Jurnal Inovasi Teknologi dan Rekayasa
Vol. 11, No. 2, July-December 2026, pp. 390-404

A Systematic Review of IoT-Enabled Predictive Drying for Arabica Coffee: Toward a Python-Based Digital Twin Framework for the Gayo Arabica Coffee

Muhammad Ihsan [a,1,*], H. Susanto [b,2], N. Fitriadi [c,3]

[a] *Faculty of Engineering, Universitas Gajah Putih, Takengon, 24552, Indonesia*
[b] *Mechanical Engineering Dept., Universitas Teuku Umar, Meulaboh, 23681, Indonesia*
[c] *Mechanical Engineering, Politeknik Aceh Selatan, Aceh, Indonesia, 23711, Indonesia*

[1] *ihsan.ccrg@gmail.com*; [2] herdisusanto@utu.ac.id; [3] nuzuli@poltas.ac.id*
** corresponding author*

ARTICLE INFO

Article history:
Published
July 30, 2026

Keywords:
Digital twin
Internet of Thinks
Drying kinetics
Arabica Coffee
Phyton simulation

ABSTRACT

Post-harvest drying is the single most decisive determinant of Arabica coffee quality, yet in the Gayo Highlands of Aceh Tengah (950 - 1,650 m.a.s.l.) it remains an experience based, weather dependent operation conducted largely on tarpaulins and open patios. High relative humidity, frequent rainfall, and low ambient temperature routinely prolong drying beyond ten days, produce non uniform final moisture content, and expose parchment to fungal colonization and ochratoxin risk, causing smallholders to fail the 12.5% maximum moisture threshold of SNI 01-2907-2008 and to forfeit specialty price premiums. This systematic literature review, conducted following PRISMA guidelines on peer reviewed publications from 2019–2026, synthesizes three research streams that have so far evolved in isolation: (i) thin layer drying kinetics of parchment coffee, (ii) Internet of Things (IoT) architectures for solar and hybrid dryers, and (iii) digital twin methodology in food and grain drying. The review finds that existing coffee drying IoT deployments are overwhelmingly reactive while validated kinetic models (modified Midilli, Page) capable of such prediction remain confined to laboratory dryers and are never coupled to field sensor streams. Digital twin research in drying has concentrated on grain and fruit, leaving coffee, and highland smallholder contexts in particular, unaddressed. From this synthesis the review derives seven research gaps and proposes SIKOPI-DT, a contextualized five-layer Python based digital twin framework in which a modified Midilli kinetic core is continuously reparametrized by low-cost IoT sensors (ESP32, SHT31, load cell) to forecast the time to target moisture content and to drive predictive supplementary-heat control under Gayo weather scenarios. Synthesized evidence indicates that such a system can plausibly reduce drying time by 30–50%, cut specific energy consumption by 15–30%, and improve moisture uniformity relative to open-sun practice. A three phase experimental validation roadmap and a smallholder economic feasibility analysis are presented, positioning this review as the theoretical foundation for a physical prototype in Aceh Tengah.

I. Introduction

Coffee is among the most economically significant agricultural commodities in the developing world, and Indonesia ranks consistently among its largest producers. Within Indonesia, the Gayo Highlands, comprising Aceh Tengah, Bener Meriah, and Gayo Lues at elevations of roughly 950 - 1,650 meters above sea level, constitute the country's principal Arabica production zone, benefiting from cool temperatures, abundant rainfall, and fertile volcanic soils that support specialty-grade cup profiles [1], [2]. Gayo Arabica holds geographical indication status and enjoys established demand in international specialty markets. Yet the economic value captured by the smallholders who produce it remains disproportionately small, and a substantial share of that shortfall is traceable not to cultivation but to post-harvest processing.

Among post-harvest unit operations, drying is decisive. The transition of wet parchment from an initial moisture content of approximately 45–55% (wet basis) down to the storage stable range of 11 - 12.5% governs microbial stability, the formation of color and flavor precursors, and the shelf life of the green bean [3], [3]. Indonesian National Standard SNI 01-2907-2008 sets a maximum moisture content of 12.5% for traded coffee beans, and this threshold is routinely enforced by exporters and collectors [4]. Drying that is too rapid or conducted at excessive temperature degrades beverage quality, consequently drying that is too slow, or that stalls under humid conditions, permits the proliferation of Aspergillus species and the accumulation of ochratoxin A. Direct field studies of solar tunnel drying of wet-processed parchment have reported fungal infection at the end of drying ranging from a few percent to more than ninety percent depending on layer thickness and dryer zone, a spread that translates directly into mycotoxin risk [4]. In Aceh Tengah, drying is still overwhelmingly performed by open-sun exposure on tarpaulins, concrete patios, or simple raised beds (para-para). This practice is inexpensive and requires no capital, but it is entirely hostage to weather. The very agroclimatic conditions that make the Gayo Highlands excellent for growing Arabica (cool air, high humidity, and frequent orographic rainfall) are precisely those that make sun shine drying slow and unreliable. Drying campaigns commonly extend beyond seven to ten days. It parchment is repeatedly rewetted by forecast rain, farmers judge dryness by biting a bean or by feel rather than by measurement and the resulting batch exhibits substantial intra batch variance in final moisture content. Comparative work on Indonesian Arabica has confirmed that raised bed and conventional sun drying differ materially in the temperature, humidity, and light regimes they impose, and consequently in the moisture trajectories and physical quality they produce [4].

The obvious engineering response is a controlled dryer, and a considerable body of work has developed solar, greenhouse effect, and hybrid solar biomass dryers for coffee [5], [6]. More recently, these have been instrumented with Internet of Things (IoT) hardware: low cost microcontrollers such as the ESP32, temperature and humidity sensors such as the DHT22 or SHT31, cloud dashboards, and in some cases fuzzy logic controllers that actuate fans and heaters [5], [5]. Such systems represent genuine progress. However, a close reading of this literature reveals a persistent limitation. Therefore these deployments are fundamentally reactive. They sense the state of the drying air and, at best, hold that air within a setpoint band. They do not model the product. They cannot answer the question the farmer actually needs answered.

The knowledge required to answer that question exists, but in a different literature. Thin-layer drying kinetics has produced established semi empirical models that describe the moisture ratio of a product as a function of time under given air conditions. For parchment Arabica specifically, the modified Midilli model has repeatedly been identified as the best fitting formulation across controlled temperature and humidity regimes, with coefficients of determination exceeding 0.99 [3], [3], [7]. These models are, in principle, exactly the predictive engine that IoT dryers lack. In practice, they are almost invariably fitted offline to laboratory data, reported as a table of coefficients, and never deployed as a live component of a control system.

The concept that unifies these two literatures is the digital twin: a virtual replica of a physical process, continuously synchronized with it by sensor data, capable of simulating forward in time and thereby of supporting prediction, optimization, and control [8], [8, 9]. Digital twin methodology has been applied to food processing generally [9], to solar drying of fruit, where a physics-based twin was used to identify trade-offs between drying time, product quality, and energy use [10], and, most extensively, to grain drying, where a recent review argues that the twin paradigm offers a route out of the opacity and experience dependence that characterize conventional practice [8, 9, 10].

This convergence defines the opportunity addressed here. The kinetic models are validated. The IoT hardware is cheap and available. The digital twin methodology is established in adjacent domains. What has not been done is to integrate them into a single, low cost, field deployable architecture adapted to the specific agroclimatic and socioeconomic realities of highland smallholder coffee reproducible computational environment such as Python, so that the resulting framework can be inspected, adapted, and extended rather than locked inside proprietary software.

Accordingly, this systematic literature review pursues six objectives: (1) to characterize the biophysical and socioeconomic constraints on coffee drying in the Gayo Highlands of Aceh Tengah; (2) to systematically synthesize validated thin-layer kinetic models for parchment Arabica and to identify the formulation most suitable as a predictive core; (3) to critically evaluate existing IoT and control architectures for coffee and analogous dryers, distinguishing reactive monitoring from predictive capability; (4) to review digital twin methodology in food and grain drying and extract transferable design principles; (5) to derive the research gaps that follow from this synthesis; and (6) to propose SIKOPI-DT, a contextualized Python-based digital twin framework, together with an experimental validation roadmap and an economic feasibility assessment for smallholder deployment in Aceh Tengah.

II. Literature

A. Coffee post-harvest processing and the centrality of drying

Wet-processed (fully washed) Arabica, the dominant method in the Gayo Highlands, proceeds through harvesting, pulping, fermentation, washing, drying, hulling, and sorting. After washing, parchment coffee carries a moisture content of roughly 45–55% on a wet basis and must be reduced to 11–12.5% for stable storage and trade [3], [4]. Drying is therefore the operation that both consumes the most time and carries the most risk. It is also the operation over which the smallholder has the least reliable control, because the energy source (solar radiation) is stochastic and, in the highlands, frequently insufficient.

The consequences of poor drying are well documented and cumulative. Incomplete drying leaves the parchment above the water-activity threshold at which Aspergillus ochraceus and related species proliferate; the resulting ochratoxin A contamination is a food-safety hazard subject to strict import limits in the European Union. Non-uniform drying within a batch means that even where the mean moisture content is acceptable, a subpopulation of beans remains hazardous, and this heterogeneity is invisible to the destructive spot-check methods used in practice. Excessive drying temperature or an excessive drying rate, conversely, damages the cell structure and degrades cup quality; studies controlling dry-bulb and dew-point temperature independently have shown that beverage quality depends on the interaction between them, not on temperature alone [3]. Drying is thus a constrained optimization, not a maximization: the objective is to reach the target moisture content as quickly as possible subject to an upper bound on temperature and drying rate. This framing is essential, and it is precisely the framing that a predictive model enables and that a reactive controller cannot express.

B. Thin layer drying kinetics of parchment coffee

Thin-layer drying theory describes the moisture ratio MR of a hygroscopic product as a decaying function of time under constant air conditions. The moisture ratio is defined as:

$$MR = \frac{M - M_e}{M_0 - M_e} \tag{1}$$

Where M is the instantaneous moisture content (dry basis), M_0 the initial moisture content, and M_e the equilibrium moisture content, itself a function of air temperature and relative humidity through a sorption isotherm. Where M_e is small relative to M_0, the simplification $MR \approx M/M_0$ is often adopted, though this is inadvisable for the humid highland conditions considered here, in which M_e

is not negligible. A family of semi empirical equations has been fitted to MR data across food products. The principal formulations encountered in the coffee literature are summarized in Table 1. Comparative studies of parchment Arabica converge with unusual consistency. Phitakwinai et al., drying parchment coffee under controlled dry bulb temperatures of 50 - 70 °C and relative humidities of 10 - 30%, fitted nine models and identified the modified Midilli formulation as the best fit, additionally deriving an Arrhenius type expression for effective moisture diffusivity as a function of temperature at each relative humidity [3]. Abdissa et al., studying wet processed parchment in a solar tunnel dryer across three dryer zones and three layer thicknesses, found the modified Midilli model best described the kinetics in most zone thickness combinations, with the Verma and two term models prevailing only in specific shallow layer cases [4]. Work on hot-air drying of coffee at 40–100 °C reported the Midilli model achieving R^2 between 0.995 and 0.9995 with RMSE below 0.21% [3]. The drying process in all cases was dominated by a falling rate period, confirming that internal moisture diffusion, rather than surface evaporation, governs the kinetics, which is why a diffusion consistent semi empirical model is appropriate.

Table 1. Principal thin-layer drying models applicable to parchment coffee

Model	Equation	Parameters	Reported performance for coffee
Newton (Lewis)	MR = exp(−kt)	k	Poor; single parameter cannot capture falling-rate curvature
Page	MR = exp(−k·t^n)	k, n	Good; frequently second-best, computationally light [3]
Henderson & Pabis	MR = a·exp(−kt)	a, k	Moderate
Logarithmic	MR = a·exp(−kt) + c	a, k, c	Moderate to good
Midilli	MR = a·exp(−k·t^n) + b·t	a, k, n, b	Excellent; R^2 0.995–0.9995, RMSE < 0.21% [4]
Modified-Midilli	MR = exp(−k·t^n) + b·t	k, n, b	Best fit for parchment Arabica across studies [3], [6]

Note: t = drying time (h); k = drying constant (h^{-n}); n, a, b, c = empirical coefficients. Coefficients are functions of air temperature and relative humidity and must be re-estimated for local conditions.

Two implications follow, and they are central to the argument of this review. First, the modified-Midilli model, expressed as

$$MR(t) = \exp(-kt^n) + bt \tag{2}$$

Defensible default predictive core for parchment Arabica, with the Page model as a lightweight fallback where computational resources on the edge device are constrained. The coefficients k, n, and b are not universal constants. They are functions of the drying air temperature and relative humidity, and therefore of the local climate and dryer design. A model fitted in a Thai or Ethiopian laboratory at 50 - 70 °C and 10 - 30% RH cannot be transplanted unmodified to a Gayo solar dryer operating at 32 - 45 °C and 55 - 80% RH. This is not a defect of the models; it is a specification for the system that should use them. The coefficients must be estimated locally, from local data. That requirement is, in essence, a specification for a digital twin.

C. Solar and hybrid dryers for coffee

Dryer configurations reported for coffee span open sun and raised bed drying, direct and indirect solar dryers; greenhouse-effect (efek rumah kaca) dryers, which are widely trialled in Indonesia; solar tunnel dryers; and hybrid systems that supplement solar gain with biomass combustion or electrical resistance heating. Reported capacities range from a single kilogram in laboratory rigs to twenty tonnes in commercial installations [3]. Hybrid solar biomass configurations are of particular relevance to the Gayo context because coffee production generates its own biomass residue in the form of pulp and parchment husk, offering an energy source that is locally available and effectively free.

The performance literature is broadly consistent: enclosed solar dryers reduce drying time relative to open-sun exposure, protect the product from rain and contamination, and raise final quality. Their weakness is equally consistent. Under prolonged cloud or rain, a passive solar dryer becomes little more than a warm humid box, and the drying rate collapses. This is exactly the failure mode most relevant to Aceh Tengah. The engineering answer is supplementary heat, however supplementary heat costs money and fuel, and the question of when to apply it, and for how long, is a scheduling problem that a reactive thermostat solves inefficiently and a predictive model solves well. The economic case for a digital twin in this setting rests substantially on this point.

D. Iot architectures for agricultural drying

IoT-instrumented dryers reported in the literature share a recognisable layered architecture, and it is useful to state it explicitly because the proposed framework extends it rather than replacing it.

1) Sensing layer

Typical instrumentation comprises air temperature and relative humidity sensors (DHT22, SHT31, or SHT85), product or plenum temperature via DS18B20 or thermocouples, solar irradiance via pyranometer or calibrated photodiode, airflow velocity, and direct product mass measurement via load cell. The rarity of the load cell is notable, because continuous gravimetric measurement is the only practical route to online moisture content estimation and hence to any predictive capability. Its absence from most deployments is a direct cause of their reactive character.

2) Connectivity layer

ESP32 and ESP8266 microcontrollers with integrated Wi-Fi dominate the reported implementations, typically transmitting to cloud platforms such as ThingSpeak, Blynk, or Firebase. For highland deployments where cellular coverage is patchy, LoRa or LoRaWAN offers long range, low power transmission from dryer to gateway, and edge first architectures that compute locally and synchronise opportunistically are markedly more robust than cloud-dependent designs.

3) Analytics and control layer

This is where the literature is thinnest and where the gap is clearest. The majority of reported coffee drying IoT systems perform threshold logic: if humidity exceeds a limit, actuate the exhaust fan; if temperature falls below a setpoint, engage the heater. A more sophisticated minority employ fuzzy logic inference to smooth control action across multiple inputs [5], [5]. Both approaches are reactive: they respond to the current measured state of the drying air. Neither maintains a model of the product, and therefore neither can forecast the moisture trajectory, estimate time-to-target, or evaluate whether a given heating action will pay for itself. Model predictive control, standard in industrial drying, is essentially absent from the smallholder coffee literature.

4) Actuation and interface layers

Actuation typically comprises exhaust fans, circulation blowers, heater relays, and, in enclosed dryers, motorised vents. Interfaces are usually web dashboards or smartphone applications; the literature on agricultural technology adoption in Indonesia consistently identifies digital literacy as a binding constraint [11], indicating that interfaces for this user population must be simple, visual, local-language, and preferably capable of pushing a single actionable recommendation rather than requiring interpretation of a time-series plot.

E. Digital twin methodology in drying

A digital twin is a virtual representation of a physical asset or process, bidirectionally coupled to it by real-time data, that supports monitoring, prediction, optimization, and control [8], [9]. The essential distinction from conventional simulation is the live coupling: the twin ingests sensor data, updates its own state and parameters, and produces forecasts that feed back into decisions on the physical system. In food processing, digital twins have been applied to refrigerated transport, thermal processing, and drying [9]. Prawiranto et al. constructed a physics-based digital twin of solar fruit drying and used it to expose the trade offs between drying time, product quality, and energy use demonstrating precisely the multiobjective reasoning that a purely reactive controller cannot perform [10]. Schemminger et al. developed a digital twin of carrot slice drying that predicted quality retention while accounting for biological variability between individual slices, illustrating that twins can address intra batch heterogeneity, the same problem that afflicts coffee batches in Aceh Tengah [8, 9]. In grain drying, digital twin research is now sufficiently developed to support a dedicated review, which frames the technology as a response to the opacity, experience-dependence, uneven quality, and poor energy efficiency of conventional practice [8, 9, 10] a description that applies to Gayo coffee drying with almost no modification. Two observations follow. The first is that the methodology is mature enough to transfer. The second is that it has not been transferred to coffee, and certainly not to smallholder coffee. Existing food-drying digital twins are predominantly physics-based, relying on computational fluid dynamics or coupled heat-and-mass-transfer partial differential equations, and are computationally heavy, expensive to construct, and dependent on proprietary solvers. A twin built on a validated semi-empirical kinetic model rather than a full physics solver would be very much less accurate in an absolute sense and very much more deployable on a Raspberry Pi in a village in Aceh Tengah. This review argues that for the decision the farmer actually faces, the second is worth more than the first.

F. Barriers to adoption in smallholder contexts

Evidence on IoT adoption among Indonesian farmers identifies a consistent set of constraints [11], [12]. Economic barriers dominate: capital is scarce, credit access is limited, and returns on technology investment are perceived as uncertain. Infrastructural barriers follow: electricity supply in highland villages is unreliable and cellular connectivity is intermittent. Human barriers are equally binding: digital literacy is low, particularly among older farmers, and trust in technology-mediated recommendations must be earned. Institutional barriers complete the picture: extension services are thinly resourced and rarely equipped to support digital tools. These constraints are not peripheral considerations to be noted and set aside. They are hard design requirements. A system that assumes reliable grid power, continuous connectivity, individual ownership, or a user comfortable with a time-series dashboard will not be adopted in Aceh Tengah, however sound its engineering.

G. Research gaps

Systematic synthesis of the four literatures reviewed above yields seven interrelated gaps:

Table 2. Research gaps identified from the systematic synthesis

No.	Gap	Description
G1	Predictive–reactive gap	IoT coffee dryers monitor and hold setpoints but do not model the product; no reported system forecasts time-to-target moisture content.
G2	Model deployment gap	Validated thin-layer kinetic models are fitted offline to laboratory data and never coupled to live field sensor streams.
G3	Commodity gap in digital twins	Digital twin drying research addresses grain and fruit; coffee is unaddressed.
G4	Agroclimatic transferability gap	Kinetic coefficients are reported for 50–100 °C, low-RH laboratory air; Gayo solar dryers operate at 32–45 °C and 55–80% RH
G5	Energy scheduling gap	The decision of when to apply supplementary heat in a hybrid dryer is treated thermostatically
G6	Uniformity gap	Intra-batch moisture heterogeneity, the principal driver of mycotoxin risk.
G7	Economic viability gap	No costed analysis exists of predictive drying control.

III. Method

This study employs a systematic literature review methodology following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines. Because the review deliberately spans three research streams that have developed independently. Drying kinetics, IoT dryer instrumentation, and digital twin methodology. The search strategy was constructed as three parallel strands with a fourth strand for the local agroecological and socioeconomic context. The synthesis is integrative rather than purely aggregative: its purpose is not only to summarise what is known within each stream but to identify what becomes possible when they are combined.

A. Inclusion and exclusion criteria

Inclusion criteria specified: (1) peer-reviewed journal articles, conference proceedings, and accredited national-journal publications appearing between 2019 and 2026, with foundational kinetic and methodological works admitted irrespective of date where they define a model still in use; (2) studies reporting quantitative drying outcomes (drying time, moisture trajectory, model coefficients and fit statistics, energy consumption, or product quality) or specifying an IoT or digital twin architecture in sufficient detail to be reproduced; (3) empirical work with an explicit methodological description; (4) publications in English or Indonesian.

Exclusion criteria eliminated: (1) purely conceptual papers with neither empirical validation nor implementable architectural specification; (2) studies of non-agricultural drying (timber, ceramics, pharmaceuticals) except where the digital twin methodology was directly transferable; (3) publications omitting model fit statistics or system specifications; (4) duplicate reports of the same dataset; (5) non-peer-reviewed commercial white papers.

B. Search strategy

A comprehensive search was conducted across Scopus, Web of Science, IEEE Xplore, ScienceDirect, SpringerLink, and Google Scholar, supplemented by Indonesian national databases (Garuda, SINTA-indexed journals) for locally published work on Gayo coffee and Indonesian dryer design. Boolean search strings were constructed for each strand:

Table 3. Search strings by literature strand

Strand	Boolean search string
S1 , Drying kinetics	("thin layer drying" OR "drying kinetics" OR "moisture ratio") AND (coffee OR "Coffea arabica" OR parchment) AND (Midilli OR Page OR "mathematical model*")
S2 , IoT dryers	("Internet of Things" OR IoT OR ESP32 OR Arduino OR LoRa) AND (dryer OR drying OR "solar dryer" OR greenhouse) AND (coffee OR "agricultural product*") AND (monitoring OR control OR fuzzy)
S3 , Digital twin	("digital twin" OR "virtual replica" OR "model predictive control") AND (drying OR dryer OR "food processing" OR grain) AND (simulation OR prediction OR optimi*ation)
S4 , Local context	(Gayo OR "Aceh Tengah" OR "Bener Meriah" OR "highland Indonesia") AND (coffee OR kopi) AND (pascapanen OR "post-harvest" OR pengeringan OR drying OR quality OR mutu)

C. Study selection and quality appraisal

Selection proceeded in three stages. Title-and-abstract screening removed clearly irrelevant records. Full-text assessment evaluated methodological quality, relevance to the review objectives, and the availability of extractable quantitative data. Quality appraisal applied a structured checklist assessing, for kinetic studies, the adequacy of the experimental design, the number of models compared, the reporting of fit statistics (R^2, RMSE, χ^2), and the validity of the moisture content determination method; and, for IoT and digital twin studies, the completeness of the architectural specification, the presence of field as opposed to purely laboratory validation, and the transparency of the control logic. Two reviewers conducted selection independently, with disagreements resolved by discussion and, where necessary, by third-party arbitration.

D. Data extraction and synthesis

Standardised extraction forms captured: (1) study characteristics (authors, year, country, design); (2) product and process conditions (coffee type and processing method, initial and final moisture content, air temperature, relative humidity, airflow, layer thickness, batch mass); (3) kinetic findings (models compared, best-fit model, coefficient values, R^2, RMSE, effective moisture diffusivity, activation energy); (4) system specifications (sensors, microcontroller, communication protocol, control logic, actuators, interface); (5) outcomes (drying time, specific energy consumption, moisture uniformity, quality indicators, fungal load); and (6) implementation factors (cost, adoption barriers, success factors). Quantitative kinetic findings were synthesised comparatively by tabulating best fit models and coefficient ranges against air conditions, permitting assessment of how coefficients vary with temperature and relative humidity, the variation that motivates online re estimation. System architectures were synthesised through structured comparison against the layered reference model of Section II-D, with each study classified by the highest level of analytical capability it demonstrated: monitoring only, threshold control, fuzzy control, or predictive control. Qualitative implementation evidence underwent thematic analysis.

E. Contextualisation for AcehTtengah

Synthesised findings were then assessed for transferability to the Gayo Highlands along five dimensions: agroclimatic conditions (ambient temperature, relative humidity, solar irradiance, rainfall frequency at 1,000 - 1,700 m.a.s.l.); operational scale (batch masses of 100 - 500 kg of wet parchment, typical of an individual smallholder or a small cooperative); infrastructure (intermittent grid power, patchy cellular coverage); user capability (limited digital literacy, preference for simple visual and voice interfaces); and institutional structure (cooperatives and farmer groups as the plausible unit of ownership). This assessment distinguished universal principles, the falling rate character of the kinetics, the layered IoT architecture, the digital twin synchronisation loop from context specific requirements, and it is this distinction that generates the framework proposed in Section IV.

IV. Results and Discussion

a. Characteristics of the reviewed literature

The systematic search identified a body of literature that is substantial in each strand individually and almost empty at their intersection, a finding that is itself the central result of this review. Kinetic studies of coffee drying are numerous and methodologically mature, converging on a small set of well-validated models. IoT dryer implementations are numerous and increasing rapidly, but are technically shallow with respect to analytics. Digital twin drying studies are fewer, growing, and concentrated on grain and fruit. Studies occupying the intersection of all three were not found. Geographically, kinetic work on coffee is dominated by Brazil, Ethiopia, Colombia, and Thailand; IoT dryer work by Indonesia, India, and Latin America; digital twin work by Europe and China. Indonesian contributions are concentrated in the IoT strand and, tellingly, rarely include a kinetic model. This distribution suggests that Indonesian research has the hardware capability and the domain access but has not yet made the analytical connection which is precisely the contribution this review seeks to enable.

b. Synthesis of kinetic evidence

Across the coffee-drying kinetic studies reviewed, the modified Midilli model was identified as the best fitting formulation in the majority of reported comparisons, with the Page model consistently the strongest lightweight alternative [3], [3], [4], [3]. Reported coefficients of determination for the modified Midilli model exceed 0.99 across drying air temperatures from 40 °C to 100 °C, with root mean square (RMS) errors typically below 0.5% [3]. Effective moisture diffusivity for parchment coffee has been reported in the range of approximately 7.8×10^{-10} to 1.5×10^{-9} $m^2 s^{-1}$, increasing with air temperature and decreasing with relative humidity, and following an Arrhenius-type dependence on absolute temperature [3].

Three findings from this synthesis directly shape the proposed framework. First, the falling-rate period dominates coffee drying, confirming internal diffusion control and validating the use of a diffusion-consistent semi-empirical model rather than a constant-rate approximation [3]. Second, the drying coefficients are strongly and systematically dependent on both temperature and relative humidity, meaning that a single fitted parameter set is valid only for the conditions under which it was fitted, and that under the variable conditions of a real solar dryer, the parameters must be treated as time-varying and re-estimated online. Third, quality is a function of the drying rate and not merely of the endpoint: at a dry-bulb temperature of 35 °C, beverage quality was found to decline as the dew point fell and the drying rate consequently rose [3]. This directly justifies imposing a maximum-drying-rate constraint within the control logic an objective that a reactive thermostat has no means of representing.

Table 4. Synthesis of reported kinetic parameters for parchment Arabica

Air condition	**Best-fit model**	**Fit quality**	**Effective diffusivity ($m^2 s^{-1}$)**
50–70 °C, RH 10–30%	Modified-Midilli	Best of 9 models compared	$7.76 \times 10^{-10} – 1.45 \times 10^{-9}$
40–100 °C, hot air	Midilli	R^2 0.995–0.9995; RMSE < 0.21%	Increasing with T
Solar tunnel, 2–6 cm layers	Modified-Midilli (most zones)	Lowest RMSE and χ^2, highest R^2	Varies by zone and thickness
Gayo solar dryer, 32–45 °C, RH 55–80%	To be determined	Not yet reported	Not yet reported

c. Synthesis of iot and control architectures

Classifying the reviewed dryer systems by analytical capability produces a distribution that is stark. The great majority provide monitoring only, or monitoring with threshold-based actuation. A minority implement fuzzy-logic control, which improves the smoothness and robustness of actuation but remains a mapping from current sensed state to current actuator command [5], [5]. No reviewed coffee-drying system maintains a model of the product state, and consequently none can produce a forecast.

The practical significance of this is worth stating plainly, because it is easily obscured by the technical sophistication of fuzzy controllers. A farmer with a fuzzy-controlled dryer knows that the air inside is 38 °C and 62% relative humidity. A farmer with a predictive dryer knows that the batch will reach 12.5% moisture at approximately 14:00 tomorrow, that rain is forecast tonight, and that running the biomass burner for three hours this evening will bring the completion forward to 08:00 and reduce the risk of a stall. The second farmer can make a decision. The first can only watch a number.

Table 5. Classification of reviewed drying systems by analytical capability

Capability level	**Description**	**Product model?**	**Can forecast?**
L0 — Manual	Open-sun; judgement by feel or bite test	No	No
L1 — Monitoring	Sensors and dashboard; no actuation	No	No
L2 — Threshold control	Rule-based fan and heater actuation on setpoints	No	No
L3 — Fuzzy control	Fuzzy inference maps multiple inputs to actuation	No	No
L4 — Predictive (digital twin)	Online kinetic model; forecasts trajectory; cost-aware scheduling	Yes	Yes

d. The SIKOPI-DT Framework

Synthesising the foregoing, this review proposes SIKOPI-DT (Sistem Informasi Kopi Digital Twin), a five-layer predictive drying framework for hybrid solar–biomass coffee dryers in the Gayo Highlands. The framework is deliberately parsimonious: every design decision is constrained by the adoption barriers established in Section II-F, and where accuracy and deployability conflict, deployability wins.

1) Layer 1 Physical and sensing layer

A hybrid solar biomass dryer chamber of 100–500 kg wet-parchment capacity is instrumented with: three to four SHT31 or SHT85 sensors measuring air temperature and relative humidity at inlet, mid-bed, and outlet; DS18B20 probes measuring product-bed temperature at two depths; a calibrated photodiode or low cost pyranometer measuring incident irradiance; and the component that distinguishes this architecture from those reviewed. A load cell (HX711 amplifier) beneath a representative sample tray of known initial mass. The load cell converts mass loss directly into instantaneous moisture content, which is the observable that closes the loop between the physical dryer and its digital twin. Its cost is negligible; its analytical value is decisive.

2) Layer 2 Edge computing and connectivity layer

An ESP32 acquires sensor data at one-minute intervals and transmits to a Raspberry Pi edge node co-located at the dryer. The edge node hosts the digital twin and is authoritative: it executes model fitting, forecasting, and control locally, and functions without connectivity. Cloud synchronisation, over LoRa to a village gateway or opportunistically over cellular, is for archival, cooperative level dashboards, and remote support only, and its loss degrades no core function. This edge-first design is a direct response to the intermittent connectivity constraint.

3) Layer 3 Digital twin core (Python)

This layer is the substance of the contribution. It comprises four components. Moisture-content estimator. Instantaneous dry-basis moisture content is computed from the load-cell mass by mass balance:

$$M(t) = \frac{m(t) - m_{\mathrm{dry}}}{m_{\mathrm{dry}}} \quad (3)$$

where m(t) is the measured sample mass and m_dry the bone-dry mass inferred from the initial mass and initial moisture content. The moisture ratio then follows from Equation (1), with the equilibrium moisture content M_e evaluated from a sorption isotherm as a function of the measured air temperature and relative humidity. Online kinetic estimator. The modified-Midilli parameters (k, n, b) of Equation (2) are re-estimated at fixed intervals by non linear least squares fitting to the moisture-ratio history accumulated within the current drying campaign, implemented in Python using

Vol. 11, No. 2, July-December 2026, pp. 400-404

scipy.optimize.curve_fit over a sliding window. Because the drying air conditions vary as cloud passes and as supplementary heat is engaged, the parameters are treated as time varying rather than fixed. This continuous re parameterisation is the mechanism by which the twin remains synchronised with the physical dryer, and it is what distinguishes a digital twin from a simulation.

Forward simulator. With current parameters and a weather forecast (irradiance, ambient temperature, relative humidity, and precipitation probability, obtained from BMKG or an equivalent API and cached locally), the twin integrates Equation (2) forward to estimate the time at which the moisture content will reach the 12.5% SNI threshold and the 11% specialty target, under alternative supplementary-heating scenarios. Implementation uses NumPy and SciPy; the computation is trivial in cost and executes in well under a second on a Raspberry Pi. Predictive controller. The controller selects a supplementary heat schedule minimising a cost function of the form:

$$J = w_1 t_{\text{finish}} + w_2 E_{\text{fuel}} + w_3 P_{\text{stall}} - w_4 Q_{\text{quality}} \quad (4)$$

where t_finish is the predicted completion time, E_fuel the supplementary energy consumed, P_stall the predicted probability of a drying stall or rewetting event given the forecast, and Q_quality a quality term penalising drying rates exceeding an empirically established safe maximum. The weights w_i are set at commissioning to reflect the cooperative's priorities a cooperative facing an imminent shipment deadline weights t_finish heavily; one with fuel scarcity weights E_fuel heavily. Equation (4) is the formal expression of what a reactive controller cannot represent: an explicit trade off between speed, cost, risk, and quality. It is, in essence, the same multi-objective reasoning that Prawiranto et al. demonstrated for solar fruit drying [10], reduced to a form that runs on a fifty dollar computer.

4) Layer 4 Actuation layer

Relay controlled outputs drive the circulation and exhaust fans, the biomass burner blower, and motorised vents. The controller issues schedules rather than instantaneous commands, and every automated action is accompanied by a plain-language explanation delivered to the interface layer [11].

5) Layer 5 User interface layer

The farmer-facing interface presents a single dominant number the predicted hours remaining to target moisture content accompanied by an icon-based status indicator and, where an action is required, one plain-language recommendation in Bahasa Indonesia or Gayo, deliverable by WhatsApp message or voice note. Time-series plots, model coefficients, and diagnostic detail are relegated to a cooperative-level or technician level view. This asymmetry is deliberate and is founded on the digital-literacy constraint: the value of a prediction to a smallholder lies not in the model that produced it but in the decision it enables.

e. Projected performance

Synthesis of the reviewed evidence permits cautious, explicitly hypothesis-framed projections for the pilot, presented in Table 6. These are projections derived from the literature, not results; establishing them empirically is the object of the experimental programmed this review is intended to justify, and they are stated in falsifiable form precisely so that the pilot can refute them.

Table 6. Projected outcomes of SIKOPI-DT against the Aceh Tengah baseline (hypotheses for experimental validation)

Indicator	**Baseline (open-sun)**	**Projected (SIKOPI-DT)**	**Basis**
Drying time to 12.5% MC	7–10 days	3–5 days	Enclosed hybrid dryer performance + predictive heat scheduling [6], [7]
Intra-batch moisture SD	High; unmeasured	Reduced; continuously monitored	Bed-level sensing + layer-thickness control [6]
Specific energy consumption	n/a (solar only) / uncontrolled biomass	15–30% below thermostatic control	Predictive vs reactive scheduling [16], [17]
Rain-rewetting events per campaign	Frequent	Near zero	Enclosure + forecast-driven vent closure
Batches meeting SNI 12.5%	Variable	> 90%	Endpoint prediction replaces bite test
Fungal load at end of drying	Up to 93% infection reported in slow solar drying [6]	Substantially reduced	Shorter time above critical water activity
Kinetic model fit (R^2)	n/a	> 0.98 under Gayo field conditions	To be established; closes gap G4

f. Economic feasibility for smallholder adoption

The framework is deliberately constructed from commodity components, and its economic case rests on that parsimony. An indicative bill of materials for the instrumentation and control system ESP32, four SHT31 sensors, two DS18B20 probes, load cell with HX711 amplifier, irradiance sensor, relay board, Raspberry Pi edge node, power supply with photovoltaic panel and battery, and enclosure falls in the region of USD 250 - 400. The dryer structure itself dominates capital cost, ranging from roughly USD 800 - 1,500 for a locally fabricated hybrid solar–biomass chamber at 300–500 kg capacity. Total system cost therefore falls approximately in the USD 1,050 - 1,900 range, which is prohibitive for an individual smallholder cultivating under one hectare but well within reach of a farmer group or cooperative serving fifteen to thirty households, the same communal ownership logic that the agricultural technology adoption literature repeatedly identifies as the enabling condition for smallholder access [11], [12].

The return arises from three sources, and it is important to be candid that their relative magnitudes are uncertain and constitute part of what the pilot must establish. The first and largest is quality driven price realisation: parchment that reliably meets the SNI moisture threshold and carries a defensible low mycotoxin claim commands a premium in the specialty channel, and avoids the price penalties and rejections that non-compliant batches attract. The second is throughput: halving the drying time doubles the number of batches a dryer can process within a harvest season, which materially changes the utilisation economics of a shared cooperative asset. The third is fuel saving: predictive scheduling burns supplementary biomass only when the forward simulation indicates it will pay for itself, rather than whenever a thermostat trips. On these grounds a payback period in the range of two to three seasons under communal ownership appears plausible, and the pilot is designed to test that estimate rather than to assume it.

g. Implementation roadmap

The framework is translated into a three-phase experimental programmed, which constitutes the empirical research the authors intend to undertake and for which this review provides the theoretical foundation.

1. *Phase 1 (Kinetic characterisation and model calibration, Months 1 - 6)*

   Laboratory and bench-scale drying of Gayo parchment Arabica across the temperature and relative-humidity envelope actually achievable in a local hybrid solar dryer (nominally 32 - 50 °C, 40 - 80% RH), with gravimetric moisture determination by oven method as reference. Six thin layer models, Henderson and Pabis, logarithmic, Midilli, and modified Midilli. Effective moisture diffusivity and activation energy are determined. The deliverable is a validated kinetic core parameterised for Gayo conditions, closing gap G4 and it is worth noting that this phase alone constitutes a publishable empirical paper.

2. *Phase 2 (Digital twin construction and dryer instrumentation, Months 7 15)*

   Fabrication of the hybrid solar–biomass dryer; installation of the sensing layer; implementation of the Python digital twin on the Raspberry Pi edge node, comprising the moisture estimator, the online curve-fitting routine, the forward simulator, and the predictive controller of Equation (4). Validation proceeds by comparing predicted against measured moisture trajectories across drying campaigns spanning the full range of Gayo weather conditions, with prediction error in time-to-target as the headline metric. The deliverable is a validated digital twin, closing gaps G1, G2, and G3.

3. *Phase 3(Comparative field trial and cooperative deployment, Months 16 - 30)*

   A controlled comparison of three treatments, open-sun drying, thermostatically controlled hybrid drying, and SIKOPI-DT predictive drying replicated across batches and weather conditions, with outcome measures comprising drying time, intra batch moisture uniformity, specific energy consumption, fungal load, cup score, and full cost accounting. Deployment with a partner cooperative in Aceh Tengah, with training of champion farmers as system technicians. The deliverables are the empirical validation of Table 6 and the economic analysis closing gaps G5, G6, and G7.

h. Limitations

Three limitations of this review should be stated. First, semi-empirical thin-layer models are correlative rather than mechanistic: they describe drying behaviour under conditions resembling those in which they were fitted, and they will extrapolate poorly to conditions outside that envelope. The online re estimation strategy mitigates this but does not eliminate it, and a physics-based twin would be more robust at a cost in complexity and hardware that this user population cannot bear. This is a deliberate and defensible trade, but it is a trade. Second, the load cell measures a representative sample tray and infers the state of the batch from it; intra-batch heterogeneity is therefore partially observed, not fully observed, and the uniformity claims in Table 6 rest on this inference. Third, the projected outcomes are synthesised from studies conducted in other agroclimatic contexts and on other commodities.

V. Conclusion

This systematic review has synthesised three literatures that have developed in isolation thin layer drying kinetics of parchment coffee, IoT instrumentation of agricultural dryers, and digital twin methodology in food and grain processing and has established that their intersection is empty. Validated kinetic models capable of predicting the moisture trajectory of Arabica parchment exist, with the modified-Midilli formulation consistently identified as the best fit across studies. Low-cost IoT hardware capable of supplying such models with live data is widely deployed on coffee dryers.

Digital twin methodology capable of coupling the two is established in grain and fruit drying. Yet no reported coffee-drying system predicts anything: the field remains, without exception, reactive.

The consequence for the smallholders of the Gayo Highlands is concrete. They dry Indonesia's premier Arabica on tarpaulins, under a climate that is simultaneously ideal for growing it and hostile to drying it, and they judge its readiness by biting it. The resulting non-uniformity, prolonged drying, and mycotoxin risk cost them access to precisely the specialty premium their coffee has earned in the cup.

In response, this review has proposed SIKOPI-DT: a five layer, Python based digital twin in which a modified-Midilli kinetic core is continuously re parameterised by low cost sensor data to forecast time-to-target moisture content and to schedule supplementary heat against an explicit trade-off between speed, energy cost, stall risk, and quality. The framework is constructed to run at the edge, without reliable connectivity, on hardware costing a few hundred dollars, under communal ownership, behind an interface that reports one number a farmer can act on.

a. Theoretical contribution

The review formalises the coupling of semi-empirical drying kinetics with IoT sensing as a digital twin, advancing dryer control from setpoint regulation to trajectory forecasting. In doing so it makes a broader methodological argument: that in resource constrained agricultural settings, a lightweight correlative twin that a cooperative can actually own and operate may deliver more realised value than a high-fidelity physics-based twin that it cannot. Fidelity is not the only currency; deployability is a currency too, and the appropriate exchange rate between them is context-dependent.

b. Practical implications

For technology providers, the review specifies an implementable architecture and identifies the load cell as the low cost component whose absence has kept coffee drying IoT reactive. For cooperatives and development practitioners, it argues that communal ownership is the viable unit of adoption and that the interface, not the model, determines whether the technology is used. For policymakers, it indicates that support for highland coffee quality is more efficiently directed at post-harvest infrastructure and rural connectivity than at cultivation inputs, since the quality loss in the Gayo value chain is concentrated after the cherry leaves the tree.

c. Future research

The immediate priority is the experimental programme specified in Section IV-G: kinetic characterisation of Gayo parchment under locally achievable drying conditions, construction and validation of the digital twin, and a controlled three-treatment field trial. Beyond it, five directions follow: integration of machine-learning residual correction atop the semi-empirical core as multi-season data accumulate; extension of the twin to predict cup quality and not only moisture content; investigation of the gender dimensions of post-harvest technology adoption, given that drying labour in Gayo is substantially performed by women; life-cycle assessment of the hybrid dryer including the carbon balance of biomass utilisation; and coupling of the twin's moisture and quality records to blockchain-based traceability, converting a control artefact into a marketable provenance claim.

The transformation this implies is not, at bottom, technological. It is epistemic: a shift from knowing whether coffee is dry by biting it to knowing when it will be dry, and what to do before then. That shift is within reach of a Raspberry Pi, a load cell, and a well-chosen equation — and this

Vol. 11, No. 2, July-December 2026, pp. 404-404

review has attempted to specify, precisely enough to be built and precisely enough to be proven wrong, how they should be assembled.

References

1. A. Hidayatullah and F. A. Prasetyo, “Kajian literatur karakteristik mutu kopi Arabika Ateng Super dan pengolahan pascapanen,” RESET: Jurnal Riset Sains dan Teknologi, vol. 1, no. 3, pp. 191–208, 2026.
2. L. Wijaya, A. Karim, and W. Hasyim, “Farmers’ perception on climate change impacts on coffee farming: A case study in the Gayo Highlands,” HABITAT, vol. 35, no. 2, pp. 122–132, 2024.
3. C. Phitakwinai, S. Thepa, and W. Nilnont, “Thin-layer drying of parchment Arabica coffee by controlling temperature and relative humidity,” Food Science & Nutrition, vol. 7, no. 9, pp. 2921–2931, 2019.
4. Z. K. Abdissa, Y. Tola, A. H. Taye, and H. H. Mohammed, “Harmonizing drying time, layer thickness, and drier zones for drying kinetics: Quality and safety of solar tunnel-dried wet-processed parchment coffee (Coffea arabica L.),” International Journal of Food Science, vol. 2023, Art. no. 6677592, 2023.
5. S. Madhu et al., “Advances in renewable energy-based drying systems for food materials: A comprehensive review of innovations, smart integration, and sustainability,” Journal of Food Process Engineering, 2025.
6. G. Prasad, S. Sarkar, and L. N. Sethi, “Solar drying technology for agricultural products: A review,” Agricultural Reviews, vol. 45, no. 4, pp. 579–589, 2024.
7. A. Midilli, H. Kucuk, and Z. Yapar, “A new model for single-layer drying,” Drying Technology, vol. 20, no. 7, pp. 1503–1513, 2002.
8. E. Purlis, “Digital twin methodology in food processing: Basic concepts and applications,” Current Nutrition Reports, vol. 13, no. 4, pp. 914–920, 2024.
9. S. Panigrahi et al., “Next-generation food drying: Specialized and smart approaches to boost efficiency and quality,” Comprehensive Reviews in Food Science and Food Safety, 2025.
10. K. Prawiranto, J. Carmeliet, and T. Defraeye, “Physics-based digital twin identifies trade-offs between drying time, fruit quality, and energy use for solar drying,” Frontiers in Sustainable Food Systems, vol. 4, Art. no. 606845, 2021.
11. M. Harisudin, E. W. Riptanti, N. Setyowati, and I. Khomah, “Determinants of the Internet of Things adoption by millennial farmers,” AIMS Agriculture and Food, vol. 8, no. 2, pp. 329–342, 2023.
12. F. Hidayati, R. Syahni, I. Suliansyah, and H. B. Tanjung, “Model teoritis adopsi inovasi pertanian: Integrasi pendekatan perilaku petani,” Rawa Sains: Jurnal Sains STIPER Amuntai, vol. 15, no. 1, 2025.