IPTM: An Integrated Privacy Threat Modelling Framework for Wearable Health Devices Combining STRIDE, LINDDUN, and Privacy Impact Assessment — A Design-Science Proposal and Evaluation Protocol

Chinonso Job[1,] Onwe, Festus Chijioke[2], Okoro Favour Ngozi[3], Ndubuisi Matthew Uhuo[4], Nwojiji Titus Uchenna[5]

[1]*Department of Computer Science, University of greater Manchester, United Kingdom.*

[2]*Information Technology Department, University of Port Harcourt, Rivers State, Nigeria.*

[3]*Department of Computer Science, Ebonyi State university, Abakaliki.*

[4,5] *Department of Computer Science, Nnamdi Azikiwe University, Awka, Anambra State, Nigeria*

doi.org/10.64643/IJIRTV13I2-206664-459

***Abstract*—Wearable health devices (WHDs) sit at the intersection of security threat modelling, privacy threat modelling, and regulatory compliance assessment, yet a companion review of the literature finds no published, empirically validated method that integrates STRIDE, LINDDUN, and Privacy Impact Assessment (PIA) into a single workflow for this domain. This paper proposes the Integrated Privacy Threat Modelling (IPTM) framework to close that gap. IPTM is a design-science artefact comprising four components: a unified data-flow representation capable of treating continuous biometric telemetry as a first-class modelling object; a threat-elicitation layer that runs STRIDE and LINDDUN categories against this representation in parallel rather than sequentially; an explicit GDPR Article 22 automated-decisionmaking threat category bridging the technical and regulatory layers; and a tiered PIA-compliance scoring module that maps elicited threats directly onto specific GDPR articles. Following the principles of design-science research, we specify IPTM's architecture and worked threat-elicitation procedure in full, and we specify a concrete, falsifiable evaluation protocol — a representative WHD use case, an expert-evaluation procedure, a comparative threat-coverage metric against the three standalone frameworks, and defined usability and compliance-mapping outcome measures — by which IPTM's central claims should be tested. Consistent with its current design-stage status, no expert evaluation or comparative threat-coverage result is reported in this paper; we close with a candid assessment of which IPTM components are well supported by precedent in the literature and which rest on an as-yet-untested integration assumption.**

***Index Terms*—Privacy threat modelling, LINDDUN, STRIDE, Privacy Impact Assessment, wearable health devices, GDPR, design science research, health IoT.**

I. INTRODUCTION

A Companion critical review of the privacy threatmodelling literature for wearable health devices (WHDs) identifies four recurring gaps: structural incompleteness of existing frameworks under continuous biometric monitoring; an inadequately mapped GDPR Article 22 automated-decisionmaking threat category; persistent usability and tooling barriers limiting practitioner adoption of privacy-specific modelling; and, most significantly, the absence of any published, empirically validated integration of STRIDE, LINDDUN, and

N. S. Abafu is with the Department of Computing, Software Engineering. Manuscript prepared for submission; corresponding author email to be inserted by the author prior to submission.

Privacy Impact Assessment (PIA) for the WHD domain specifically [1]–[4]. This paper proposes the Integrated Privacy Threat Modelling (IPTM) framework as a candidate response to that gap, designed directly against the four design requirements the companion review derives from the literature.

IPTM is presented as a *design-science artefact* in the sense formalised by Hevner et al. [5] and operationalised by Peffers et al.'s design-science

research methodology [6]: its architecture and worked procedure are fully specified in this paper, but the artefact has not yet been evaluated against the expert-review and comparative-coverage protocol we specify in Section V. We state this plainly here, in the introduction, because IPTM's contribution at this stage is the design and its accompanying falsifiable evaluation protocol, not a set of expert-evaluation results.

A. Contributions

- A four-component IPTM architecture (Section III) unifying continuous-flow data representation, parallel STRIDE/LINDDUN elicitation, an explicit GDPR Article 22 threat category, and tiered PIA-compliance scoring.
- A worked threat-elicitation procedure applying IPTM to a representative WHD use case (Section IV), illustrating the integration mechanism concretely rather than only abstractly.
- A concrete, falsifiable evaluation protocol (Section V) specifying the expert-panel design, comparative threatcoverage metric, and usability and compliance-mapping outcome measures by which IPTM's central claims can be tested in follow-on empirical work.
- A candid, component-by-component assessment (Section VI) of which parts of IPTM are well supported by existing precedent and which rest on an untested integration assumption.

II. RELATED WORK

IPTM's individual mechanisms each have direct precedent in the literature, though, as the companion review establishes, not in combination. STRIDE's data-flow-diagram-based elicitation method, formalised in the security engineering literature

Continuous-Flow Data Representation (windowed DFD extension)
↓
Parallel STRIDE / LINDDUN Elicitation (per data-flow segment)
↓
GDPR Article 22 Threat Category (automated-decision flag)
↓
Tiered PIA Compliance Scoring (article-mapped severity score)

Fig. 1. Conceptual workflow of the four IPTM components. Each data-flow segment in the continuous-flow representation is elicited against both STRIDE and LINDDUN categories in parallel; any segment feeding an automated health-risk inference is additionally flagged against GDPR Article 22; all elicited threats are then scored and mapped to specific GDPR articles by the compliance module.

[7], provides the structural basis IPTM extends to continuous telemetry. LINDDUN's seven-category privacy taxonomy [3] and its lightweight LINDDUN GO variant [4] together motivate IPTM's tiered-depth elicitation design, discussed in Section III-D. Robles-Gonzalez et al.'s extension of LINDDUN′ to identification and authentication processes demonstrates that targeted, domain-specific LINDDUN extensions are a viable design pattern, which IPTM follows in constructing its continuous-flow extension [8]. Ekdahl and Nyman's data-flowbased GDPR validation methodology is the closest existing precedent for IPTM's compliance-scoring module, though it was not designed to consume STRIDE/LINDDUN threat output as its input [9]. Popoola's doctoral work combines MLbased privacy-violation prediction with blockchain-mediated consent in a related but architecturally distinct system, targeting a smart-home/health-ecosystem context rather than the WHD continuous-telemetry case IPTM targets [10]; Bugeja et al.'s PRASH framework similarly demonstrates adaptive, context-aware privacy risk scoring for smart homes, a design pattern IPTM's compliance-scoring module adapts [11].

III. THE IPTM ARCHITECTURE

IPTM comprises four components, illustrated in Figure 1 and specified in turn below.

A. Continuous-Flow Data Representation

Standard data-flow diagrams (DFDs), as used in both STRIDE and LINDDUN elicitation, represent data flows as discrete edges between processes, external entities, and data stores [3], [7]. IPTM extends this representation with a *windowed flow* construct: a data flow edge F is annotated with a window function $w(F) = (\tau,\rho)$, where τ is the nominal sampling interval of the underlying biometric signal and ρ is the retention period of any buffered window of that signal at the

receiving node. This annotation makes explicit, for every flow in the diagram, whether the data crossing it is a discrete event or a continuous stream with a defined accumulation window— directly addressing the structural-incompleteness gap identified

TABLE I IPTM PARALLEL ELICITATION TEMPLATE (ILLUSTRATIVE STRUCTURE, NOT POPULATED RESULTS)

Flow Segment	STRIDE Categories Triggered	LINDDUN Categories Triggered
On-device sensor → companion app	Tampering, Spoofing	Linkability, Identifiability
Companion app → cloud ingestion	Information disclosure, DoS	Detectability, Disclosure of information
Cloud → automated risk-scoring model	Elevation of privilege	Unawareness, Noncompliance

in the companion review, since elicitation in the next stage is then required to consider the aggregate risk of the accumulated window ρ, not only the risk of a single data point.

B. Parallel STRIDE/LINDDUN Elicitation

For each annotated flow F, IPTM requires elicitation against both the six STRIDE categories and the seven LINDDUN categories in a single pass, rather than as two sequential, independently scoped exercises. This is a procedural rather than a technical innovation: its purpose is to ensure that a flow identified as, for example, vulnerable to Tampering (STRIDE) is simultaneously assessed for whether that same vulnerability also creates a Linkability or Detectability risk (LINDDUN), which sequential application does not structurally guarantee. Table I illustrates the elicitation template.

C. GDPR Article 22 Threat Category

Any flow segment that feeds an automated inference step producing a decision or risk score about the data subject without meaningful human review is flagged with an explicit Art22-Automated-Decision marker, distinct from and in addition to its STRIDE/LINDDUN categorisation. This directly operationalises the GDPR-mapping gap identified in the companion review [2]: rather than relying on a generic LINDDUN "Non-compliance" categorisation to indirectly capture automated-decision risk, IPTM elicits it as its own category with its own mandatory mitigation requirement (meaningful human review, contestability, or explicit data-subject consent, per Article 22(2)).

D. Tiered PIA Compliance Scoring

Each elicited threat, once categorised under STRIDE, LINDDUN, and (where applicable) Article 22, is scored on a severity scale and mapped to the specific GDPR article(s) it implicates (e.g. Article 5 data minimisation, Article 25 privacy by design, Article 32 security of processing). Following the usability precedent set by LINDDUN GO [4], IPTM offers two elicitation depths: a *lightweight pass*, intended for nonspecialist development teams, which applies the template in Table I at the level of major system components; and a *full pass*, intended for compliance-critical components or formal audit, which applies the same template at the level of individual data-flow edges. This tiering directly addresses the usability-barrier gap identified in the companion review, while preserving the option of full analytical depth where compliance risk warrants it.

IV. WORKED EXAMPLE: APPLYING IPTM TO A REPRESENTATIVE WHD USE CASE

To illustrate the elicitation mechanism concretely, consider a representative WHD use case: a wrist-worn device that continuously samples heart-rate variability (HRV) and accelerometer data, streams this to a companion smartphone application over Bluetooth Low Energy, which in turn forwards aggregated windows to a cloud service running an automated arrhythmiarisk-scoring model. Applying the continuous-flow representation (Section III), the device-to-app flow is annotated with a window of $\tau \approx$ 1s sampling and $\rho \approx$ 24h on-device buffering; the app-to-cloud flow is annotated with $\rho \approx$ 30 days cloud retention. Parallel elicitation (Table I) flags the device-to-app flow for Tampering and Spoofing (STRIDE) and Linkability and Identifiability (LINDDUN), since

a 24-hour HRV buffer is itself a re-identifying behavioural signature even absent any direct identifier. The cloud risk-scoring step is flagged Art22Automated-Decision (Section III-C), triggering the mandatory human-review mitigation requirement. The compliancescoring module then maps the 30-day cloud retention finding to GDPR Article 5(1)(e) (storage limitation) and the automated risk-scoring finding to Article 22, producing a structured, article-referenced output rather than a generic risk rating. We emphasise that this worked example illustrates the elicitation mechanism; it is not a substitute for the expert-evaluated application specified in Section V, and the specific window parameters used here are illustrative rather than measured.

V. PROPOSED EVALUATION METHODOLOGY

Because IPTM has not yet been evaluated, this section specifies the protocol by which its central claims should be tested, following design-science evaluation practice [5], [6].

A. Use Case and Comparative Baselines

IPTM should be applied to at least one real or realistic WHD system (the arrhythmia-monitoring use case in Section IV is a candidate) and its threat-elicitation output compared against the output of STRIDE, LINDDUN, and PIA applied independently and sequentially to the same system, by the same analyst team, under a controlled time budget.

B. Comparative Threat-Coverage Metric

The primary outcome measure is threat-coverage recall: the proportion of a pre-established reference threat catalogue for the use case (constructed independently, e.g. by an expert panel not involved in applying any of the four methods being compared) that each method identifies. This directly tests the central claim that integration improves coverage relative to independent, sequential application, rather than assuming it.

C. Expert Evaluation Panel

Following the purposive-sampling design appropriate for specialist-population studies, a panel of healthcare IT security practitioners, privacy/data-protection officers, and WHD system designers should independently apply IPTM to the use case and complete a structured post-task evaluation covering perceived usability, perceived completeness relative to their own domain expertise, and time-to-completion relative to the comparative baselines in Section V.A. Inter-rater agreement on threat severity scoring should be reported (e.g. via Krippendorff's alpha) to assess whether IPTM's scoring module produces consistent output across analysts, addressing the assumption-inconsistency concern documented for LINDDUN by Van Landuyt and Joosen [3].

D. Compliance-Mapping Validity

Because IPTM's distinguishing claim relative to STRIDE and LINDDUN alone is its explicit GDPR article mapping, this mapping's validity should itself be independently checked by a data-protection-law specialist not involved in IPTM's design, verifying that each automated article mapping produced by the compliance-scoring module (Section III-D) is legally defensible and not merely plausible-sounding.

E. No Claim Without All Four Measures

Consistent with the practice recommended in the companion review, no claim of IPTM's effectiveness should be considered supported by future work unless results are reported against all four measures above (comparative coverage, expert-rated usability, inter-rater scoring consistency, and independently checked compliance-mapping validity); reporting threat-coverage recall alone, for instance, would not by itself validate the compliance-mapping component's legal accuracy.

VI. DISCUSSION: COMPONENT-LEVEL ASSESSMENT

Best supported by precedent:the parallel STRIDE/LINDDUN elicitation mechanism. Both individual frameworks are independently well validated [1], [12], and the procedural innovation of eliciting them in parallel against a shared flow representation, rather than sequentially, is a comparatively low-risk design choice; the open empirical question is the magnitude of the coverage improvement it produces, which Section V.B is designed to measure directly. Moderately supported: the GDPR Article 22 threat category. Azam et al.'s survey establishes the gap this component addresses [2], and Ekdahl and Nyman's data-flow-based

compliance methodology establishes that GDPR-article-level mapping from a technical data-flow model is achievable in principle [9]; what is not yet established is whether IPTM's specific mapping rules are legally robust across the variety of automated-inference architectures WHDs use, which Section V.D is designed to test.

Least supported: the continuous-flow window annotation (τ,ρ) and its downstream effect on elicitation outcomes. No cited precedent specifies or validates this particular annotation scheme; it is, at present, a design proposal motivated by the structural-incompleteness gap rather than an empirically validated modelling construct. This is the component most likely to require revision once the worked-example procedure (Section IV) is applied by an independent analyst panel rather than illustrated by the framework's own designer, and we flag it as the highest-priority target for the evaluation work in Section V.

Unaddressed by this paper: cross-jurisdictional compliance scope. IPTM's compliance-scoring module is specified against GDPR only. WHDs frequently operate across jurisdictions with materially different health-data regulation (e.g. HIPAA in the United States), and extending the compliancemapping module beyond GDPR is identified as future work rather than attempted here.

VII. CONCLUSION

This paper has proposed IPTM, a four-component designscience artefact integrating STRIDE, LINDDUN, and GDPRmapped Privacy Impact Assessment for wearable health devices, directly addressing the four gaps identified in a companion critical review of the privacy threat-modelling literature. We have specified IPTM's architecture, illustrated its elicitation mechanism through a worked example, and— consistent with its current design-stage status—specified rather than executed the evaluation protocol by which its central claims should be tested: comparative threat-coverage recall against the three standalone frameworks, expert-rated usability and inter-rater scoring consistency, and independently checked legal validity of its GDPR compliance mapping. Section VI's component-level assessment identifies the continuous-flow window annotation as IPTM's least-validated design element and therefore the appropriate starting point for that future empirical work.

STATEMENT ON VALIDATION STATUS

IPTM has not been applied by an independent analyst or expert panel, and no comparative threat-coverage, usability, or compliance-mapping-validity result has been measured at the time of writing. The worked example in Section IV illustrates the elicitation mechanism for expository purposes only and does not constitute an evaluation. No claim in this paper should be read as an empirical result for IPTM.

REFERENCES

[1] M. Aijaz, M. Nazir, and M. N. A. Mohammad, "Threat modelling and assessment methods in the healthcare-it system: A critical review and systematic evaluation," *SN Computer Science*, vol. 4, no. 6, p. 714, 2023.

[2] N. Azam, L. Michala, S. Ansari, and N. B. Truong, "Data privacy threat modelling for autonomous systems: A survey from the GDPR's perspective," *IEEE Transactions on Big Data*, vol. 9, no. 2, pp. 388–414, 2023.

[3] D. Van Landuyt and W. Joosen, "A descriptive study of assumptions made in LINDDUN privacy threat elicitation," in *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, 2020, pp. 1280– 1287.

[4] K. Wuyts, L. Sion, and W. Joosen, "LINDDUN GO: A lightweight approach to privacy threat modeling," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp.302–309.

[5] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design science in information systems research," *MIS Quarterly*, vol. 28, no. 1, pp. 75–105, 2004.

[6] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A design science research methodology for information systems research," *Journal of Management Information Systems*, vol. 24, no. 3, pp. 45–77, 2007.

[7] R. J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2010.

[8] A. Robles-Gonzalez, J. Parra-Arnau, and J. Forn´ e, "A LINDDUN-based´ framework for privacy

threat analysis on identification and authentication processes," *Computers & Security*, vol. 94, p. 101755, 2020.

[9] A. Ekdahl and L. Nyman, "A methodology to validate compliance to the GDPR," *arXiv preprint*, 2019, arXiv:1901.08387.

[10] O. J. Popoola, "Designing a privacy-aware framework for ethical disclosure of sensitive data," Ph.D. dissertation, Sheffield Hallam University, 2025.

[11] J. Bugeja, A. Jacobsson, and P. Davidsson, "PRASH: A framework for privacy risk analysis of smart homes," *Sensors*, vol. 21, no. 19, p. 6399, 2021.

[12] V. Vakhter, B. Soysal, P. Schaumont, and U. Guler, "Threat modeling and risk analysis for miniaturized wireless biomedical devices," *IEEE Internet of Things Journal*, vol. 9, no. 15, pp. 13338–13352, 2022.