



Principled Mining, Forecasting, and Monitoring of Honeybee Time Series with EBV+

MST SHAMIMA HOSSAIN, University of California, Riverside, California, USA

CHRISTOS FALOUTSOS, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

BORIS BAER, HYOSEUNG KIM, and VASSILIS J. TSOTRAS, University of California, Riverside, California, USA

Honeybees, as natural crop pollinators, play a significant role in biodiversity and food production for human civilization. Bees actively regulate hive temperature (homeostasis) to maintain a colony's proper functionality. Deviations from usual thermoregulation behavior due to external stressors (e.g., extreme environmental temperature, parasites, pesticide exposure) indicate an impending colony collapse. Anticipating such threats by forecasting hive temperature and finding changes in temperature patterns would allow beekeepers to take early preventive measures and avoid critical issues. In that case, how can we model bees' thermoregulation behavior for an interpretable and effective hive monitoring system? In this article, we propose the *principled* Electronic Bee-Veterinarian Plus (EBV+) method based on the thermal diffusion equation and a novel "sigmoid" feedback-loop (P) controller for analyzing hive health with the following properties: (i) it is *effective* on multiple, real-world beehive time sequences (recorded and streaming), (ii) it is *explainable* with only a few parameters (e.g., hive health factor) that beekeepers can easily quantify and trust, (iii) it issues *proactive* alerts to beekeepers before any potential issue affecting homeostasis becomes detrimental, and (iv) it is *scalable* with a time complexity of $O(t)$ for reconstructing and $O(t \times m)$ for finding m cuts of a sequence with t time-ticks. Experimental results on multiple real-world time sequences showcase the potential and practical feasibility of EBV+. Our method yields accurate forecasting (up to 72% improvement in RMSE) with up to 600 times fewer parameters compared to baselines (ARX, seasonal ARX, Holt-winters, and DeepAR), as well as detects discontinuities and raises alerts that coincide with domain experts' opinions. Moreover, EBV+ is scalable and fast, taking less than 1 minute on a stock laptop to reconstruct 2 months of sensor data.

CCS Concepts: • **Mathematics of computing** → **Time series analysis**; • **Information systems** → **Data mining**; • **Applied computing** → **Agriculture**;

Additional Key Words and Phrases: Honeybees, Forecasting, Event-detection

Associate Editor: Jason Wang

This work is supported by the U.S. Department of Agriculture's National Institute of Food and Agriculture grants (2024-67022-43695 and 2024-67021-43696), UC Multicampus Research Programs and Initiatives (MRPI) award (# M21PR2306), and UCR Delfino Agriculture Innovation Seed Fund.

Authors' Contact Information: Mst Shamima Hossain (corresponding author), University of California, Riverside, California, USA; e-mail: mhoss037@ucr.edu; Christos Faloutsos, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA; e-mail: christos@cs.cmu.edu; Boris Baer, University of California, Riverside, California, USA; e-mail: boris.baer@ucr.edu; Hyoseung Kim, University of California, Riverside, California, USA; e-mail: hyoseung@ucr.edu; Vassilis J. Tsotras, University of California, Riverside, California, USA; e-mail: vtsotras@ucr.edu.



This work is licensed under Creative Commons Attribution International 4.0.

© 2025 Copyright held by the owner/author(s).

ACM 1556-472X/2025/5-ART100

<https://doi.org/10.1145/3719014>

ACM Reference format:

Mst Shamima Hossain, Christos Faloutsos, Boris Baer, Hyoseung Kim, and Vassilis J. Tsotras. 2025. Principled Mining, Forecasting, and Monitoring of Honeybee Time Series with EBV+. *ACM Trans. Knowl. Discov. Data.* 19, 5, Article 100 (May 2025), 30 pages.
<https://doi.org/10.1145/3719014>

1 Motivation and Background

How can we design better tools for beekeepers to effectively monitor honeybee hive health status and receive proactive real-time notifications to take preventive action to avoid colony collapse?

Motivation. Without honeybees, humanity will suffer. Honeybees are vital pollinators. Their pollination services are required to produce more than 80 crops or about a third of what we eat [13]. The yearly value of bee pollination has been estimated to be up to \$550 billion globally and \$29 billion for the United States alone [32]. However, beekeepers have reported that they have lost almost half of their colonies in 2023, with a substantially higher colony loss of 48.2% from April 2022 to April 2023 in the United States [24, 28]. The reasons for the decline are multiple, including parasites, pesticides, and, most recently, climate change [19, 27].

Background. As part of the 2nd CIBER Bee Health Conference held at UC Riverside in 2022 [38], practitioners were very emphatic on the need for a reliable and explainable hive monitoring system to help them understand the hive’s internal condition for timely action. Thus, we need to stay away from black-box methods and instead use parsimonious, explainable models; such methods have few parameters and can easily justify anticipated extremities to domain experts.

Like a human body, hive temperature gives valuable information about hive health. Bees tightly regulate their hive temperature (homeostasis) by heating during cold periods and cooling during the heat. This is known as *thermoregulation*, resulting in hive temperatures staying between 33 – 36°C when it is healthy [23, 25]. Environmental stressors (e.g., parasites, pesticide exposure, extreme weather conditions, food shortages, drought) can cause bees to lose their ability to maintain the hive’s homeostasis [33]. Thus, increased fluctuation of hive temperature can be seen as a first-order response indicating an impending colony collapse.

Method Overview. We propose a method, **Electronic Bee-Veterinarian Plus (EBV+)**, specifically designed to monitor hive temperature. We applied EBV+ to multiple real-time temperature sequences collected from different hive settings (Figure 1(a), details in Section 3) and achieved reliable forecasts for several days ahead with a good fit and an improvement of at least 680 points in terms of **Akaike Information Criterion (AIC)** [9] over baselines (Figure 1(c), details in Section 5.2); a lift of up to 72% in the area under the curve over baselines (Figure 6); the computation time of less than 1 minute for fitting 2 months of data; and segmentation (Figure 1(b)) and proactive alerts (Figure 1(d)) coinciding with external disruptive events affecting bees’ thermoregulation ability (details in Sections 5.1 and 5.3).

Contributions. Specifically, the main contributions of this article are summarized as follows:

- *Principled:* We propose a method, EBV+, comprising (i) $EBV_{\text{model}+}$: a novel model based on the thermal diffusion equation, control theory, and our novel “*sigmoid*” P-controller idea (Section 4.2); (ii) $EBV_{\text{fit\&cut}+}$: a segmentation algorithm (Section 4.5) based on $EBV_{\text{model}+}$ to estimate the number of cut-points where each cut represents a different hive condition; and (iii) $EBV_{\text{fcast\&warn}+}$: a proactive alerting algorithm to warn the beekeepers about deteriorating




Fig. 1. EBV+ at work: (a) deployed sensors, (b) reconstruction and event detection, (c) EBV+ wins on forecasting hive core temperature, (d) proactive notification.

hive health (Section 4.6). As all honeybees follow the same mechanism for thermoregulation, our method is universally applicable.

- *Effective:* We use the proposed $EBV_{fit\&cut+}$ and $EBV_{fcast\&warn+}$ for (i) fitting, (ii) forecasting, (iii) segmentation, and (iv) proactive notification to detect unexpected events. It performs well in terms of forecasting accuracy.
- *Explainable:* We can explain the current hive condition with the parameters (i.e., hive health factor in Section 4.2) measured and discontinuity (offline and online) detected by $EBV_{fit\&cut+}$ and $EBV_{fcast\&warn+}$.
- *Proactive:* Based on the forecasted value of hive temperature using $EBV_{fcast\&warn+}$, we can issue proactive alerts to beekeepers about deteriorating hive health before potential issues become critical.
- *Scalable:* Our method is scalable and hands-off, requiring no parameter tuning. It applies to all hives irrespective of hive size and bee population.
- *Informative:* Our experimental findings coincide with domain experts' expectations: (i) bees in treated hives are less affected by intensive hive inspection; (ii) treated colonies have a higher health factor (i.e., healthier) to maintain thermoregulation than control ones; and (iii) heating a hive is easier for honeybees than cooling it down [17].

Reproducibility. Our implementation and the full dataset are open sourced at <https://github.com/rtenlab/EBeeVet>.

This is an extended version of our prior work [26] with the following new contributions: (i) we propose a thermal diffusion and “sigmoid” P-controller-based model, EBV_{model+} , to add a

Table 1. EBV+ wins, Fulfilling all the Specifications

Property \ Method	Prior hive monitoring [15, 29]	Traditional time series [16, 39]	Deep learning techniques [10]	EBV [26]	EBV+ (Proposed)
Domain-informative	✓			✓	✓
Effective		?	?	✓	✓
Explainable	?			✓	✓
Scalable		✓	?	✓	✓
Proactive alerts		?	?	?	✓

“?” means “maybe, depending on implementation”.

practical upper limit on bee thermoregulation ability and solve the convergence issue in extreme scenarios in the previous EBV_{model} model; (ii) we propose a bi-level optimization method for global parameter estimation for hives of the same geographic region, genotype, and hive setup, and the hive-specific parameters estimation to quantify and compare hive health across different hives; and (iii) we propose a proactive notification algorithm, $EBV_{\text{fcast\&warn+}}$, with additional experimentation demonstrating its effectiveness to alert beekeepers about deteriorating hive health in real time.

2 Related Work

Beehive Monitoring and Analysis. Past systems collect in-hive microclimate parameters like temperature, humidity, and CO_2 [6, 12, 30, 35], sound signals [14, 34], and visual bee activities using IR LEDs or cameras installed at the hive entrance [11, 18]. Efforts have been made mostly to analyze sound signals to detect swarming events [14, 34]. Attempts to analyze other in-hive data such as temperature to detect anomalies [15, 29] are limited to *arbitrary thresholds* defined by domain specialist, rule-based, or simple autoencoders. All these approaches fail to provide causal relationships with colony health, lack robust evidence for analysis, and suffer from poor sensing quality or intermittent interruptions due to environmental conditions.

Traditional Time Series Forecasting. Time series have attracted huge research interest [3–5, 7, 8, 20, 22, 37]. Classical methods, such as (i) **Autoregression (AR)**-based methods with exogenous variable [16], e.g., ARX/ARMAX/ARIMAX and seasonal ARX/ARMAX/ARIMAX, and (ii) exponential smoothing models [39], e.g., Holt-Winters, can be used for forecasting, but they are neither domain-informative nor explainable. The hyperparameters of these models do not give any insights into the hive’s current condition.

Deep Learning Techniques. Such recent methods include Autogluon-Timeseries [10] (an open source autoML library for time-series forecasting), DeepAR [31], Temporal Fusion Transformer [21], and so on. The main drawback of these methods is the lack of explainability, despite recent attempts [36].

Contrast to the Existing Literature. Existing hive health monitoring systems are supervised in nature, based on either rule-based [15] or black-box methods [34]. Our prior work, EBV [26], was the first attempt to address this problem in an unsupervised manner. However, one limitation of our prior work is its inability to handle scenarios where hive strengths do not converge, potentially leading to an unrealistic estimation of infinite strength values. It also lacked the ability to provide proactive alerts to beekeepers. *EBV+* solves these issues and satisfies all the desirable properties listed in Table 1.

3 Dataset and Experimental Setup

We have used a dataset collected from August to September 2021 (late Summer) from an apiary in UC Riverside, CA, USA. This period in that region has very hot and dry climate conditions, with peak temperature exceeding 40°C. Such conditions can cause significant heat stress but do not lead to colony collapse as long as bees can manage in-hive temperature.

A total of 10 single-box homogeneous Langstroth hives, each consisting of 10 deep frames, were used. Hives were divided equally into two groups for experimental purposes: *control hives* and *treated hives*. The selection of hives for each group was done randomly. Five control hives received no treatment throughout the experiment period. The remaining five hives were treated in response to extreme temperatures by placing approximately 6 kg of ice on top of them when the environment temperature exceeded 35°C (details in Section 5.1, Figure 4).

For data collection, we used a commercial multi-sensor device, Nordic Thingy:52 [41], which has wireless communication to transfer collected data. We recorded temperature per 20-minute intervals for 2 months (August–September 2021). Every 2 weeks throughout that period, hives were opened twice: (i) for health evaluation and taking out sensor devices for battery recharging and (ii) for reinstalling recharged devices the next day. While all hives maintained the core temperature within 33 – 36°C, the control ones had a much higher deviation over several days, implying that they were affected by heat stress (Figure 7).

Sensor Placement. We kept each sensor device inside a box so that bees were not able to cover the hardware with wax or propolis, affecting its sensing ability. Each hive had two sensor devices placed in the following locations (Figure 1(a)):

- (1) *Peripheral Area*: On top of the outer frame, which typically coincides with an area used for honey storage.
- (2) *Core Area*: On the central frame near the brood area, where developing eggs, larvae, and pupae, as well as the queen, are present. The temperature of this area is crucial to hive health as developing broods are susceptible to heat stress, and healthy hives maintain the core area temperature between 33 – 36°C [23, 25].

We also installed one sensor device in an empty Langstroth box. As this empty hive had no bees inside regulating the core temperature, we consider the external temperature, $\Theta_{ext}(t)$, to be the outer surface temperature of this hive.¹

4 Proposed Method: EBV+

Here, we outline our principled models for reconstruction and forecasting with justification. First, we present our basic model, EBV_{model} (Sections 4.1 and 4.1.1 [26]), based on physics and control theory. Then, we describe the advanced version, EBV_{model+} (Sections 4.2 and 4.2.1), that overcomes the limitations of the basic EBV_{model} model, as pointed out in Section 2. Table 2 summarizes the symbols and their definitions used in the article.

4.1 Basic Model: EBV_{model}

We first present our basic model, EBV_{model} , to monitor and forecast the hive core temperature $\Theta(t)$, given the outside temperature $\Theta_{ext}(t)$. To capture thermoregulation ability, our model introduces the novel terms called “hive strengths” defined as follows:

Definition 1 Hive Strengths (s_c : *Cooling*, s_h : *Heating*). These parameters refer to the daily health level of bees in a single hive and are inversely proportional to the fluctuation in $\Theta(t)$. The lesser the fluctuation in $\Theta(t)$, the stronger the hive is, and vice versa.

¹This is the surface of the outer cover for protecting the outer frame (peripheral area) from the weather.

Table 2. Symbols and Definitions

Symbols	Definitions
$\mathcal{D}$	Temperature dataset
$\mathcal{M}$	Model in Equation (2)
$\mathcal{L}(\mathcal{D} \mathcal{M})$	Likelihood of $\mathcal{D}$ given $\mathcal{M}$
θ_{ideal}	Ideal core temperature
$\Theta_{ext}(t)$	External temperature
$\Theta_{peri}(t)$ and $\Theta(t)$	Hive's peripheral and core temperatures
$\theta_{ext}(t), \theta_{peri}(t)$, and $\theta(t)$	Relative external, hive peripheral, and hive core temperatures to ideal core temperature, e.g., $\theta(t) = \Theta(t) - \theta_{ideal}$
$\hat{\Theta}(t)$	Fitted/forecasted $\Theta(t)$
$\Theta_{adj}(t)$	Second surface temperature based on hive type
s_c, s_h	Hive strength of cooling and heating
r_c, r_h	Thermoregulation capacity limit (cooling and heating)
h	Hive health factor

Here, the values of (s_c, s_h) give a quantitative measure of how good the bees are at thermoregulation, that is, at maintaining temperature $\Theta(t)$ given $\Theta_{ext}(t)$. Good ability to thermoregulate is a sign of hive health; thus, high values of s_c and s_h can be used to compare health among different hives having the same setup, genotype, and geographic region. A healthy hive will have higher (s_c, s_h) and lower fluctuation in temperature $\Theta(t)$. Based on the strength parameters, our model is given by the following lemma.

MODEL 1 (EBV_{model}). *Given the external temperature $\Theta_{ext}(t)$ and the strengths s_h and s_c of the hive, the hive core temperature, $\Theta(t)$, obeys:*

$$\frac{\partial \theta(t)}{\partial t} = \begin{cases} \theta_{ext}(t) + \theta_{adj}(t) - 2\theta(t) - s_c\theta(t), & \text{if } \theta_{ext}(t) \geq 0. \\ \theta_{ext}(t) + \theta_{adj}(t) - 2\theta(t) + s_h\theta(t), & \text{otherwise.} \end{cases} \quad (1)$$

where

$$s_c, s_h > 0 \quad \text{and} \quad \theta_{adj}(t) = \begin{cases} \theta_{ext}(t), & \text{if hive type = control} \\ \theta_{peri}(t), & \text{if hive type = treated.} \end{cases}$$

Next, we provide the justification of our naive model, EBV_{model} (Equation (1)), by using first principles: Thermal diffusion and Control theory.

4.1.1 Justification. Here, we consider all temperatures relative to the hive's ideal core temperature, θ_{ideal} , as bees' efforts to keep the core temperature, $\Theta(t)$, within the normal range depends on its deviation from θ_{ideal} . It is worth noting that the ideal temperature θ_{ideal} is within 33 – 36°C, but the exact value can vary slightly within this range, from day to day and from hive to hive, due to several reasons, such as hive size, population, and season. Below, we give a detailed step-by-step explanation for our basic model, EBV_{model} .

Step 0: Overview. We consider three key areas of the hive to capture bees' thermoregulation mechanism:

- (1) Hive's core area with temperature, $\theta(t)$
- (2) Hive's top surface (adjacent to ice) with temperature, $\theta_{adj}(t)$

- (3) Other surfaces with temperature, $\theta_{ext}(t)$: for simplicity, we assume other surfaces are homogeneous as they directly face external temperature.

Step 1: Physics—Thermal Diffusion. The first term in Equation (1) is due to heat transfer: for an empty hive, the rate of change in hive core temperature, $\frac{\partial\theta(t)}{\partial t}$, obeys the thermal diffusion model [1],

$$\nabla^2\theta = \frac{\partial^2\theta}{\partial x^2} + \frac{\partial^2\theta}{\partial y^2} + \frac{\partial^2\theta}{\partial z^2} \propto \frac{\partial\theta}{\partial t}$$

Heat transfer occurs from the areas of higher temperature to the areas of lower temperature. Also, heat transfer considering a linear plane will capture the same effect as heat transfer in a 3D hive box (homogeneous outer surfaces of hives). Thus, in our hive model:

- Heat transfers from the top surface to the core area at a rate proportional to $\theta_{adj}(t) - \theta(t)$
- Heat transfers from other surfaces to the core area at a rate proportional to $\theta_{ext}(t) - \theta(t)$

The rate of change in core temperature, $\frac{\partial\theta(t)}{\partial t}$, is the net heat flow into the core area: $\frac{\partial\theta(t)}{\partial t} \propto (\theta_{adj}(t) - \theta(t)) + (\theta_{ext}(t) - \theta(t))$, which can be rewritten as $\frac{\partial\theta(t)}{\partial t} \propto \theta_{ext}(t) + \theta_{adj}(t) - 2\theta(t)$.

Step 2: Control Theory—(split) P-controller. The second term is the feedback loop [40]: when the hive core temperature $\Theta(t)$ is away from ideal core temperature θ_{ideal} , bees try to bring it to zero with strengths s_c and s_h and their efforts are proportional to $\theta(t)$, i.e., if bees are left alone,

$$\frac{\partial\theta(t)}{\partial t} \propto \begin{cases} -s_c\theta(t), & \text{if } \theta_{ext}(t) \geq 0. \\ +s_h\theta(t), & \text{otherwise.} \end{cases}$$

Notice that we propose to have a *split* in the standard proportional (P) controller: the bees have different ability to cool (s_c : when $\Theta_{ext}(t)$ goes above θ_{ideal}) and different ability to heat (s_h : when $\Theta_{ext}(t)$ goes below θ_{ideal})—this is exactly what we call a *split* P-controller. We have followed the sign convention: positive work for heating and negative work for cooling. Here, we chose to use a P-controller as it appropriately replicates bees' thermoregulation behavior. Biologically, while bees (and P-controllers) aim to maintain $\theta(t) = 0$, they face continuous offsets from $\theta_{ext}(t)$ and $\theta_{adj}(t)$. This results in a dynamic regulation process: $\theta(t)$ shifts, making bees keep adjusting their efforts proportionally. This allows us to observe that $\theta(t)$ hovers close to (but not exactly at) the equilibrium point in a healthy state. Our P-controller mirrors this actual behavior of bees in nature.

Using the superposition principle for Step 1 (thermal diffusion) and Step 2 (split P-controller), we can combine the right-hand side of the two equations that make it the same as Equation (1).

4.1.2 Stability Analysis. From Equation (1), our basic model, EBV_{model} , has two switching modes:

$$\frac{\partial\theta(t)}{\partial t} = \begin{cases} \theta_{ext}(t) + \theta_{adj}(t) - 2\theta(t) - s_c\theta(t), & \text{for cooling} \\ \theta_{ext}(t) + \theta_{adj}(t) - 2\theta(t) + s_h\theta(t), & \text{for heating} \end{cases}$$

To linearize the system, let us define $x = -s_h$ or s_c .

Then, $\frac{\partial\theta(t)}{\partial t} = \theta_{ext}(t) + \theta_{adj}(t) - (2 + x)\theta(t)$.

If we apply the Laplace transform, the above equation becomes: $s\theta(s) = \theta_{ext}(s) + \theta_{adj}(s) - (2 + x)\theta(s)$.

We get the transfer function (ratio of output to input) by rearranging the above equation, which shows how the system responds to input temperatures ($\theta_{ext}(t)$, $\theta_{adj}(t)$) in the Laplace domain:

$$\frac{\theta(s)}{\theta_{ext}(s) + \theta_{adj}(s)} = \frac{1}{s + 2 + x}$$

Here, the pole is: $s + 2 + x = 0 \Rightarrow s = -(2 + x)$. If the real part is negative, the system is stable. Hence, the system is stable if $-(2 + x) < 0 \Rightarrow s_c > -2$ and $s_h < 2$. In our hive model, the system is always stable while cooling ($\because s_c > 0$) and stable while heating if $0 < s_h < 2$. However, in practice, healthy bees tend to maintain the hive core temperature within a suitable range for survival even when s_c and s_h fall outside these conditions.

The theoretical equilibrium point can be computed by setting $\frac{\partial\theta(t)}{\partial t} = 0$ and solving for $\theta(t)$:

- For cooling, $\theta(t)$ at equilibrium = $(\theta_{ext}(t) + \theta_{adj}(t))/(2 + s_c)$
- For heating, $\theta(t)$ at equilibrium = $(\theta_{ext}(t) + \theta_{adj}(t))/(2 - s_h)$

Note that this theoretical equilibrium point is never 0; rather it is always a fraction of the unregulated temperature $(\theta_{ext}(t) + \theta_{adj}(t))/2$. Also notice, if the external temperature has a (daily) periodicity, the hive temperature will, too, with a smaller amplitude. This is exactly what we observed in our hives, and this is the reason that we propose the simple, P-controller, as opposed to a more complicated one.

4.1.3 Limitations. There are a few limitations of the *split* P-controller that we want to address with our advanced model, EBV_{model+} , in the next section.

- (1) *Abrupt switching between modes.* Bees cannot switch between two modes (heating and cooling) abruptly² (Figure 2(a)): the transition should be smooth.
- (2) *No defined boundary for bee strength (s_c, s_h).* The *split*-P controller response in EBV_{model} is proportional to $\theta(t)$ with no defined upper-limit on bee strengths (Figure 2(a)). But bees cannot increase their strength with increasing $\theta(t)$. They will eventually get exhausted at one point as they possess finite strength.

Since EBV_{model} does not impose any boundary on the solver during optimization, it may find an unexpectedly larger value (can be considered infinity) for (s_c, s_h) in case of extreme conditions: if $\Theta(t)$ is equal to θ_{ideal} , i.e., there is no change in the relative hive core temperature $\theta(t)$ with respect to time t , Equation (1) becomes,

$$\frac{\partial\theta(t)}{\partial t} = 0 = \theta_{ext}(t) + \theta_{adj}(t)$$

For parameters to converge in this scenario, either of the following conditions should hold: (i) $\theta_{ext}(t) = \theta_{adj}(t) = 0$, or (ii) $\theta_{ext}(t) = -\theta_{adj}(t)$, both of which are unfortunately unrealistic.

4.2 Advanced Model: EBV_{model+}

To address the aforementioned problem of the basic model, we now present our advanced version, EBV_{model+} . Unlike the basic model that uses strength parameters, the advanced model uses the novel “health factor” coupled with the thermoregulation capacity limits.

Definition 2 Thermoregulation Capacity Limits (r_c : *Cooling*, r_h : *Heating*). Bees cannot have infinite thermoregulation ability as it is limited by the number of bees in a hive. Without loss

²The biological reason behind this is that heating means muscular contractions of honeybees to turn carbohydrates into energy, whereas cooling requires evaporation of water in nectar inside the hive. As each individual bee provides a unit of heating or cooling that adds up to the colony level, it makes more sense to model the transition between these two mechanisms as a gradual change than an abrupt one.




Fig. 2. $EBV+$ vs. EBV : relative hive temperature, $\theta(t)$ vs. controller response. (a) EBV_{model} . (b) EBV_{model+} .

of generality, we assume the same thermoregulation capacity limits (r_c and r_h) apply across hives with the same setup, genotype, and geographic region.

Definition 3 Hive Health Factor (h). The idea behind the daily health factor h is similar to daily hive strengths (s_c, s_h): the lesser fluctuation in $\Theta(t)$, the higher h (the stronger the hive) is and vice versa. The difference is that it is a relative representation of the hive's thermoregulation capacity limits (r_c, r_h).

Here, thermoregulation capacity limits (r_c, r_h) map to the maximum possible thermoregulation ability (controller response; Figure 2(b)) when bees manage hive temperature, $\Theta(t)$ infinitely away from θ_{ideal} . Bees cannot go beyond a certain thermoregulation limit to maintain $\Theta(t)$ irrespective of how far $\Theta(t)$ is from the ideal scenario. Health factor (h) scales the controller response to match the fluctuation in $\Theta(t)$ and is proportional to how well bees manage the fluctuation in $\Theta(t)$. Higher h corresponds to a stronger hive (less fluctuation in $\Theta(t)$) and vice versa. Using the capacity limits and health factor, the following lemma gives our advanced model, EBV_{model+} .

MODEL 2 (EBV_{model+}). Given the external temperature $\Theta_{ext}(t)$, the maximum thermoregulation capacity limits r_h and r_c , and the hive health factor h , the hive core temperature, $\Theta(t)$, obeys:

$$\frac{\partial \theta(t)}{\partial t} = \theta_{ext}(t) + \theta_{adj}(t) - \underbrace{2\theta(t) - h\sigma(r_c, r_h; \theta(t))}_{\text{controller response}} \quad (2)$$

$$\text{where } \sigma(r_c, r_h; \theta(t)) = \frac{r_c r_h (1 - e^{-\theta(t)})}{r_h + r_c e^{-\theta(t)}} \quad (3)$$

Figure 2 illustrates the difference in controller response between the basic model (EBV) and the advanced model ($EBV+$). After a certain increase or decrease in $\theta(t)$, bees reach their maximum cooling capacity (r_c) or heating capacity (r_h) and continue functioning at this limit even if $\theta(t)$ increases or decreases beyond that point. We justify our advanced model EBV_{model+} (Equation (2)) using the first principles of physics and control theory.

4.2.1 Justification. The first three terms in Equation (2) follow the thermal diffusion model for an empty hive (same as the basic model in Section 4.1). The last term is our “sigmoid” P-controller response³: we limit the controller gain within the range of $[-r_h, r_c]$ with a smooth transition between two modes (heating and cooling). Also, the lower the health factor h is, the weaker the bees are: $\Theta(t)$ fluctuates more. At any time t , hive health factor h will be a fraction ($0 < h \leq 1$) of

³Sigmoid-based controllers have been studied for control applications, such as [2].

the maximum possible controller response, $\sigma(r_c, r_h; \theta(t))$, i.e., if bees are left alone,

$$\frac{\partial \theta(t)}{\partial t} = -h\sigma(r_c, r_h; \theta(t)) = -h \frac{r_c r_h (1 - e^{-\theta(t)})}{r_h + r_c e^{-\theta(t)}}$$

The health factor h and the transition mode between two temperature control modes represent how healthy a hive is: the quicker the transition is, and the higher h is, the healthier a hive is. Here, we summarize these two factors in one parameter, h , assuming the transition rate of 1 for simplicity.

Our *sigmoid* P-controller is derived as follows: Consider a sigmoid function $f(x)$ that limits the system output x in the range of $[0, 1]$ with a transition rate of 1:

$$f(x) = \frac{1}{1 + e^{-x}}$$

Based on this, we define the *sigmoid* P-controller by:

$$\sigma(\theta(t)) = \frac{A}{1 + e^{-(\theta(t)-C)}} + B \quad (4)$$

$$\text{where } \lim_{\theta(t) \rightarrow \infty} \sigma(\theta(t)) = r_c, \lim_{\theta(t) \rightarrow -\infty} \sigma(\theta(t)) = -r_h, \text{ and } \lim_{\theta(t) \rightarrow 0} \sigma(\theta(t)) = 0$$

Here, A , B , and C are the scaling, vertical, and horizontal shifting constants of the *sigmoid* P-controller, respectively. We have followed the sign convention to define the limits: positive work for heating and negative work for cooling. When the hive core temperature, $\Theta(t)$, is infinitely away from θ_{ideal} , bees reach their maximum thermoregulation capacity limit r_c and r_h . If $\Theta(t)$ is the same as θ_{ideal} , bees will stop heating or cooling the hive until the offsets from $\theta_{ext}(t)$ and $\theta_{adj}(t)$ cause $\theta(t)$ to shift, making bees start thermoregulation again. Solving Equation (4) with the above-defined limits gives the feedback controller term in Equation (3).

We use the superposition principle to combine these two effects, thermal diffusion and *sigmoid* P-controller, that makes the equation the same as Equation (2). Our goal is to find the values of cooling and heating capacity limits r_c and r_h , the health factor h , and the ideal core temperature θ_{ideal} that best reconstructs $\Theta(t)$ given $\Theta_{ext}(t)$ and $\Theta_{adj}(t)$.

4.3 Proposed Optimization Objective

In this section, we outline our optimization objective functions to find out the parameters for reconstructing $\Theta(t)$. First, we present a basic objective function to minimize the reconstruction error. Next, we take into account the dependence of current hive health h on its past values to model the daily change in h . Although the objective functions are defined for parameter estimation using EBV_{model+} , they can also be used for EBV_{model} with minor modifications.

Method 1 (Basic). If the health factor of a single hive on any day is independent of the previous day's health factor and also independent of other hives of similar specification, our optimization objective will be to minimize the reconstruction error as follows:

$$\underset{s_c, s_h}{\text{minimize}} L = \sum_{c=1}^C \sum_{t=1}^N (\Theta_c(t) - \hat{\Theta}_c(t))^2 \quad (5)$$

Here, C and N correspond to the total number of hives and time ticks in the experimental period, and $\hat{\Theta}_c(t)$ is the reconstructed temperature sequence for the c th hive using Equation (2), i.e., $1 \leq c \leq C$.

Method 2 (Advanced). Changes in hive temperature $\Theta(t)$ (e.g., human intervention, heat stress) correspond to changes in the health factor h . Intuitively, the current daily value of h depends on its

previous day's value unless some events affect bees so severely that they cause drastic changes in $\Theta(t)$. To model the change in h , we assume the following:

$$h(d) \approx h(d - 1) \quad (6)$$

Here, d corresponds to the current daily value of h . Now, we formalize this soft constraint in our optimization objective as follows:

$$\boxed{\text{minimize}_{r_c, r_h, h_c(d), \theta_{ideal,c}(d)} L = \sum_{c=1}^C L_{c_1} + \lambda L_{c_2}} \quad (7)$$

$$\text{where } L_{c_1} = \sum_{t=1}^N \|\Theta_c(t) - \hat{\Theta}_c(t)\|^2, \text{ and } L_{c_2} = \sum_{d=1}^{N_d} \|h_c(d) - h_c(d - 1)\|^2$$

Here, L_{c_1} and L_{c_2} are the L2 norm of the reconstruction error and the difference in consecutive h values for a single hive. N_d corresponds to the total number of days in the experimental period. $\lambda (> 0)$ is a half-range-based multiplier to balance the influence of $\Theta(t)$ and h as they both have different scales. We define λ as follows:

$$\lambda = \frac{\text{Half range of } \Theta(t) * N_s * N_d}{\text{Half range of } h * N_d} = \frac{\text{Half range of } \Theta(t) * N_s}{\text{Half range of } h} \quad (8)$$

Here, $N_s = \frac{N}{N_d}$ is the number of samples per day. For λ , we will consider the half range of $\Theta(t)$ for healthy hives [23, 25] that varies in the range of [33, 36] and half range of h that varies in the range of $(0, 1]$ for all hives (details in Section 5). As we use constant values for the half-range of $\Theta(t)$ and h , λ only depends on N_s . We use the same N_s for all c for all experiments, so λ will be a constant multiplier.

To get rid of the dimensionality curse and have a flexible optimizer choice (details in Section 4.3.1) while estimating parameters using EBV_{model+} , we use block coordinate descent optimization technique and formulate it as a bi-level optimization problem:

- At the upper level, we want to calculate the values of r_c and r_h that hold for hives with the same genotype, hive setup, and geographic location (Section 4.2). Here, we consider fixed lower level decision variables (h, θ_{ideal}).
- Only daily parameters, i.e., hive health factor h and ideal core temperature θ_{ideal} , will differ across hives based on their thermoregulation ability. At the lower level, we aim to minimize the reconstruction error of $\Theta(t)$ and the daily change in h for individual hives considering fixed upper-level decision variables (r_c, r_h).

Thus, the problem to solve for Equation (2) will be as follows:

$$\text{minimize}_{r_c, r_h} L = \sum_{c=1}^C L_c \quad (9)$$

$$\text{s.t. minimize}_{h_c(d), \theta_{ideal,c}(d)} L_c = L_{c_1} + \lambda L_{c_2} \quad (10)$$

where $\hat{\Theta}_c(t)$ is subject to Equation (2) and $0 < r_c \wedge r_h \leq \infty$;

$$0 < h_c(d) \leq 1; \quad \theta_{ideal}^{lower} \leq \theta_{ideal,c}(d) \leq \theta_{ideal}^{upper}; \quad d = 1, 2, \dots, D (= \text{no. of days})$$

The constraints of the search space for the optimization problem are discussed in detail in Section 4.4. This optimization objective function, coupled with the advanced model EBV_{model+} , helps address the limitations of the basic model EBV_{model} .




Fig. 3. Bi-level optimizer converges in finite time as opposed to single-level optimizer.

Algorithm 1: $\text{EBV}_{\text{fit+}}$

- 1 $\mathcal{R} \leftarrow \text{EBV}_{\text{gfit+}}(\mathcal{D}_p)$ /*Global Hive Parameter $\{r_c, r_h\}$ */
 - 2 $\{\hat{\Theta}, \mathcal{H}\} \leftarrow \text{EBV}_{\text{hfit+}}(\mathcal{D}, \mathcal{R})$ /*Fixed $\mathcal{R}$ for Hive-Specific Parameters*/
-

4.3.1 Why Bi-Level Optimization? The main reason for proposing a bi-level approach instead of a single-level one is to avoid convergence issues due to the high number of decision variables (dimensionality curse). Suppose we have data from $C (= 5)$ hives for $N_d (= 60)$ days. In such a scenario, we will have a very high-dimensional space (602 dimensions): $300 = (5 * 60)$ for $h + 300 = (5 * 60)$ for $\theta_{ideal} + 2$ for (r_c, r_h) . Any non-linear optimizer will either get stuck at the local minima (e.g., **Levenberg-Marquardt (LM)**) or take forever to converge (e.g., Nelder-Mead) due to the very high-dimensional space. Dividing into sub-problems with lower dimensional space using bi-level optimization helps with this issue.

A secondary reason is to have the flexibility to customize the optimizer for different levels. As mentioned in Section 4.1.1, we have used Nelder-Mead for the upper level and LM for the lower level optimization. While the LM optimizer works well for lower level optimization, it gets stuck to the initial condition (local minima) for upper level optimization. Nelder-Mead's approach to exploring using a simplex helps with the issue as it can explore the space more broadly. We can also use Nelder-Mead for the lower level optimization, but it will be comparatively slow. For our data, the bi-level optimizer converges to the optimal solution (r_c, r_h, h) in 10 hours with a reconstruction error (**Root Mean Square Error (RMSE)**) of 0.61 (Figure 3) as opposed to any non-linear single-level optimizer that takes infinite time to find a solution. Here, we use only data from the control hives (details in Section 5).

4.4 Reconstruction Algorithm: $\text{EBV}_{\text{fit+}}$

How can we minimize L to reconstruct $\Theta(t)$ and estimate the optimum global (r_c, r_h) as well as hive-specific (h, θ_{ideal}) parameters to be used for both segmentation and proactive warning?

Overview. Algorithm 1 outlines our proposed reconstruction algorithm, $\text{EBV}_{\text{fit+}}$. At first, the algorithm calls $\text{EBV}_{\text{gfit+}}$ (Section 4.4.1), which takes only a subset of the entire recorded dataset $\mathcal{D}$, $\mathcal{D}_p$, to estimate the global hive parameters $\{r_c, r_h\}$ (Line 1). We selected $\mathcal{D}_p$ such that it captures global hive parameters in a regular hive condition. To ensure that we use only the data from control hives that received no ice treatment in our experiments. The reason is that we assume global parameters remain unaffected by external treatments (e.g., ice treatment, pesticide treatment) if bee genotype, hive structure, and geographic location remain the same. So we can use those parameters to estimate $\{h, \theta_{ideal}\}$ for both treated and untreated hives. Next, Algorithm 1 uses $\text{EBV}_{\text{hfit+}}$ (Section 4.4.2) to obtain the set of hive-specific daily parameters $\{h, \theta_{ideal}\}$ (Line 2), which is required for the offline segmentation on the recorded dataset $\mathcal{D}$ (Section 4.5) and the online proactive notification on the daily streaming data $\mathcal{D}_s$ (Section 4.6). We use fixed $\mathcal{R}$ estimated using $\text{EBV}_{\text{gfit+}}$ for this step as $\mathcal{R}$

Algorithm 2: $EBV_{\text{gfit+}}$

Input: $\mathcal{D}_p = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_p$
Output: $\mathcal{R} = \{r_c, r_h\}$

```

1  $\mathcal{R} \leftarrow \mathcal{R}_{INIT}, \mathcal{H}_c \leftarrow \mathcal{H}_{INIT}$  /*initialization;  $\mathcal{H}_c = \{h, \theta_{ideal}\}_c$  for all days*/
2 /*Lower Level Optimization: Fixed  $\mathcal{R}$  from UpL */
3 Function LowL( $\mathcal{D}_p, \mathcal{R}$ ):
4   for each hive  $c \in C_p$  do
5     while not converged do
6       Solve  $\hat{\Theta}_c$  for  $\{\mathcal{R}, \mathcal{H}_c\}$  using Eqn. (2) //  $\mathcal{R}$  from UpL &  $\mathcal{H}_c$  from LowL's prev. iter.
7       Calculate  $L_c$  using Eqn. (10) //  $L_c = L_{c1} + \lambda L_{c2}$ 
8       Update  $\mathcal{H}_c$  using Non-linear Optimizer
9      $\{L, \hat{\Theta}, \mathcal{H}\}_p \leftarrow \{L, \hat{\Theta}, \mathcal{H}\}_p \cup \{L, \hat{\Theta}, \mathcal{H}\}_c$ 
10  return  $\{L, \hat{\Theta}, \mathcal{H}\}_p$ 
11 /*Upper Level Optimization: Fixed  $\mathcal{H}_p$  from LowL */
12 Function UpL( $\mathcal{D}_p$ ):
13   while not converged do
14      $L_{p, \leftarrow} = \text{LowL}(\mathcal{D}_p, \mathcal{R})$  //  $\mathcal{R}$  from UpL's prev. iter.
15     Calculate  $L$  using Eqn. (9) //  $L = \sum_{p=1}^C L_p$ 
16     Update  $\mathcal{R}$  using Non-linear Optimizer
17     /*Early Stopping:  $\delta = 0.01$ */
18     if  $\Delta \mathcal{R} < \delta$  for  $e$  epochs then
19        $\mathcal{R} \leftarrow \mathcal{R}_{prev}$ 
20       break
21  return  $\mathcal{R}$ 

```

estimated for a subset of hives is applicable for other hives of the same specification as well. Below, we explain the details of the two sub-algorithms, $EBV_{\text{gfit+}}$ and $EBV_{\text{hfit+}}$.

4.4.1 Global Hive Parameter Estimation: $EBV_{\text{gfit+}}$. Algorithm 2 ($EBV_{\text{gfit+}}$) takes $\mathcal{D}_p = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_p$ as an input and returns $\mathcal{R} = \{r_c, r_h\}$ as an output:

- $\mathcal{D}_p = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_p \subset \mathcal{D} = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}$ where $\mathcal{D}$ is the data from all experimental hives. Here, $\Theta_{\mathcal{E}} = \{\Theta_{\mathcal{E}_1}, \Theta_{\mathcal{E}_2}, \dots, \Theta_{\mathcal{E}_C}\}$, $\Theta_{\mathcal{P}} = \{\Theta_{\mathcal{P}_1}, \Theta_{\mathcal{P}_2}, \dots, \Theta_{\mathcal{P}_C}\}$, and $\Theta = \{\Theta_1, \Theta_2, \dots, \Theta_C\}$ are sets of external, peripheral, and core temperature sequences for all hives over the experiment period, and $\Theta_{\mathcal{E}_{p,c}}$, $\Theta_{\mathcal{P}_{p,c}}$, and $\Theta_{p,c}$ are the sets of daily recorded temperature sequences from the c th hive; i.e., $c \leq C$, C is the total number of experimental hives.
- $\mathcal{R} = \{r_c, r_h\}$ is the set of global hive parameters (maximum thermoregulation capacity limit) for hives of the same specification.

As mentioned in Section 4.3, we formulate our problem as bi-level optimization and use the following notation. Random initialization is used for initializing all parameters used in optimization (Line 1). This step uses a subset of the entire dataset to avoid over-fitting. We also discard the days when $\Theta(t) = \theta_{ideal}$ throughout the day for convergence. While we have used two different optimizers for both steps, any other non-linear optimizer can be used. Here, we keep upper level decision variables fixed while optimizing the lower level and vice versa until the early stopping criteria for (r_c, r_h) is met. We give a detailed explanation below:

Lower Level Optimization. (Lines 2–10) The function $\text{LowL}(\mathcal{D}_p, \mathcal{R})$ takes the subset of the entire dataset $\mathcal{D}_p = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_p$ and updated $\mathcal{R}$ from the function $\text{UpL}(\mathcal{D}_p)$ as inputs and returns $\{L, \hat{\Theta}, \mathcal{H}\}_p$ as output. $\{L, \hat{\Theta}, \mathcal{H}\}_p$ are the sets of minimum loss, reconstructed hive temperature sequence, and hive-specific daily parameters for the fixed input $\mathcal{R}$ value from $\text{UpL}(\mathcal{D}_p)$.

Here, we consider each hive individually. For the c th hive, we want to minimize the reconstruction error, $\Theta_{r_{\mathcal{E}}}(t) = \Theta_c(t) - \hat{\Theta}_c(t)$, and consider the intuition that the daily change in health factor, h depends on the previous day's h by minimizing the loss, L_c . Here, $\Theta_c(t)$ and $\hat{\Theta}_c(t)$ are the actual and reconstructed core temperatures for the c th hive. $\mathcal{H}_c = \{H_{c_1}, H_{c_2}, \dots, H_{c_n}\}$ is the set of estimated daily parameters to achieve L_c . Here, $H_{c_i} = \{h, \theta_{ideal}\}_{c_i}$ is the set of health factor and ideal temperature on the i th day, i.e., $i \leq n$. We use the LM optimization for faster convergence considering the following constraints for the search space $\mathcal{H}_s = \{h_s, \theta_s\}$ (Line 3).

- (1) $h_s = [h_{lower}, h_{upper}] = (0, 1]$ is the search space for bees' health factor h . Here, we consider $h_{lower} = 0.1 (> 0)$; chosen experimentally.
- (2) $\theta_s = [\theta_{ideal}^{lower}, \theta_{ideal}^{upper}]$ is the search space for θ_{ideal} that can be anywhere between $33 - 36^\circ\text{C}$ for healthy hives and go outside this range for unhealthy ones. We assumed the range to be $(33 - 36) \pm 2^\circ\text{C}$.

We first find values of $\hat{\Theta}_c(t)$ for fixed $\mathcal{R}$ from UpL and updated $\mathcal{H}_c$ from LowL's previous iteration (Line 6). Next, we calculate the loss, L_c for $\hat{\Theta}_c(t)$ and $\mathcal{H}_c$ using Equation (10) and update $\mathcal{H}_c$ using the LM optimizer (Lines 7–8). We repeat these three steps until we find such $\mathcal{H}_c$ that gives the minimum L_c . We keep the variables from UpL block fixed for these steps.

Upper-Level Optimization. (Lines 11–21) The function $\text{UpL}(\mathcal{D}_p)$ takes the data of subset hives, $\mathcal{D}_p = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_p$, as input and returns the global hive parameters, $\mathcal{R} = \{r_c, r_h\}$, as output.

Here, we consider that bees of the same genotype, geographical region, and hive setup will have the same upper thermoregulation limit, $\mathcal{R}$, to maintain $\Theta(t)$. We want to find a $\mathcal{R}$ value that will minimize the overall loss, L , for the input hives. We use fixed hive-specific parameters, $\mathcal{H}_p$ from the lower level, LowL and $\mathcal{R}$ from UpL's previous iteration to calculate loss, L_p for individual (Line 14). Then, we update $\mathcal{R}$ based on calculated total loss, L for all hives using Equation (9) (Lines 15–16). We repeat these steps until the early stopping criteria is met for optimal $\mathcal{R}$. We use the Nelder–Mead optimization considering the following constraints for the search space, $\mathcal{R}_s = \{\mathcal{R}_C, \mathcal{R}_H\}$ (Line 1): $\mathcal{R}_C \wedge \mathcal{R}_H = [0, r_{\infty}] = [0, \infty]$ are the search space for bees' maximum cooling limit, r_c , and heating limit, r_h .

Early Stopping. (Lines 17–21) To reduce the computation time and achieve faster convergence, we add the early-stopping constraint to the upper level optimization:

“If the change in the global hive parameter, $\mathcal{R}$, is less than $\delta(\rightarrow 0^+)$ for e epochs, we use the value of $\mathcal{R}$ from the previous epoch, $\mathcal{R}_{prev}$, as the optimal value.”

4.4.2 Hive-Specific Parameter Estimation: $\text{EBV}_{\text{hfit+}}$. Algorithm 3 ($\text{EBV}_{\text{hfit+}}$) takes the global hive parameter $\mathcal{R} = \{r_c, r_h\}$ (estimated using $\text{EBV}_{\text{gfit+}}$) and a dataset $\mathcal{D} = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}$ (either the recorded dataset or streaming data) as inputs and returns $\{\hat{\Theta}, \mathcal{H}\}$, which includes the set of reconstructed sequences $\hat{\Theta}$ and hive-specific daily parameters $\{h, \theta_{ideal}\}$ for the input hives as output. We use random initialization for $\mathcal{H}$ (Line 3). We aim to find such value of $\mathcal{H}$ that best reconstructs the hive temperature, Θ for the input data sequence and minimizes the distance between consecutive h values. We follow the same steps (Lines 3–5) as LowL of $\text{EBV}_{\text{gfit+}}$ and repeat until we find minimum L_c . We used the LM optimizer for this part, while other non-linear optimizers can also be used. As we have the estimated $\{r_c, r_h\}$ using $\text{EBV}_{\text{gfit+}}$ and the current value of h depends on its previous day's value, we can estimate the h in case of $\Theta(t) = \theta_{ideal}$.

Algorithm 3: EBV_{hfit+}

Input: $\mathcal{R}, \mathcal{D} = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}$ //Fixed $\mathcal{R}$ from EBV_{gfit+}; $\mathcal{D}$ is recorded/streaming data
Output: $\{\hat{\Theta}, \mathcal{H}\}$

```

1  $\mathcal{H} \leftarrow \mathcal{H}_{INIT}$  /*initialization*/
2 while not converged do
3   Solve  $\hat{\Theta}$  for  $\mathcal{R}$  &  $\mathcal{H}$  using Eqn. (2) //  $\mathcal{R}$  from EBVgfit+ and  $\mathcal{H}$  from EBVhfit+'s prev. iter.
4   Calculate  $L_c$  using Eqn. (10) //  $L_c = L_{c_1} + \lambda L_{c_2}$ 
5   Update  $\mathcal{H}$  using Non-linear Optimizer

```

4.5 Segmentation Algorithm: EBV_{fit&cut+}

Overview. Based on EBV_{model+}, we develop a segmentation algorithm to find out a set of segments and their parameters by penalizing model complexity. The cut-points spot discontinuities corresponding to events affecting bees' thermoregulation ability. We use AIC to measure the goodness of fit of our algorithm. We aim to minimize the following equation:

$$AIC = -2\ln\mathcal{L}(\mathcal{D}|\mathcal{M}) + 2(m + (m + 1)p) \quad (11)$$

$$\text{where } \mathcal{L}(\mathcal{D}|\mathcal{M}) = \prod_{i=1}^N f(\Theta[i] - \hat{\Theta}[i])$$

Here, $\mathcal{L}(\mathcal{D}|\mathcal{M})$ is the likelihood of data $\mathcal{D}$ given the model $\mathcal{M}$ described in Equation (2) and m, p ($= 2$) represent the number of cuts and parameters for each segment, respectively. Also, $f(\cdot)$ is the pdf of a continuous normal random variable, and N is the total time ticks.

We want to find the combination of segments and corresponding parameters, $\{h, \theta_{ideal}\}$, that minimizes Equation (9) and outputs the best-fitted core temperature, $\hat{\Theta}(t)$. Algorithm 4, EBV_{fit&cut+}, gives the implementation of EBV_{model+} for curve fitting and segmentation. It takes $\mathcal{D}_c = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_c$ as input and $\mathcal{CP}_c = \{cp_1, cp_2, \dots, cp_m\}_c$, $\mathcal{X}_c = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_{m+1}\}_c$ as outputs:

- $\mathcal{D}_c = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_c$ is the recorded temperature dataset for the c th hive, i.e., $c \leq C$, where C is the total number of experimental hives.
- $\mathcal{CP}_c = \{cp_1, cp_2, \dots, cp_m\}_c$ is the set of cut-point locations, and $\mathcal{X}_c = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_{m+1}\}_c$ is the set of $\{h, \theta_{ideal}\}_c$ for the c th hive, where m is the number of cut points and n is the number of daily sequences, i.e., $m \leq n$.

For this part, we use the following notations.

- $\mathcal{L} = \{l_1, l_2, \dots, l_n\}$ is the set of likelihoods for different cut-points positions, and $\mathcal{L}_{min}$ is the minimum of $\mathcal{L}$.
- $\mathcal{P}_{RMSE_{min}}$ and $\mathcal{P}_{\mathcal{L}_{min}}$ are the parameters $\mathcal{X}_c = \{h, \theta_{ideal}\}_c$ for each segment, which give the minimum RMSE and AIC for $\mathcal{D}_c$.

For segmentation, we follow a greedy approach. Let us begin with the simplest case: a single cut-point search with two segments (i.e., $\mathcal{S}_{c_1}$ and $\mathcal{S}_{c_2}$) for a given dataset $\mathcal{D}_c$. Each segment will have its respective parameters, $\mathcal{X}_{c_i} = \{h, \theta_{ideal}\}_{c_i}$, for best reconstruction. We start with the cut-point position, $cp = 1$, and iterate till n . For each cut-point, we will have $\{(\mathcal{S}_{c_1}, \mathcal{X}_{c_1}), (\mathcal{S}_{c_2}, \mathcal{X}_{c_2})\}$ and $\mathcal{L}$ for $\mathcal{D}_c$. We will choose $\mathcal{CP}_c = \{cp_1\}_c$ with $\mathcal{L}_{min}$. How about a second one? We will repeat the above process, excluding cp_1 from our search. Then, we will have $\mathcal{CP}_c = \{cp_1, cp_2\}_c$ and repeat the same process for the next one till we find reconstruction with the minimum AIC.

Reconstruction using more segments will be closer to the original reconstruction: we get the best fit if each sequence is reconstructed independent of the other, i.e., n sets of $\mathcal{X}_c$ for n sequences.

Algorithm 4: $EBV_{\text{fit\&cut+}}$

Input: $\mathcal{D}_c = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_c$
Output: $\{\mathcal{CP}_c, \mathcal{X}_c\}$

```

1  $j \leftarrow 1, \mathcal{CP}_c \leftarrow \{0, n\}, \mathcal{X}_c \leftarrow \emptyset$ 
2 while  $j < n$  do
3   for each point,  $p \in \{\mathcal{D}_c - \mathcal{CP}_c\}$  do
4      $cp \leftarrow \{\mathcal{CP}_c\} \cup p$  //set of cut-points
5     for  $i = 0$  to  $\text{length}(\mathcal{CP}_c)$  do
6       Create Segment  $s[i] \leftarrow \mathcal{D}_c[cp[i] : cp[i + 1]]$ 
7       Fit and choose  $\mathcal{P}_{\text{best}} \leftarrow \mathcal{P}_{\text{RMSE}_{\min}}$ 
8       Calc.  $\mathcal{L}(\Theta_c(t), \hat{\Theta}_c(t))$  /*Eqn. (11)*/
9   Choose  $\mathcal{L}_{\min}(\Theta_c, \hat{\Theta}_c(t))$ , and  $\mathcal{P}_{\mathcal{L}_{\min}}$ 
10   $\{\mathcal{CP}_c, \mathcal{X}_c\} \leftarrow \{\{\mathcal{CP}_c\} \cup p_{\mathcal{L}_{\min}}, \{\mathcal{X}_c\} \cup \mathcal{P}_{\mathcal{L}_{\min}}\}$ 
11   $j \leftarrow j + 1$ 

```

But more segments also increase the chances of detecting cuts with no significant changes in $\mathcal{X}_c$. In segmentation, we aim to reduce the model complexity by reducing the number of $\mathcal{X}_c$ s. Thus, instead of considering the minimum AIC, we use the combination of segments and parameters that gives 10% higher AIC than the minimum one [37]. The experimental results in Section 5.1 demonstrate that our assumption is reasonable.

4.6 Warning Algorithm: $EBV_{\text{fcast\&warn+}}$

How do we proactively detect changes that affect bees' thermoregulation ability and hive health and alert beekeepers in real time for timely action?

Overview. In this section, we develop an algorithm, $EBV_{\text{fcast\&warn+}}$, that can issue proactive warnings if there is any change in the bees' thermoregulation ability or if the health factor h becomes too low based on the forecasting using $EBV_{\text{model+}}$. Algorithm 5 demonstrates the implementation of $EBV_{\text{fcast\&warn+}}$ for issuing a proactive warning. It takes $\mathcal{D}_{i_c}[P_{-M} : P] = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_{i_c}$ and $\mathcal{D}_{s_c} = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_{s_c}$ as inputs and returns whether to be silent or raise a flag or warning:

- $\mathcal{D}_{i_c}[P_{-M} : P] = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_{i_c}$ is the recorded temperature sequences for the past M days from the present day, P , for the c th hive.
- $\mathcal{D}_{s_c} = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_{s_c}$ is the next day's streaming temperature sequence for the c th hive.
- If a *flag* is raised, it corresponds to a significant degradation in forecasting accuracy: a negative change in h . On the other hand, if a *warning* is raised, it corresponds to weaker bees and very low h : lower than a threshold. There can be silence after the flag/warning is lifted.

Forecasting (One Day Ahead). (Lines 1–4) Given the past M days of recorded temperature for the c th hive, $\mathcal{D}_{i_c}[P_{-M} : P]$, how do we forecast the next day's core temperature, $\hat{\Theta}_{f_c}(t)$?

We first calculate the hive-specific daily parameters, $\mathcal{H}_{i_c} = \{h, \theta_{\text{ideal}}\}_{i_c}$, and the reconstructed sequence, $\hat{\Theta}_{i_c}$, using the estimated $\mathcal{R}$ from $EBV_{\text{fit+}}$ (Algorithm 1). Then, we forecast $\hat{\Theta}_{f_c}$ for the next day assuming the hive-specific parameters remain the same in the next day, i.e., $\mathcal{H}_{f_c} = \mathcal{H}_{i_c}[P]$, and using $EBV_{\text{model+}}$ (Equation (2)).

Anomaly Detection. (Lines 5–8) Given the next day's hive temperature forecast, $\hat{\Theta}_{f_c}$, now we want to detect changes in the thermoregulation pattern in real time to warn beekeepers proactively.

First, let us define what an anomaly means in our context. While forecasting, we assumed $\mathcal{H}_{f_c} = \mathcal{H}_{i_c}[P]$ to forecast $\hat{\Theta}_{f_c}$ for the next day. If some factor affects the bees' thermoregulation

Algorithm 5: EBV_{fcast&warn+}

Input: (1) $\mathcal{D}_{i_c}[P_{-M} : P] = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_{i_c}$ (2) $\mathcal{D}_{s_c} = \{\Theta_{\mathcal{E}}, \Theta_{\mathcal{P}}, \Theta\}_{s_c}$
Output: Flag/Warning/Silence

```

1 /*One day ahead Forecast:  $\hat{\Theta}_{f_c}$ */
2  $\{\hat{\Theta}_{i_c}, \mathcal{H}_{i_c}\} \leftarrow EBV_{\text{hfit+}}(\mathcal{S}, \mathcal{D}_{i_c}[P_{-M} : P])$ 
3  $\mathcal{H}_{f_c} \leftarrow \mathcal{H}_{i_c}[P]$ 
4 Solve  $\hat{\Theta}_{f_c}$  using Eqn. 2
5 /*Anomaly Detection*/
6 for sliding window,  $w$  do
7   Check  $(|\hat{\Theta}_{f_c, w} - \Theta_{s, w}|) > s * \sigma_{\Theta_{AE}}$ 
8   if True for  $\frac{w}{k}$  consecutive pts then  $P.T. = \text{True}$  /*P.T. = Persistence Test*/
9 /*Proactive Notification*/
10 while True do
11   if  $h_{f_c} \leq \beta$  for  $\geq x$  days then
12     Raise a warning /*bees are too weak*/
13   else
14     Calculate  $h_{s_c}$  using Eqn. 2
15     if  $(h_{s_c} - h_{i_c}[P]) < 0) \wedge (P.T. = \text{True})$  for  $\geq x$  days then
16       Raise a flag /*strength decreased*/
17   if flag/warning discontinues then
18     Silent flag/warning for  $y$  days
19     break

```

ability and our assumption does not hold, the deviation between the forecasted and actual sequences will be large, and we will consider that sequence anomalous. Intuitively, the larger the deviation of the forecasted data sequence from the actual one, the more anomalous that sequence is. To detect anomalies, we introduce a test named *Persistence* test based on the deviation of the forecasting from the actual value of $\Theta(t)$:

“For a sliding forecasting window size of $w(\leq N_s)$, if $\frac{w}{k}$ consecutive points fall out of s standard deviations of the reconstruction error of the input sequences used for forecasting, the forecasted sequence is considered anomalous.”

While we can use any other threshold, we follow some common practices to detect an anomaly:

- We use the standard deviation of the absolute residuals of the past input sequences, $\Theta_{AE} = |\hat{\Theta}_{i_c} - \Theta_{i_c}|$, as an estimate of the deviation of the future residuals. Statistically, two standard deviations of the reconstructed residual of the input sequence represent 95% of the residual data from the reconstructed sequence. Thus, it is safe to assume any forecasted point as anomalous if it falls out of this range.
- We choose the window size $w = \frac{N_s}{2}$ experimentally but any value within the range of $[N_s/3, N_s]$ can also be used, where N_s is the number of samples per day. Here, we used a sliding window to detect changes in bee health, h , as early as possible.
- If a few consecutive forecasted points fall out of the above range, we can consider the whole sequence anomalous. Based on our experiments, we choose the number of consecutive points $\frac{w}{4}$ ($= 25\%$ of w).

Proactive Notification. (Lines 9–19) Now, we use the forecasted $\hat{\Theta}_{fc}$ and the *Persistence* test to raise a flag/warning in real time:

- *Warning.* If the forecasted value of h_{fc} falls below a threshold β for x consecutive days, it indicates very weak bees. We issue a warning to alert beekeepers to take action as soon as possible. We define the threshold β based on the lower limit of the search space, h_{min} for estimating h .
- *Flag.* An anomalous forecasted sequence detected by the *Persistence* test indicates a significant change in thermoregulation ability, Δh_{fc} , in bees. If *Persistence* holds *True* and the change $\Delta h_{fc} < 0$ for x consecutive days, we raise a flag to inform beekeepers of the change in h .

Based on domain experts' opinion, we set $x = 2$ days: if both conditions are *True* for two consecutive days, alert beekeepers. For real-time notification, we consider the following factors:

- When a flag/warning discontinues, there will be a silence period of y days. We do not perform any persistence test during that period or check if h_{fc} is too low. We assume that beekeepers will take preventive action based on the issued flag/warning, and it takes some time for bees to adapt to the change from the action taken. Based on experts' opinions, we assume $y = 3$ days.
- A warning is more severe than a flag: h is too low. If a warning is issued, we will not issue any flag (negative change in h_{sc}) until that warning discontinues and the silence period is over.
- If we encounter missing data points during an ongoing warning due to the nature of real-world data collection, we continue issuing the warning. The reason is we assume $h_{fc} = h_{ic}[P]$; even without next-day data, we can treat the next day's h_{fc} being the same as the present-day $h_{ic}[P]$.
- If there are any missing data points during an ongoing flag, we discontinue the flag and start the silence period. We cannot perform the persistence test if we do not have streaming $\Theta(t)$.

5 Empirical Evaluation

This section presents the effectiveness of our proposed method, EBV+, using our collected datasets and experiments designed to ensure the following:

- Q1 *Effectiveness:* How accurately does our proposed method EBV+ detect events affecting hive health (offline and in real time) and forecast the core temperature, $\Theta(t)$? (Sections 5.1, 5.2 and 5.3)
- Q2 *Explainability:* Can EBV+ explain the change in hive health and the effect of corresponding events in detected segments? (Section 5.4)
- Q3 *Scalability:* Does our model perform linearly in time? (Section 5.5)
- Q4 *Observations:* Do the observations based on our experiments coincide with the domain experts' expectations? (Section 5.6)

Experimental Setup. All experiments are done on a Dell XPS 15 laptop with an Intel i9 processor and 32 GB memory. We used averaged hourly data for all experiments as the change in temperature is noticeable if considered hourly. There are missing data points due to the nature of real-world data collection, e.g., sensors not working properly, beekeepers' intervention for inspection. However, no interpolation is done since EBV+ is robust to handling missing values. While estimating the parameters, we use the following specifications:

- For Algorithm 2, $EBV_{\text{fit}+}$, we only used data from control hives that received no ice treatment in our experiments. We assume that all hives of the same hive setup, bee genotype, and geographic location have the global parameter, $\mathcal{R}$. So, the $\mathcal{R}$ calculated for control hives will also work for treated hives. The results of all our experiments support this assumption to be true.

—To solve Equation (7), we need the multiplier λ , which we calculate using Equation 8 as follows:

$$\lambda = \frac{\text{Half range of } \Theta(t) * N_s}{\text{Half range of } h} = \frac{2 * 24}{0.5} = 96$$

The values are determined as follows: (i) As per expert opinion, $\Theta(t)$ varies in the [33, 36] range for healthy hives [23, 25]. So we use, half range of $\Theta(t) = \frac{\Theta_{\max}(t) - \Theta_{\min}(t)}{2} = \frac{36 - 33}{2} = 1.5 \approx 2$; (ii) Using averaged hourly data, we will have $N_s = 24$ per day, and (iii) h varies in the [0.1, 1] range, so we use, half range of $h = \frac{h_{\max} - h_{\min}}{2} = \frac{1 - 0.1}{2} = 0.45 \approx 0.5$.

Using all these specifications, we estimate the value of $\mathcal{R} = \{r_c, r_h\} = \{80, 221\}$ which we use to estimate the value of h for all our experimental hives. Due to space constraints, we present one example hive from each hive setting in this section, but our findings apply to all hives used during the experiment.

Validation Methods. We employed multiple standard validation techniques to evaluate the performance and reliability of our proposed method. Our primary metrics for model validation are RMSE and AIC (details in Section 5.2) for forecasting $\Theta(t)$ and accuracy and precision (details in Section 5.3) for proactive notification followed by unexpected events affecting bees' thermoregulation ability. RMSE provides a direct measure of our model's accuracy and allows straightforward comparisons across different methods and hive setups (Figure 5). AIC measures evaluate the goodness of fit while considering model complexity (Figure 1(c)). Our validation approach involves using fixed time windows for forecasting the hive temperature, $\Theta(t)$. Specifically, we use 7 days of temperature sequences to forecast the next 14 days of hive temperature. This approach allows us to assess our model's performance over extended periods by capturing both short-term fluctuations and long-term trends. We apply this method across various hive setups (control and treated) to ensure robust validation of our model's capabilities. The following subsections detail our experimental conditions and results.

5.1 Q1—Effectiveness: Event-Detection

Experimental Setup. Frames were taken out of hives twice in 2 days during every inspection to recharge and replace sensors installed inside (Section 3). It takes more time than usual⁴ and thus induces more stress to bees and larvae. We refer to this type of inspection as intensive compared to regular ones. All hives were opened on the same days: 17–18 August, 1–2 September, and 15–16 September (marked by open red boxes in figures: *intensive inspection*). For this part, we also looked at the treated hives that received ice treatments during hot days (marked by blue boxes with ice symbols, e.g., Figure 4). We want to see if $\text{EBV}_{\text{fit\&cut+}}$ is able to detect changes in $\Theta(t)$ following the inspections.

- (1) *Intensive hive inspections:* For maximal hive performance, bees need to remove larvae affected during inspection and reestablish homeostasis. Maintaining thermoregulation and typical colony activities simultaneously for stressed hives can lead to a noticeable fluctuation in $\Theta(t)$ within 9–10 days following the inspection.
- (2) *Effect of ice treatment:* Domain experts expect ice treatment during hot days will help reduce stress and support thermoregulation following intensive inspections in the long run.

⁴Normal hive inspection takes around 10 minutes and is done once for 2–3 weeks. Our inspection took 15–20 minutes for each hive as we removed the frame with the sensor device (and brood) installed to take out the device for recharge. We opened the hives and removed the same frames the next day to reinstall the recharged sensors. Removing brood frames from the hive and exposing them to the environment stresses the honeybee colony and can cause dead larvae.




Fig. 4. *EBV+ is effective and explainable*: accurate segments: $EBV_{fit\&cut+}$ fits (red line) the data (blue line) well and finds cut-points (brown vertical line) followed by unexpected events (red boxes: intensive hive inspection, sensor recharge and replace). Notice that the baselines are not applicable (cannot do segmentation). (a) Control hive (high heat stress). (b) Treated hive (low heat stress).

Results. $EBV_{fit\&cut+}$ detected $n = 2$ and $n = 1$ cut-points and corresponding parameters for each segment (brown vertical line in Figure 4) for the control and treated hives, respectively. The reconstructed curve (red line) by $EBV_{fit\&cut+}$ is well aligned with the recorded $\Theta(t)$ (blue line).

For the control hive used as an example here (Figure 4(a)), $EBV_{fit\&cut+}$ detected cut-points on 28 August and 9 September: 11 and 8 days following the first and second inspection, respectively. We suspect that it is correlated with hive cleaning (such as removing affected larvae, repairing damaged combs, and/or resealing of the brood chamber with wax and/or propolis) and re-establishing homeostasis after intensive inspection. No cut-point was detected after the inspections on 15–16 September, possibly due to the recorded data length after the inspection being too short to detect discontinuities.

For the treated hive, $EBV_{fit\&cut+}$ detected a cut-point (Figure 4(b)) on 26 August: 9 days following the first inspection. No cut-points were detected after the next two inspections. This indicates that bees in the treated hive were good at managing the stress resulting from intensive inspections. We strongly suspect that this is a result of the positive cumulative effect of ice treatment over time.

In summary, our algorithm $EBV_{\text{fit\&cut+}}$ detected a total of 9 cut-points for 15 inspections (3 per hive) of control hives and 5 cut-points for 15 inspections of treated ones. There were no false positives, and the detected events were linked to life history occurrences (Section 5.4).

5.2 Q1—Effectiveness: Forecasting

Experimental Setup. For each trial, we use 7 days of temperature sequences (endogenous variable: $\Theta(t)$, and exogenous variable: $\Theta_{ext}(t)$, $\Theta_{adj}(t)$) as inputs and forecast $\Theta(t)$ hourly for the next 14 days, given $\Theta_{ext}(t)$ for the forecasting period. The forecasted core temperature is compared against the actual core temperature using two metrics: (i) RMSE = $\sqrt{\frac{\sum_{i=1}^K (\Theta[i] - \hat{\Theta}[i])^2}{K}}$, where K is the number of forecasted time-ticks and $\hat{\Theta}[i]$ is the forecasted value at i th time-tick and (ii) AIC = $-2\ln\mathcal{L}(\mathcal{D}|\mathcal{M}) + 2P$, where $\mathcal{L}(\mathcal{D}|\mathcal{M})$ is the likelihood of data $\mathcal{D}$ given a model $\mathcal{M}$, chosen either from the baselines (described below) or EBV+ and P is the number of parameters for that model. RMSE measures the forecasting error, whereas AIC measures the goodness of fit in terms of forecasting accuracy and model complexity. The lower the value of RMSE and AIC, the better the model is.

Baselines. To demonstrate the effectiveness of EBV+ in forecasting, we compared it to the state-of-the-art time series forecasting models ARX, seasonal ARX [16], Holt-Winters [39], DeepAR from Autogluon [10], and our basic method EBV. We followed standard practice to determine (i) the order (= 1) of ARX and seasonal ARX using AIC [9], and (ii) the hyperparameters of the Holt-Winters method using additive trend and seasonality with non-linear optimization. For DeepAR, we used 2 RNN layers with 40 cells for each and a learning rate of 0.001. As Autogluon (DeepAR) requires the input sequence length to be greater than twice the length of the forecasting period, we recursively forecast 14 days ahead for 7 days of input data.

For EBV, we use estimated parameters (s_c, s_h, θ_{ideal}) from input sequences along with $\Theta_{ext}(t)$ for the next 14 days to forecast $\Theta(t)$ for that time period. For EBV+, we estimate the parameters (h, θ_{ideal}) from input sequences and estimated values of (r_c, r_h) to forecast $\Theta(t)$ for the next 14 days.

To calculate AIC, we determine the number of parameters for all methods except DeepAR for each trial: EBV+ = 2 (h, θ_{ideal}); EBV = 3 (s_c, s_h, θ_{ideal}); ARX(1) = 5 (regressors' coefficients = $\{a_0, a_1, b_0, b_1\}$, variance = σ^2); seasonal ARX(1) = 6 (regressors' coefficients = $\{a_0, a_1, b_0, b_1\}$, seasonal coefficient = s_{24} , variance = σ^2), Holt-Winters = 29 (smoothing coefficients (level, trend, seasonality) = $\{\alpha, \beta, \gamma\}$, addition coefficients (level, trend, seasonality) = $\{l, t, s_1, s_2, \dots, s_{24}\}$). For seasonal ARX and Holt-Winters, we consider that hive core temperature $\Theta(t)$ follows daily seasonality. For DeepAR, we have a fixed number of parameters independent of the number of trials (≈ 0.1 million).

Results. Figure 5 demonstrates examples of the 14-day-ahead forecasting of $\Theta(t)$ (11–24 August 2021) using 7 days' temperature sequences (3–10 August 2021) as inputs. The magenta, black, and cyan lines represent the input $\Theta(t)$, $\Theta_{ext}(t)$, and $\Theta_{peri}(t)$ (for the treated hive), respectively. As can be seen, the forecasted core temperature sequence (dashed red line) is very close to the actual one (blue line), demonstrating the effectiveness of our method.

Accuracy. Figure 6 shows the number of parameters vs. average forecasting RMSE of all sequences from control (right panel) and treated (left panel) hives for all methods. EBV+ shows an average improvement of up to 72% over baselines in terms of accuracy and has up to 600 times fewer parameters than baselines. The average AICs for all experimental hives are also reported in Figure 1(c). Each point corresponds to a (hive-id, forecasting method) pair. We show that almost all points fall in the upper left region of the red line (average AICs of EBV+), indicating better forecasting AIC of EBV+ than baselines.




Fig. 5. *EBV+ forecasts $\Theta(t)$ well*: The 14 days ahead forecasting (dotted red line) is close to the actual core temperature (solid blue line). (a) Control hive. (b) Treated hive.




Fig. 6. *EBV+ wins against baseline in terms of core temperature forecasting accuracy (up to 72% improvement) with lesser parameters (up to 600 times)*. Error bars show 1 standard deviation.

EBV also does good forecasting (second-best), but our advanced method, EBV+, has an additional advantage over any other method with fewer parameters: the parameters of baselines (except EBV) are not explainable, whereas EBV+ can better explain bee health status with h . Also, the intuition of EBV+ conforms to domain experts' expectations: increased fluctuation in core temperature $\Theta(t)$ indicates decreased bee health (h) to maintain $\Theta(t)$ (Figure 7).

5.3 Q1—Effectiveness: Proactive Notification

Experimental Setup. For each trial for a single hive, we use 7 days of temperature sequences ($\Theta_{ext}(t)$, $\Theta_{peri}(t)$, $\Theta(t)$) as inputs in a sliding window manner and shift the input window by 1 day for the next trial. Then, we forecast $\Theta(t)$ hourly for the next day, assuming that we have $\Theta_{ext}(t)$ for that day. While performing this experiment, we consider the following:

- To perform the *Persistence* test, we assume to have incoming data per hour, and we do the test every hour based on $\Theta(t)$ from past 12 hours (window size, $w = 12$). We determine whether the sequence used for the test is anomalous based on the standard deviation of the residual of input data: $\sigma_{\Theta_{AE}} = \sqrt{\frac{\sum_{i=1}^K (\mathcal{E}[i] - \mu_{\mathcal{E}})^2}{K-1}}$, where $\mathcal{E}[i] = \hat{\Theta}[i] - \Theta[i]$ is the residual of the i th sample of the sequence based on the reconstructed value $\hat{\Theta}[i]$ and actual value $\Theta[i]$, and $\mu_{\mathcal{E}}$ is the average of all residuals from input sequence with K time-ticks.




Fig. 7. *EBV+ notifies proactively: EBV_{fc&warn+} forecasts (dashed green line) the next day's core temperature (solid blue line) well and raises flags (orange vertical lines) followed by unexpected events (red boxes: intensive hive inspection, sensor recharge, and replace). It also raises alarms (red vertical lines) if h_{fc} falls below a threshold: very weak bees.* (a) Control hive (high heat stress). (b) Treated hive (low heat stress).

— As mentioned in Section 4.4, the lower limit of the search space h_{min} for estimating h is 0.10. Based on the domain experts' opinion, we assume $\beta = 0.15$.

Intensive hive inspections are stressful for bees, especially for broods, as they affect the thermoregulation ability (Section 5.1). We want to see if $EBV_{fc&warn+}$ can issue proactive warnings and detect changes in $\Theta(t)$ in real time following inspections.

- (1) *Flag*: Reestablishing homeostasis every time after intensive inspection can be tough for severely stressed hives, leading to a drop in thermoregulation ability.
- (2) *Warning*: If forecasted $h_f < \beta (= 0.15)$, bees are too weak to maintain $\Theta(t)$, and beekeepers should take actions to restore bees' health, h .

The effectiveness of $EBV_{\text{fcast\&warn+}}$ to raise flags and warnings followed by unexpected events (i.e., hive opening, sensor recharge and replace) is measured using the following two metrics: TP and FP are true positive and false positive, respectively.

$$\text{Accuracy} = \frac{\text{Total flags \& warnings raised corresponding to hive openings}}{\text{Total Hive Openings w/o Missing Data}}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} = \frac{\text{Total flags \& warnings raised corresponding to hive openings}}{\text{Total flags \& warnings raised}}$$

Results. In summary, our algorithm $EBV_{\text{fcast\&warn+}}$ raised 5 flags and 6 warnings for control hives (accuracy: 85.7% and precision: 92.3%) and 11 flags and only 1 warning for treated ones (accuracy: 90.9% and precision: 83.3%). In Figure 7, we demonstrate two examples: one from control and one from treated hive.

$EBV_{\text{fcast\&warn+}}$ raised one flag (orange vertical line) and one warning (red vertical line) for the control hive and three flags for the treated hive (Figure 7). The 1 day ahead forecasted sequence, $\hat{\Theta}(t)$ and $\hat{h}$ (dashed green line), by $EBV_{\text{fcast\&warn+}}$ using 7 days' data and the estimated daily parameter, h (solid magenta line), aligns well with the actual 1 day ahead $\Theta(t)$ and estimated h (solid blue line).

For the control hive (Figure 7(a)), $EBV_{\text{fcast\&warn+}}$ raised a flag on 22 August noon: 5 days after the first inspection, and the flag continued till 7 am on 26 August (marked by an orange area). As mentioned in Section 5.1, we suspect it to be correlated with hive cleaning and re-establishing homeostasis after an inspection. After that, there was a silence period of 3 days (marked by a gray area). A warning was issued on 1 September at midnight, which continued till our experiment's end, 27 September (marked by a red area). We strongly suspect it to be related to the previously issued flag. As bees were already exhausted from the first inspection, their health dropped below the threshold after the second one, and they could never recover from that. For the treated hive (Figure 7(b)), $EBV_{\text{fcast\&warn+}}$ raised flags on 24 August, 5 September, and 26 September: 7, 4, and 11 days after the inspections. The first two flags continued for 1 day (till 25 August at 9 pm and 6 September at 7 pm), while the third one continued for 1.5 days (till the experiment ended). There was a silent period after the first two flags were discontinued. We raise no warning for the treated hive, and we suspect that it is the positive effect of the ice treatment. Although hive openings also impacted treated hives, the severity was less than the control ones.

Although more flags were raised for treated hives (11 compared to 5 for control ones), (a) they were comparatively short: bees in treated hives recover quickly; (b) there were no warnings following the flags (except one treated hive that had comparatively higher fluctuation than other treated ones since the very beginning). To test the practical feasibility of our algorithm in a real-time setting, we have also deployed our algorithms $EBV_{\text{fit+}}$ and $EBV_{\text{fcast\&warn+}}$ at an ongoing field experiment with six homogeneous Langstroth hives (similar hive setup to the experimental ones) using an *InfluxDB* server. Data collection and processing are done using the same pipeline as before, and the dashboard is created using *Graphana*, an open source observability platform. We have demonstrated real-time health tracking and warning notification results for one hive in Figure 8 for the past 30 days (20 April–21 May 2024). The upper left figure is 1-day ahead forecasting (green line) that aligns well with recorded $\Theta(t)$ (blue line). For forecasting, we have assumed that the external temperature $\Theta_{\text{ext}}(t)$ has the same pattern as the previous day (any other weather forecasting API can also be used for a more accurate forecast). The upper right, lower left, and lower right figures demonstrate changes in daily health factors, raised flags, and warnings. Flags and warnings are presented using binary representation (0: no flag/warning; 1: continued flag/warning). There was a false, short warning on 29 April, as nothing serious was reported by the beekeeper. The second flag on 8 May 2024, was due to the disturbance created for the hive inspection on 8 May and the




Fig. 8. *EBV+ at work in real-time*: (a) upper left: forecasts (green line) the next day's core temperature (blue line) well; (b) upper right: estimates daily health factor (blue line); (c) lower left: raises flags (orange line) followed by unexpected event; and (d) lower right: raises alarms (red line) if h_{fc} falls below a threshold: very weak bees.

mite treatment done on 10 May. Finally, the third flag on 15 May coincides with our prior suspicion of the correlation between the hive inspection and re-establishing hive homeostasis, proving the feasibility of EBV+.

5.4 Q2—Explainability

As Section 5.1 mentions, each cut-point and the change in parameters between adjacent segments (Figure 4) are indications of events affecting thermoregulation and hive health. From baselines mentioned in Section 5.2, we choose ARX(1) as an example to demonstrate its interpretability in Figure 9, and we do not notice patterns among coefficients (a_0, a_1, b_0, b_1) of the model. While plotting, we have skipped the variance of residual values, σ^2 , to focus only on coefficients. Similar results are observed for other traditional models and DeepAR (thousands of parameters). Hence, we conclude that none of the baselines except EBV and EBV+ are explainable.

Effect of Intensive Hive Inspection. According to domain experts, the effect of intense hive inspections is reflected through the change in $\Theta(t)$ within a few days following the inspection. For the control hive (Figure 4(a)), both h and θ_{ideal} went down by 7% and 1°C after the first cut-point. It indicates bees' struggle to manage the cleanliness and $\Theta(t)$ following the hive inspection and sensor replacement. After the second cut-point, both h and θ_{ideal} again went down by 11% and 2°C , indicating severely affected thermoregulation and weak bees. The issued flag after the first inspection and the issued warning during the second inspection that continued till the end of the experiment (Figure 7(a)) also indicated a severely stressed hive.

Effect of Ice Treatment on Thermoregulation. Figure 7 shows the forecasted curves (dashed green line: *top graphs*) for the recorded $\Theta(t)$ (solid blue line: *top graphs*) assuming that daily parameter h will remain same for 1-day ahead forecasting (dashed green and solid blue lines: *bottom graphs*).




Fig. 9. *EBV+ is explainable*: $EBV_{fit\&cut+}$ explains events via parameter change in adjacent segments. Notice that baselines (e.g. $ARX(1)$) have no visible patterns to explain.

In the bottom graphs, we can observe (i) the low health factor ($h < 0.25$; Figure 7(a)) in control hives compared to treated ones ($h > 0.5$; Figure 7(b)); (ii) comparatively drastic change in hive health factor, h and increased fluctuation in $\Theta(t)$ of the control one, following intensive inspection.

For the treated hive in Figure 4(b), $EBV_{fit\&cut+}$ detects only one cut-point (only after the first inspection) with a decrease of 2.7% and 1°C in h and θ_{ideal} , respectively. No parameter changes were detected for the following two events. Also, as shown in Figure 7(b), $EBV_{fcast\&warn+}$ raises only three flags (no warnings). This indicates that bees were healthy enough to combat the aftereffect of intensive inspections and control the fluctuation in $\Theta(t)$.

These results coincide with domain experts' expectations: ice packs during hot days will help bees regulate $\Theta(t)$ and manage the aftereffects of intensive inspections.

5.5 Q3–Scalability

Our proposed model, EBV_{model+} , is fast and takes only less than 1 minute on average to process and fit 2 months of data for a well-defined search space on the laptop described above. We conclude with the following lemmas.

PROPOSITION 5.1. *Reconstruction using EBV_{model+} is linear in time, i.e., for t time ticks, the time complexity of reconstruction is $O(t)$.*

PROOF. We use the LM optimization, and the time for convergence is linear with the number of time ticks. $\square$

All (solid) lines in Figure 10 (time to reconstruct 2 months of data for experimental hives in logarithmic scale) are parallel to the reference line of linear growth (dashed blue line), indicating scalability.




Fig. 10. *EBV+* is linear with input size: growth parallel to diagonal.




Fig. 11. *EBV+* spot differences: health factor, h vs. the pdf of h . It corroborates that bees in treated hives are stronger. (a) Pdf of h .

PROPOSITION 5.2. *Segmentation using $EBV_{fit\&cut+}$ is linear in time, i.e., the time complexity for finding m segments in a given dataset $\mathcal{D}$ with t time ticks is $O(t \times m)$.*

PROOF. For a time sequence with t time ticks, we use a greedy algorithm to find the m cut-points. Each cut-point requires one scan of the sequence $O(m)$, and thus, the complexity is $O(t \times m)$. $\square$

5.6 Q4—Observations

Effect of Intensive Hive Inspection. As described in Sections 5.1 and 5.3, our algorithm $EBV_{fit\&cut+}$ detected 80% more cut-points and $EBV_{fcast\&warn+}$ raised 80% more warnings in total for control hives than for treated ones for the same number (3 per hive) and type of (intensive) inspection. That is:

OBSERVATION 1. *Control (= untreated) hives suffer more from intensive hive inspections.*

That is, control hives show more discontinuities and less control over their temperature.

Effect of Ice Treatment on Thermoregulation. We demonstrate the distribution of health factor of different hives over the experimental period as a plot of health factor, h vs. the pdf of h in Figure 11. Every data point corresponds to a (density of h , range of h) pair. Here, we have divided the values

of h into 10 bins, each with a range of 0.1. Notice that the first bin $([0, 0.1])$ is empty as we assumed the lower limit of h , $h_{min} = 0.1$. Notice that bees in control (= untreated) hives have higher density in the lower region: $[0.1, 0.3]$ compared to treated hives (dense in the upper region: $[0.3, 1]$).

OBSERVATION 2. *Bees in treated hives are stronger, i.e., they have better thermoregulation ability than the ones in control hives.*

Difference in Thermoregulation Ability Heating and Cooling. In Figure 11, we have mentioned the maximum heating and cooling limit (r_c, r_h) calculated using Algorithm 1. Notice that, $r_c (= 80) > r_h (= 221)$, coinciding with domain experts' expectation:

OBSERVATION 3. *Heating a hive is easier than cooling it down.*

6 Conclusions, Significance, and Impact

We proposed an unsupervised method, EBV+, for easier understanding, effective monitoring, and reliable analysis of hive health by beekeepers to take timely preventive measures.

EBV+ has all desirable properties any beekeeper would want in an effective hive monitoring system: (a) it is principled, based on physics (thermal diffusion equations for heat transfer) and control theory (feedback-loop “sigmoid” P-controllers for realistic upper bounds on bees' thermoregulation ability); (b) it can interpret the change in thermoregulation with only a handful of parameters (i.e., hive health factor); (c) it is effective in fitting, summarizing, forecasting, and detecting events proactively affecting hive health; (d) scalable; and (e) informative, agreeing with experts' expectations.

Significance. Human civilization largely depends on honeybees for food production. Reliable hive monitoring tools with trustworthy explanations are one of the pressing priorities in the current world. A method for analyzing honeybees' thermoregulation behavior with confirmed analytic capabilities will be a game changer for the global beekeeping industry. In such a scenario, our advanced method, EBV+, will significantly attract the broad interest of beekeepers, growers, and researchers studying all facets of pollinator health worldwide.

Impact. A preliminary version of our model received positive feedback from local beekeepers at the UC Riverside CIBER Bee Health conference in 2022 [38]. We are in constant contact with beekeepers and co-organize yearly events with CIBER to keep them updated about our endeavor. Our implementation and the full dataset is available on GitHub: <https://github.com/rtenlab/EBeeVet>

Acknowledgments

We extend special thanks to Jamison Scholer for his help managing the bee hives and Mohsen Karimi for his help collecting the data.

References

- [1] J. Crank. 1975. *The Mathematics of Diffusion*. Oxford University Press.
- [2] Abdullah Ateş, Barış Baykant Alagöz, Celalettin Yeroğlu, and Hafız Alisoy. 2015. Sigmoid based PID controller implementation for rotor control. In *Proceedings of the 2015 European Control Conference (ECC)*. IEEE, 458–463.
- [3] Byoung-Kee Yi, H. V. Jagadish, and Christos Faloutsos. 1998. Efficient retrieval of similar time sequences under time warping. In *Proceedings of the 14th International Conference on Data Engineering*, 201–208.
- [4] Byoung-Kee Yi and Christos Faloutsos. 2000. Fast time sequence indexing for arbitrary Lp norms. In *Proceedings of the 26th International Conference on Very Large Data Bases*, 385–394.
- [5] Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, and Hoang Anh Dau. 2016. Matrix profile I: All pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets. In *Proceedings of the 2016 IEEE 16th International Conference on Data Mining (ICDM)*, 1317–1322.

- [6] Chau-Chung Songa, Geng-Yi Lin, Chi-Chung Peng, and Chung-Wen Hung. 2022. Development of intelligent beehive and network monitoring system for bee ecology. In *Proceedings of the 27th International Conference on Artificial Life and Robotics*, 641–644.
- [7] Christos Faloutsos, M. Ranganathan, and Yannis Manolopoulos. 1994. Fast subsequence matching in time-series databases. *ACM SIGMOD Record* 23, 2 (May 1994), 419–429.
- [8] Christos Faloutsos, M. Ranganathan, and Yannis Manolopoulos. 1994. Fast subsequence matching in time-series databases. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data*, 419–429.
- [9] Clifford M. Hurvich and Chih-Ling Tsai. 1993. A corrected Akaike information criterion for vector autoregressive model selection. *Journal of Time Series Analysis* 14 (1993), 271–279.
- [10] Caner Turkmen, Nick Erickson, and Oleksandr Shchur. 2022. Easy and Accurate Forecasting with AutoGluon-TimeSeries. Retrieved from <https://aws.amazon.com/blogs/machine-learning/easy-and-accurate-forecasting-with-autogluon-timeseries/>
- [11] Cheng Yang, John Collins, and Mark Beckerleg. 2017. A model for pollen measurement using video monitoring of honey bees. *Sensing and Imaging* 19 (Dec 2017), 1–29.
- [12] Douglas Santiago Kridi, Carlos Giovanni N. de Carvalho, and Danielo G. Gomes. 2016. Application of wireless sensor networks for beehive monitoring and in-hive thermal patterns detection. *Computers and Electronics in Agriculture* 127 (May 2016), 221–235.
- [13] Elisabeth Eilers, Claire Kremen, Sarah Smith Greenleaf, Andrea K. Garber, and Alexandra-Maria Klein. 2011. Contribution of pollinator-mediated crops to nutrients in the human food supply. *PLoS One* 6 (Jun. 2011), 1–6.
- [14] Fiona Edwards Murphy, Bruno Srbnovski, Michele Magno, Emanuel M. Popovici, and Pádraig M. Whelan. 2015. An automatic, wireless audio recording node for analysis of beehives. In *Proceedings of the 2015 26th Irish Signals and Systems Conference (ISSC)*. IEEE, Carlow, Ireland, 1–6.
- [15] Fiona Edwards-Murphy, Michele Magno, Pádraig M. Whelan, John O'Halloran, and Emanuel M. Popovici. 2016. b+WSN: Smart beehive with preliminary decision tree analysis for agriculture and honey bee health monitoring. *Computers and Electronics in Agriculture* 124 (2016), 211–219.
- [16] George E. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung. 2015. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, USA.
- [17] Hasila Jarimi, Emmanuel Tapia-Brito, and Saffa Riffat. 2020. A review on thermoregulation techniques in honey bees' (*Apis Mellifera*) beehive microclimate and its similarities to the heating and cooling management in buildings. *Future Cities and Environment* 6 (Aug. 2020), 1–7.
- [18] Joe-Air Jiang, Chien-Hao Wang, Chi-Hui Chen, Min-Sheng Liao, Yu-Li Su, Wei-Sheng Chen, Chien-Peng Huang, En-Cheng Yang, and Cheng-Long Chuang. 2016. A WSN-based automatic monitoring system for the foraging behavior of honey bees and environmental factors of beehives. *Computers and Electronics in Agriculture* 123 (Apr. 2016), 304–318.
- [19] Johanne Brunet and Fabiana P. Fragoso. 2024. What are the main reasons for the worldwide decline in pollinator populations? *CABI Reviews* (Jan. 2024). Retrieved from <https://www.cabidigitallibrary.org/doi/10.1079/cabreviews.2024.0016>
- [20] Leman Akoglu and Christos Faloutsos. 2010. Event detection in time series of mobile communication graphs. In *Proceedings of the 27th Army Sci Conference*. Retrieved from https://www.andrew.cmu.edu/user/lakoglu/pubs/EVENTDETECTION_AkogluFaloutsos.pdf
- [21] Bryan Lim, Sercan O. Arik, Nicolas Loeff, and Tomas Pfister. 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting* 37, 4 (2021), 1748–1764.
- [22] Lexiang Ye and Eamonn Keogh. 2009. Time series shapelets: A new primitive for data mining. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 947–956.
- [23] Marco Kleinhenz, Brigitte Bujok, Stefan Fuchs, and Jürgen Tautz. 2003. Hot bees in empty broodnest cells: Heating from within. *The Journal of Experimental Biology* 206 (Dec. 2003), 4217–4231.
- [24] Matthew L. Forister, Emma M. Pelton, and Scott H. Black. 2019. Declines in insect abundance and diversity: We know enough to act now. *Conservation Science and Practice* 1 (2019), e80.
- [25] Markus Petz, Anton Stabentheiner, and Karl Crailsheim. 2004. Respiration of individual honeybee larvae in relation to age and ambient temperature. *Journal of Comparative Physiology B* 174 (Oct. 2004), 511–518.
- [26] Mst. Shamima Hossain, Christos Faloutsos, Boris Baer, Hyoseung Kim, and Vassilis J. Tsotras. 2024. EBV: Electronic bee-veterinarian for principled mining and forecasting of honeybee time series. In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*. SIAM, Philadelphia, PA, 298–306.
- [27] Monica Vercelli, Silvia Novelli, Paola Ferrazzi, Giada Lentini, and Chiara Ferracini. 2021. A qualitative analysis of beekeepers' perceptions and farm management adaptations to the impact of climate change on honey bees. *Insects* 12, 3 (Mar. 2021), 228.
- [28] Nathalie Steinhauer, et al. 2024. United States Honey Bee Colony Losses 2022-2023: Preliminary Results from the Bee Informed Partnership. Retrieved from <https://beeinformed.org/2023/06/22/united-states-honey-bee-colony-losses-2022-23-preliminary-results-from-the-bee-informed-partnership/>

- [29] Padraig Davidson, Michael Steininger, Florian Lautenschlager, Konstantin Kobs, Anna Krause, and Andreas Hotho. 2020. Anomaly detection in beehives using deep recurrent autoencoders. arXiv:2003.04576. Retrieved from <https://arxiv.org/abs/2003.04576>
- [30] Stefania Cecchi, Susanna Spinsante, Alessandro Terenzi, and Simone Orcioni. 2020. A smart sensor-based measurement system for advanced bee hive monitoring. *Sensors* 20, 9 (Apr. 2020), 2726.
- [31] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting* 36, 3 (2020), 1181–1191.
- [32] Simon G. Potts, Vera Imperatriz-Fonseca, Hien T. Ngo, Marcelo A. Aizen, Jacobus C. Biesmeijer, Thomas D. Breeze, Lynn V. Dicks, Lucas A. Garibaldi, Rosemary Hill, Josef Settele, et al. 2016. Safeguarding pollinators and their values to human well-being. *Nature* 540, 2 (Dec. 2016), 220–229.
- [33] Simone Tosi, Fabien J. Démares, Susan W. Nicolson, Piotr Medrzycki, Christian W. W. Pirk, and Hannelie Human. 2016. Effects of a neonicotinoid pesticide on thermoregulation of African honey bees (*Apis mellifera scutellata*). *Journal of Insect Physiology* 93–94 (Oct.–Nov. 2016), 56–63.
- [34] Tony Zhang, Szymon Zmyslony, Sergei Nozdrenkov, Matthew Smith, and Brandon Hopkins. 2021. Semi-supervised audio representation learning for modeling beehive strengths. arXiv:2105.10536. Retrieved from <https://arxiv.org/abs/2105.10536>
- [35] V. Pandimurugan, Rahul Mandviya, Abhinav Gadgil, Kumar Prakhar, and Aditya Datar. 2021. IoT based smart beekeeping monitoring system for beekeepers in India. In *Proceedings of the 2021 4th International Conference on Computing and Communications Technologies (ICCCCT)*, 65–70.
- [36] Waddah Saeed and Christian Omlin. 2023. Explainable AI (XAI): A systematic meta-survey of current challenges and future opportunities. *Knowledge-Based Systems* 263 (2023), 110273.
- [37] Yasuko Matsubara, Yasushi Sakurai, and Christos Faloutsos. 2014. AutoPlait: Automatic mining of co-evolving time sequences. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, 193–204.
- [38] Center for Integrative Bee Research (CIBER). 2022. Ciber Bee Health Conference at UC Riverside. Retrieved from <https://ciber.ucr.edu/bee-health-conference-2022>
- [39] James D. Hamilton. 1994. *Time Series Analysis*, Vol. 2. Princeton University Press.
- [40] Norman S. Nise. 2019. *Control Systems Engineering*. John Wiley & Sons.
- [41] Nordic Semiconductor. 2024. Nordic Thingy:52. Retrieved from <https://www.nordicsemi.com/Products/Development-hardware/Nordic-Thingy-52>

Received 2 June 2024; revised 21 December 2024; accepted 4 February 2025