoking topography measurements.

using both storage and BLE states under an opportunistic rather than a continuous data collection method. In section 4.1, we explained the coarse thermo-temporal segmentation we used to filter out irrelevant and highly unlikely smoking frames. We use the same algorithm to switch the state of SmokeMon firmware from sensing only to storage or BLE. If the frame is unlikely to be smoking, we do not activate the storage or the BLE state and just use the sensing only state. We activate the storage or the BLE state only when the algorithm detects a smoking segment. We use the in-wild data collected to simulate a realistic scenario of state change triggered by our algorithm. Then use the state power consumption that we estimated for each state earlier, shown in table 3.

Table 8 shows the percentage of discarded frames from each participant. On average, we achieved a reduction of 71% frames. Using the number of discarded frames we calculate the added battery life (in hours) to SmokeMon when BLE and Storage State are activated opportunistically using our algorithm. On average, when our algorithm is used to activate the BLE state opportunistically, we increased the battery life by 28% (4.1 hours) while streaming the data using BLE. Similarly, we improved the battery life by 6% (1.07 hours) on average when the storage state is activated using our algorithm. For instance, our algorithm identified 3% of the collected data as smoking-related to P9, which means that we only stream (BLE state) or store (storage state) this portion of the data. As a result, the system is in the sensing only state for most of the time (97%), which is the most power-efficient state.

### 6.4 SmokeMon Feedback From In-wild Participants

Below is a summary of the participants' response to the post-experiment questionnaire mentioned earlier in section 5.2.2.

**General feedback.** When asked about their thoughts about wearing SmokeMon to collect information about their activity, most participants (9 out of 11) expressed positive comments that affirm the unobtrusive design of SmokeMon. Participants report that it was *"comfortable"* (P18), *"easy"* (P16) and *"not a problem "* (P11) to wear SmokeMon. Also, participants report that they *"forgot"* (P12, P10) about the device, *"didn't notice it"* (P13), and that it *"did not get in the way"* (P9). Two participants (P17, P19) report that the system was *"bulky."* The size of the system made one of those participants self-conscious and embarrassed to wear it in public (P17). When asked about improvements that can be made to SmokeMon, three participants reported that they prefer a *"smaller"*, *"flat"*, and *"a more inconspicuous design"*. Also, three participants suggested improving the magnetic attachment mechanism to make it easier to use and more stable. One participant suggested attaching the device to clothing directly rather than using magnets. In general, the experience with the initial prototype of SmokeMon is favorable. We used inexpensive off-the-shelf components to develop SmokeMon to show the method's effectiveness. In

Table 8. Further energy optimizations that can be achieved with SmokeMon. P10 is not included in this analysis because ratio of smoking events to nonsmoking events is not similar to other participants due to the participants not following the study protocol.

| P# | % Discarded Frames | Added Battery Life | |
|---|---|---|---|
| | | BLE State | Storage State |
| P9 | 0.97 | 5.62 | 1.46 |
| P10 | - | - | - |
| P11 | 0.63 | 3.63 | 0.94 |
| P12 | 0.80 | 4.65 | 1.21 |
| P13 | 0.86 | 4.99 | 1.30 |
| P14 | 0.48 | 2.78 | 0.72 |
| P15 | 0.45 | 2.60 | 0.68 |
| P16 | 0.58 | 3.33 | 0.87 |
| P17 | 0.73 | 4.21 | 1.09 |
| P18 | 0.69 | 3.99 | 1.04 |
| P19 | 0.90 | 5.18 | 1.35 |
| **Mean** | 0.71 | 4.10 | 1.07 |

future iterations, the comments about the attachment style and the size of the system can be addressed by building a custom Printed Circuit Board (PCB) instead of utilizing off-the-shelf components.

**Participants comfort.** We asked the participant if any situations where wearing the device made them feel uncomfortable. None of the participants mentioned any discomforting situation. One participant (P14) noted, "I had a moment of hesitation the first time I used the restroom while wearing the device, but then I recalled what the researcher showed me re: what the device sees and doesn't see, and after that, I wasn't concerned." We also asked participants to report situations where they took off or covered the device to further understand concerns that SmokeMon can instigate. Participants reported taking off the device while sleeping and showering, which was expected as we told them to take it off during these times. One participant (P13) reported that they would take off the device if they wanted to "lie in positions where it would get in the way." Two participants (P19, P11) mentioned that they covered the sensor with a jacket or scarf when it was too cold.

We also asked participants if SmokeMon changed how they go about their daily activities in general and specifically while smoking. Eight participants mentioned that the device did not interfere with their daily activities. The remaining three (P10, P13, P17) mentioned that they were more conscious about their everyday activities that can interfere with the device (e.g., napping). Most participants (8 out of 11) reported that SmokeMon did not interfere with their smoking behavior. The remaining three participants (P10, P17, P18) reported interference. However, further inspection of their comments illustrates that the interference was caused by their misinterpretation of the study protocol rather than the system itself. For example, two participants did not keep track of the time they wore SmokeMon, so they wore it for more time and smoked at times which they usually do not smoke to finish the 10 hours of data collection. P10 thought that the protocol required switching on the device while smoking and turning it off otherwise, causing her to report that SmokeMon interfered with her smoking activity as she now has to remember to switch on the device.

**Bystander reaction.** We asked participants if the device made anyone around them uncomfortable or if anyone asked them to turn off the device. Two participants mentioned that people they know were curious about the device and what it records, which prompted participants to explain that the device captures only thermal images. None of the bystanders expressed any concerns. Instead, bystanders seem to be worried about capturing RGB or

audio data; as (P12) noted, "I had one friend who inquired about what the device was recording, and he seemed satisfied when I told him it wasn't recording audio or video."

**Attitude toward Data sharing.** After viewing their data, we asked participants how concerned they would be if their data were released to the following groups: other researchers, family, friends, doctors, dietitians, and the general public. Fig. 10 shows that most participants are "not at all concerned" with sharing their data with researchers, doctors, dietitians, and the general public. Two participants reported more concerns when their data is shared with their friends or family but not the general public, which may suggest sensitivity of sharing smoking behavior data with people they are familiar with as opposed to just sharing thermal images about them. Overall, these results show the promise of low-resolution thermal images providing sufficient rich data that people are willing to share with the public for the greater good.

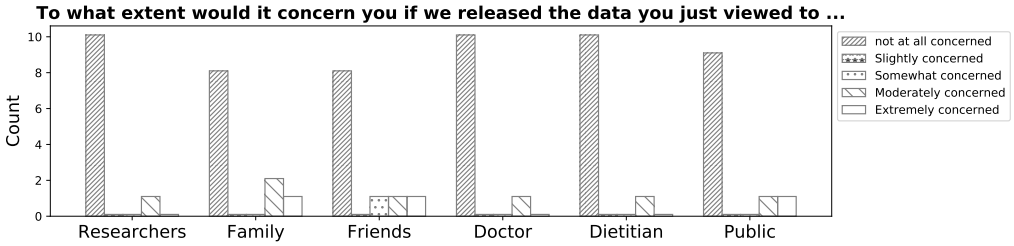


Fig. 10. Overall, participants are not at all concerned with sharing their data.

## 7 DISCUSSION AND FUTURE WORK

This work explores the idea of using a chest-worn, low-resolution thermal camera to reliably and unobtrusively capture smoking topography. The results suggest the considerable promise of thermal cameras in detecting smoking topography and smoking sessions among confounding activities while remaining energy-efficient, comfortable to wear, and low-burden. Below we discuss the main takeaways from our results, limitations, and future directions.

**SmokeMon detects smoking sessions reliably.** As shown in the results, we are able to detect smoking sessions with high recall and precision. Similarly, smart lighters can detect smoking sessions and the frequency of smoking [21, 37]. However, if the user is using another lighter or if they are sharing the lighter with someone else, it is hard to verify if the user was smoking or not. In the case of SmokeMon, we are able to verify if the wearer is smoking or not, both automatically using the SmokeMon algorithm or by manually inspecting the thermal images.

**SmokeMon detects smoking puffs using a single high-information sensing modality that is comfortable to wear.** Puff behaviors can vary between and within subjects, making the puff detection task very challenging. For example, as we have seen in our dataset, people alternate between the hand they use for smoking, can hold the cigarette in their mouth for a long time, and can engage in other non-smoking hand gestures while holding the cigarette. To overcome the puff detection challenge, prior work has either (1) instrumented multiple low-information sensing modalities located on different parts of the body [22, 36], which may decrease adherence to wearing the system, or (2) used an uncomfortable single high-information modality sensor (e.g., RGB cameras) [20]. Like RGB cameras, SmokeMon can capture challenging smoking cases without requiring additional sensor instrumentation. However, data collected by SmokeMon did not cause privacy concerns or discomfort to their wearer or the bystanders around them.

**SmokeMon unobtrusively extract smoking topography.** SmokeMon is the first system that attempts to utilize thermal sensors to extract smoking topography unobtrusively without requiring contact with the cigarette or interfering with the smoking habits. We show that, on average, SmokeMon performs as well as current smoking topography devices and ground truth. However, we believe that smoking topography results can be further improved by adding sensors to measure breathing patterns, especially when the cigarette is occluded by the user's hand or when the user is holding the cigarette in their mouth without the hand. Researchers are developing chest sensors that can measure respiration rate and volume [14]. In the future, such sensors can be placed in the magnetic pad of SmokeMon, which has direct contact with the user's skin.

## 7.1 Limitations

**Ground truth of smoked substance:** While we asked participants to label their smoking sessions, we did not ask them to identify the substance being smoked in their smoking log. This led to a confusing variability in smoking topography that affected model learning. For example, smoking topography and hand position differ between marijuana and cigarette smoking behaviors.

**Generlizability:** Although we show promising results, our dataset is limited, and we can not claim generalizability. An additional long-term study is needed to confirm the utility of thermal imaging under various conditions. For example, we notice that model performance suffers when participants smoke in a cold environment. This confusion is due to the imbalance of environmental temperatures in our dataset, as the vast majority of participants smoked in warm, indoor environments.

**Chest-worn devices sensitive to wear position:** Although we personalized device placement and sensor angle for each participant and added magnets for extra stability, we encountered data in which the sensor position was not ideal (e.g., device tilting). Although our data augmentation method accounts for such movements, it does not handle an extreme shift in position (i.e., when the device is too close to the head, resulting in an image that does not capture anything but the head). In such cases, we believe a trigger should be implemented on the device to alert the participants when the sensor is positioned incorrectly. Another approach is to increase the FoV of the camera. Unfortunately, changing the thermal camera lens is not as straightforward or as cheap as using a fisheye lens, as in the case of light-based RGB cameras. Changing the thermal image FoV requires knowledge in optics and money to manufacture the lens. However, in prior work, researchers have explored using multiple thermal cameras to increase the FoV [8]. In the case of SmokeMon, adding a second thermal camera above the current will double the vertical FoV. Adding a second camera will reduce battery life if both cameras are always on. Therefore, to reduce power consumption, the firmware should put the second camera to sleep and only trigger it to wake up during a possible smoking session (e.g., when the first camera observes very high-temperature pixels).

## 7.2 Future Directions

**Identifying between different smoking sessions:** While our data collection focuses on cigarette smoking, we noticed that many people smoke different substances and use different paraphernalia in addition to smoking cigarette tobacco. In some data, we observed smoking patterns from participants using electronic cigarettes (i.e., 'vapes') and water pipes (i.e., 'bongs'). Although e-cigarettes do not produce a thermal signature, the puffing behavior and the thermal reflection of the device are captured in the thermal image. Future studies in which data is collected from different smoking substances and paraphernalia would add another dimension to smoking topography; beyond its current fixation on tobacco cigarettes.

**Optimizing the pipeline to run on device:** As mentioned in section 4.2, our puff detection pipeline benefits from using a MobileNet v2, which is capable of running on edge at a very high speed. However, deploying these

models on-device often requires model compression and quantization, reducing accuracy. In the future, we aim to try our trained model on-device and, based on its real-time performance, optimize our model's architecture.

**Explore different thermal sensor positions beyond the chest:** In addition to the chest, we will try positioning thermal sensors on different parts of the body, including the head (e.g., glasses, hat). For example, we may be able to address challenges such as occlusion by utilizing a downward view. Other possible positions include the shoulders, where the cigarette's tip may be more visible across varying environments and body postures.

## 8 CONCLUSION

We present the development and evaluation of SmokeMon; a wearable system designed to unobtrusively measure smoking topography in free-living environments. SmokeMon uses deep learning to infer features of smoking episodes from the heat signatures captured by SmokeMon, allowing fine-grained measurement of cigarette smoking without any physical contact to the hand or cigarette. SmokeMon's battery life allows for continuous passive sensing throughout an entire waking day. A free-living study demonstrates the ability of SmokeMon to accurately detect smoking episodes and its topography in natural settings. Participants report high acceptability of the device and the data it collects. Overall, the results of this study show SmokeMon to have great potential in the nascent practice of measuring smoking topography.

## REFERENCES


[1] 2020. Smoking & Tobacco Use – Fast Facts. https://www.cdc.gov/tobacco/data_statistics/fact_sheets/fast_facts/index.htm. Accessed: 2021-02-04.

[2] 2020. Smoking Topography Study 2018. https://clinicaltrials.gov/ct2/show/NCT03498053. Accessed: 2021-02-04.

[3] 2020. Tobacco. https://www.who.int/news-room/fact-sheets/detail/tobacco. Accessed: 2021-02-04.

[4] 2021. CReSS Pocket. https://www.borgwaldt.com/en/products/smoking-vaping-machines/smoking-topography-devices/cress-pocket.html. Accessed: 2021-02-04.

[5] Yomna Abdelrahman, Alireza Sahami Shirazi, Niels Henze, and Albrecht Schmidt. 2015. Investigation of material properties for thermal imaging-based interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 15–18.

[6] Yomna Abdelrahman, Eduardo Velloso, Tilman Dingler, Albrecht Schmidt, and Frank Vetere. 2017. Cognitive heat: exploring the usage of thermal imaging to unobtrusively estimate cognitive load. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 33.

[7] Fatema Akbar, Ayse Elvan Bayraktaroglu, Pradeep Buddharaju, Dennis Rodrigo Da Cunha Silva, Ge Gao, Ted Grover, Ricardo Gutierrez-Osuna, Nathan Cooper Jones, Gloria Mark, Ioannis Pavlidis, et al. 2019. Email Makes You Sweat: Examining Email Interruptions and Stress Using Thermal Imaging. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 668.

[8] Rawan Alharbi, Chunlin Feng, Sougata Sen, Jayalakshmi Jain, Josiah Hester, and Nabil Alshurafa. 2021. HeatSight: Wearable Low-power Omni Thermal Sensing. In *2021 International Symposium on Wearable Computers*. 108–112.

[9] Rawan Alharbi, Tammy Stump, Nilofar Vafaie, Angela Pfammatter, Bonnie Spring, and Nabil Alshurafa. 2018. I Can't Be Myself: Effects of Wearable Cameras on the Capture of Authentic Behavior in the Wild. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 3, Article 90 (sep 2018), 40 pages. https://doi.org/10.1145/3264900

[10] Rawan Alharbi, Tammy Stump, Nilofar Vafaie, Angela Pfammatter, Bonnie Spring, and Nabil Alshurafa. 2018. I can't be myself: effects of wearable cameras on the capture of authentic behavior in the wild. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 1–40.

[11] Marielle C. Brinkman, Hyoshin Kim, Jane C. Chuang, Robyn R. Kroeger, Dawn Deojay, Pamela I. Clark, and Sydney M. Gordon. 2015. Comparison of True and Smoothed Puff Profile Replication on Smoking Behavior and Mainstream Smoke Emissions. *Chemical Research in Toxicology* 28, 2 (2015), 182–190. https://doi.org/10.1021/tx500318h arXiv:https://doi.org/10.1021/tx500318h PMID: 25536227.

[12] Thomas A. Burling, Maxine L. Stitzer, George E. Bigelow, and Andrew M. Mead. 1985. Smoking topography and carbon monoxide levels in smokers. *Addictive Behaviors* 10, 3 (Jan. 1985), 319–323. https://doi.org/10.1016/0306-4603(85)90014-0

[13] Youngjun Cho, Nadia Bianchi-Berthouze, Nicolai Marquardt, and Simon J Julier. 2018. Deep Thermal Imaging: Proximate Material Type Recognition in the Wild through Deep Learning of Spatial Surface Temperature Patterns. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 2.

[14] Michael Chu, Thao Nguyen, Vaibhav Pandey, Yongxiao Zhou, Hoang N Pham, Ronen Bar-Yoseph, Shlomit Radom-Aizik, Ramesh Jain, Dan M Cooper, and Michelle Khine. 2019. Respiration rate and volume measurements using wearable strain sensors. *NPJ digital medicine*

2, 1 (2019), 1–9.

[15] Ana Florescu, Roberta Ferrence, Tom Einarson, Peter Selby, Offie Soldin, and Gideon Koren. 2009. Methods for quantification of exposure to cigarette smoking and environmental tobacco smoke: focus on developmental toxicology. *Therapeutic drug monitoring* 31, 1 (2009), 14.

[16] Seungyeop Han, Rajalakshmi Nandakumar, Matthai Philipose, Arvind Krishnamurthy, and David Wetherall. 2014. GlimpseData: Towards continuous vision-based personal analytics. In *Proceedings of the 2014 workshop on physical analytics*. ACM, 31–36.

[17] Ronald I Herning, Reese T Jones, John Bachman, and Allan H Mines. 1981. Puff volume increases when low-nicotine cigarettes are smoked. *Br Med J (Clin Res Ed)* 283, 6285 (1981), 187–189.

[18] Fang Hu, Peng He, Songlin Xu, Yin Li, and Cheng Zhang. 2020. FingerTrak: Continuous 3D hand pose tracking by deep learning hand silhouettes captured by miniature thermal cameras on wrist. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2 (2020), 1–24.

[19] Masudul Imtiaz, Delwar Hossain, Volkan Senyurek, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2019. Wearable Egocentric Camera as a Monitoring Tool of Free-Living Cigarette Smoking: A Feasibility Study. *Nicotine & tobacco research : official journal of the Society for Research on Nicotine and Tobacco* 22 (11 2019). https://doi.org/10.1093/ntr/ntz208

[20] Masudul H Imtiaz, Delwar Hossain, Volkan Y Senyurek, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2019. Wearable Egocentric Camera as a Monitoring Tool of Free-Living Cigarette Smoking: A Feasibility Study. *Nicotine & Tobacco Research* (11 2019).

[21] Masudul Haider Imtiaz, Raul I Ramos-Garcia, Volkan Yusuf Senyurek, Stephen Tiffany, and Edward Sazonov. 2017. Development of a multisensory wearable system for monitoring cigarette smoking behavior in free-living conditions. *Electronics* 6, 4 (2017), 104.

[22] Masudul H Imtiaz, Raul I Ramos-Garcia, Shashank Wattal, Stephen Tiffany, and Edward Sazonov. 2019. Wearable sensors for monitoring of cigarette smoking in free-living: A systematic review. *Sensors* 19, 21 (2019), 4678.

[23] Takayuki Kawashima, Yasutomo Kawanishi, Ichiro Ide, Hiroshi Murase, Daisuke Deguchi, Tomoyoshi Aizawa, and Masato Kawade. 2017. Action recognition from extremely low-resolution thermal image sequence. In *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE, 1–6.

[24] Eric Larson, Gabe Cohn, Sidhant Gupta, Xiaofeng Ren, Beverly Harrison, Dieter Fox, and Shwetak Patel. 2011. HeatWave: thermal imaging for surface user interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2565–2574.

[25] Eun Lee, Jennifer Potts, Andrew Waters, Eric Moolchan, and Wallace Pickworth. 2003. Smoking topography: Reliability and validity in dependent smokers. *Nicotine & tobacco research : official journal of the Society for Research on Nicotine and Tobacco* 5 (10 2003), 673–9. https://doi.org/10.1080/1462220031000158645

[26] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. 2019. The Jester Dataset: A Large-Scale Video Dataset of Human Gestures. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. 2874–2882. https://doi.org/10.1109/ICCVW.2019.00349

[27] Matthew Louis Mauriello, Brenna McNally, and Jon E Froehlich. 2019. Thermporal: An Easy-to-Deploy Temporal Thermographic Sensor System to Support Residential Energy Audits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 113.

[28] Matthew Louis Mauriello, Leyla Norooz, and Jon E Froehlich. 2015. Understanding the role of thermography in energy auditing: current practices and the potential for automated solutions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1993–2002.

[29] Matthew Louis Mauriello, Manaswi Saha, Erica Brown Brown, and Jon E Froehlich. 2017. Exploring novice approaches to smartphone-based thermographic energy auditing: A field study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 1768–1780.

[30] Melissa Mercincavage, Joshua L Karelitz, Catherine L Kreider, Valentina Souprountchouk, Benjamin Albelda, and Andrew A Strasser. 2021. Comparing video observation to electronic topography device as a method for measuring cigarette puffing behavior. *Drug and Alcohol Dependence* 221 (2021), 108623.

[31] Abhinav Parate, Meng-Chieh Chiu, Chaniel Chadowitz, Deepak Ganesan, and Evangelos Kalogerakis. 2014. Risq: Recognizing smoking gestures with inertial sensors on a wristband. In *Proceedings of the annual international conference on Mobile systems, applications, and services (Mobisys)*.

[32] RJ Robinson, EC Hensel, PN Morabito, and KA Roundtree. 2015. Electronic cigarette topography in the natural environment. *PloS one* 10, 6 (2015), e0129296.

[33] Daniel JO Roche, Spencer Bujarski, Emily Hartwell, ReJoyce Green, and Lara A Ray. 2015. Combined varenicline and naltrexone treatment reduces smoking topography intensity in heavy-drinking smokers. *Pharmacology Biochemistry and Behavior* 134 (2015), 92–98.

[34] Kathryn C Ross and Laura M Juliano. 2016. Smoking through a topography device diminishes some of the acute rewarding effects of smoking. *Nicotine & Tobacco Research* 18, 5 (2016), 564–571.

[35] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2019. MobileNetV2: Inverted Residuals and Linear Bottlenecks. arXiv:cs.CV/1801.04381

[36] Edward Sazonov, Paulo Lopez-Meyer, and Stephen Tiffany. 2013. A wearable sensor system for monitoring cigarette smoking. *Journal of studies on alcohol and drugs* 74, 6 (2013), 956–964.

[37] Philipp M Scholl and Kristof Van Laerhoven. 2017. Lessons learned from designing an instrumented lighter for assessing smoking status. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*. 1016–1021.

[38] Sougata Sen, Vigneshwaran Subbaraju, Archan Misra, Rajesh Krishna Balan, and Youngki Lee. 2015. The case for smartwatch-based diet monitoring. In *2015 IEEE international conference on pervasive computing and communication workshops (PerCom workshops)*. IEEE, 585–590.

[39] Volkan Senyurek, Masudul Imtiaz, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2019. Cigarette Smoking Detection with An Inertial Sensor and A Smart Lighter. *Sensors* 19, 3 (2019). https://doi.org/10.3390/s19030570

[40] Volkan Senyurek, Masudul Imtiaz, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2019. Cigarette smoking detection with an inertial sensor and a smart lighter. *Sensors* 19, 3 (2019), 570.

[41] Volkan Y Senyurek, Masudul H Imtiaz, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2019. Smoking detection based on regularity analysis of hand to mouth gestures. *Biomedical signal processing and control* 51 (2019), 106–112.

[42] Muhammad Shoaib, Hans Scholten, Paul JM Havinga, and Ozlem Durmaz Incel. 2016. A hierarchical lazy smoking detection algorithm using smartwatch sensors. In *IEEE International Conference on e-Health Networking, Applications and Services (Healthcom)*.

[43] Andrew L Skinner, Christopher J Stone, Hazel Doughty, and Marcus R Munafò. 2019. StopWatch: The preliminary evaluation of a smartwatch-based system for passive detection of cigarette smoking. *Nicotine and Tobacco Research* 21, 2 (2019), 257–261.

[44] Andrew A Strasser, Wallace B Pickworth, Freda Patterson, and Caryn Lerman. 2004. Smoking topography predicts abstinence following treatment with nicotine replacement therapy. *Cancer Epidemiology and Prevention Biomarkers* 13, 11 (2004), 1800–1804.