Call for Papers

ACM-BCB 2027 Industry and Applied Track

18th ACM Conference on Bioinformatics,
Computational Biology, and Health Informatics
Boston, MA USA. August 10-13, 2027

Background

Health and biomedical informatics now runs in production. Diagnostic models sit inside radiology workflows, language models draft clinical documentation, remote monitoring pipelines carry data from wearables into care pathways, and population health analytics inform how payers and public health agencies allocate resources.

Most of these systems began in research settings, where controlled evaluation is essential for understanding a method and the effect of design choices on its performance. Moving a system from that setting into clinical or operational use raises a different and equally demanding set of problems: integration with clinical workflow and legacy infrastructure, performance under distribution shift, regulatory and privacy constraints, the cost of inference at scale, human oversight and escalation, revalidation over time, and the organisational work of getting clinicians to trust and use the system.

We invite submissions that demonstrate the application of systems in real-world settings, irrespective of whether the data involved is proprietary or the system is publicly available.

Topics

We invite submissions describing innovations and implementations across bioinformatics, computational biology, and health informatics that are relevant to real-world deployment. The focus is on papers that advance understanding of, and demonstrate effective handling of, practical issues arising when these technologies operate in non-trivial real-world systems: applications deployed outside controlled environments such as laboratories, classrooms, or experimental crowd-sourced setups. Applications that use established rather than state-of-the-art methods are in scope.

There is no requirement that the system be built by a for-profit company. Health systems, academic medical centres, public health agencies, non-profits, and research groups working with operational partners are all welcome.

A. System design, efficiency, maintainability, and scalability

Including but not limited to:

- Continuous maintenance, monitoring, and revalidation of deployed clinical systems
- Detecting and responding to data and population drift
- Efficient training and inference under clinical latency and cost constraints
- Enabling infrastructure for large-scale health data processing
- Handling unexpected clinician or patient behaviour
- Human-in-the-loop design, escalation paths, and clinical oversight
- Integration with electronic health records, imaging systems, and legacy infrastructure

- Negative results related to real-world deployment

B. Novel applications and use cases

Including but not limited to:

- Best practices and lessons learned from deployment programmes
- Case studies from design through clinical go-live
- Clinical decision support, triage, and recommendation systems in operational use
- Computational phenotyping and population health analytics at health systems and payers
- Description of an application or system
- Design of application-relevant clinical or biomedical datasets
- Development under constraints of model size, data availability, or compute
- Large language model applications in documentation, coding, triage, and patient communication
- Previously unaddressed problems in operational health informatics
- Wearable and remote monitoring pipelines in operational programmes

C. Methods for deployed systems

Including but not limited to:

- Bias, fairness, and equity auditing in deployed clinical models
- Interactive and clinician-facing system design
- Interpretability and explanation in clinical contexts
- Online learning and continual model updating in regulated settings
- Prospective, retrospective, and silent-mode evaluation methodologies
- Robustness across sites, scanners, vendors, and populations
- Trustworthiness, safety, and harm mitigation in deployed systems

D. Data, privacy, governance, and regulatory practice

Including but not limited to:

- De-identification, synthetic data, and disclosure risk in practice
- Federated and privacy-preserving analytics across institutions
- Health data spaces, interoperability, terminology mapping, and standards implementation
- Multimodal health data harmonisation at scale
- Regulatory validation, post-market surveillance, and audit of clinical AI
- Secure enclaves and controlled research environments

Opinion and vision papers concerning real-world deployment in health and biomedicine are also welcome.

Submission Categories

Submissions must clearly identify one of the following three categories:

Deployed Describes a system solving a non-trivial real-world problem in operational use. The paper may address the problem and its actual use cases, its significance (opportunity size,

value proposition, intended end state), method design and formulation, trade-offs in the solution, deployment challenges, and lessons learned.

Emerging Describes the development of a system solving a non-trivial real-world problem. The system need not be deployed, or close to it, but there must be evidence that development is intended for real-world deployment. Papers describing enabling infrastructure for large-scale deployment of health informatics techniques also belong here.

Discovery Reports results obtained from applying these techniques in real-world health scenarios that yield actionable insight. Such findings should point to promising directions in their application area, leading to further system or care improvements. For example, an analysis of nursing documentation may reveal that a particular charting pattern precedes deterioration, enabling an earlier alerting policy.

Evaluation and Decision Criteria

Submissions are reviewed single-blind: reviewers may see author names and affiliations, while reviewer identities are not disclosed to authors. Submissions are assessed on technical quality, practical significance, clarity, potential impact, and relevance to health and biomedical practice.

Submissions should emphasize real-world implementation, the development of systems intended for real-world use, or insight derived from real-world health data with clear practical impact. The track welcomes industrial-strength systems, commercial applications, and industry experience, including case studies, deployments, lessons learned, and problem statements. Proprietary or commercial data, results, systems, and code are welcome; authors are not required to release confidential assets. Reviewers will take legitimate confidentiality, privacy, regulatory, and operational constraints into account. Where data or code cannot be shared, authors should describe the cohort or data source, preprocessing, system or model configuration, evaluation protocol, and limitations in as much detail as is feasible. Reproducibility is valued, but inability to release protected data or proprietary code will not by itself count against a submission. Claims should still be supported by appropriate evidence, and limitations should be stated transparently.

Given the nature of the work expected in this track, authors should pay particular attention to evaluation methodology, and state clearly whether evaluation was retrospective, prospective, silent-mode, or in live clinical use, along with the population, period, and measured endpoints.

Fundamentally new machine-learning methodology is not required. Novelty may lie in the application, system or engineering design, usability, integration, business or clinical use case, dataset, interdisciplinary contribution, deployment experience, or actionable lessons learned. A well-executed, well-measured application of an established method is in scope.

Important Dates

All deadlines are 11:59pm Anywhere on Earth (AoE).

Milestone	Date
Full paper submission	March 26, 2027
Extended abstract submission	April 28, 2027
Final acceptance notification	May 28, 2027
Camera-ready	June 25, 2027
Conference	August 10-13, 2027

Submission Requirements

Originality. We invite original, full-length industry and applied track papers not previously published, accepted for publication, or under consideration elsewhere. Prior appearance as a preprint or technical report is permitted; to preserve anonymity we recommend not citing the preprint in the submission.

Submission formats and length. Authors may submit either (1) a full paper of up to 8 pages including references, will be included in conference proceedings if accepted or (2) a 2–4-page extended abstract including references, will be included in workshop proceedings if accepted. Both formats should use the double-column ACM SIG conference format and the ACM Master Article Template. Do not modify the style files or use templates from other conferences. An optional appendix may follow the bibliography for examples, preprocessing decisions, model parameters, pseudocode, user-study details, or additional error analysis; submissions must be understandable without the appendix. The extended-abstract format is intended for focused industry case studies, deployment reports, lessons learned, emerging applications, and problem statements that may not require a full-length paper.

Single-blind review. Submissions should include author names and affiliations. Reviewers may know the authors' identities, but reviewer identities will not be disclosed to authors. Authors may identify their company, institution, system, or product when relevant, while avoiding confidential information they are not authorized to disclose.

Final version and publication. Accepted submissions will receive instructions for preparing the camera-ready version, including any permitted revisions and applicable ACM publication requirements. Authors should ensure they have permission to publish any included company, product, clinical, or institutional information.

Submission system. Papers are submitted through the ACM-BCB 2027 Industry and Applied Track OpenReview submission page.

Authorship. The author list should include all and only those individuals who made substantial contributions. The author list must be final at full paper submission; subsequent changes require Program Chair approval. ACM's authorship policy applies.

Writing assistance. Use of AI-based writing assistants must comply with ACM's policies on authorship and on the use of generative AI tools.

Presentation. Accepted papers are presented orally or as posters, determined by the program committee on the nature rather than the quality of the work, with no distinction drawn in the proceedings. All accepted papers must be presented at the conference, and at least one author must register by the early registration deadline. The Industry and Applied Track runs in parallel with the research track.

Multiple Submission Policy

We will not consider any paper under review at a journal or another conference at the time of submission, and submitted papers must not be submitted elsewhere during the ACM-BCB 2027 review period. This covers refereed and archival conferences and workshops. We will not consider papers that overlap significantly in content or results with work published or to be published elsewhere.

Authors submitting more than one paper to ACM-BCB 2027 must ensure submissions do not overlap significantly in content or results. **Identical or closely related work submitted to multiple ACM-BCB 2027 tracks, for example to both the regular track and this track, will be treated as duplicate submission and rejected without review.**

Where a group of largely overlapping co-authors submits related papers, the authors must ensure mutual citation and discussion between them, and supply anonymised PDFs of the concurrent submissions as supplementary material.

Ethics Policy

Authors must comply with ACM's Code of Ethics and with ACM Publications Policies, including the **ACM Publications Policy on Research Involving Human Participants and Subjects**. Work involving patient data, clinical deployment, or human subjects must state the ethical and regulatory basis on which it was conducted, including ethics or IRB approval and consent or waiver arrangements where applicable.

Additional space beyond the page limit is allowed for an ethical considerations or broader impact statement. Such a statement is not mandatory, but papers working with sensitive health data or on sensitive clinical tasks that do not discuss these issues will not be accepted; equally, including the section does not by itself guarantee acceptance. Papers flagged for ethical concerns may receive conditional acceptance, with the camera-ready re-reviewed to confirm concerns have been addressed. Alleged violations of ACM Publications Policies will be investigated by ACM and may result in retraction and other penalties.

Contact

Industry and Applied Track Co-Chairs:

- Ping Chen, University of Massachusetts Boston
- Linsey Pang, PayPal AI

General Chairs: Wei Ding, Fusheng Wang

Frequently Asked Questions

Is this track only for people working in industry?

No. The track welcomes the entire ACM-BCB community. Anyone working on deployed or deployment-bound health informatics systems is invited to submit, and everyone is welcome to attend the sessions.

What counts as a real-world application?

A system deployed, or clearly intended for deployment, outside controlled environments such as laboratories, classrooms, or experimental crowd-sourced setups. Users are typically clinicians, patients, operational staff, or public health practitioners rather than researchers.

Can students submit?

Yes, provided the work matches this call.

My data is protected patient data and I cannot release it. Can I still submit?

Yes. Proprietary, commercial, and protected data and results are welcome. You do not need to release confidential data or code. Explain the data source and cohort, preprocessing, system or model configuration, evaluation protocol, and limitations to the extent permitted. Lack of public data release will not by itself count against the submission; claims should still be supported by appropriate evidence.

Can I submit to the regular research track instead?

Yes, the regular track is open to industry authors. This track exists for papers whose focus differs from regular research track reviewing criteria. Submitting the same or closely related work to both tracks is not permitted.

How do I choose between the two tracks?

Papers whose contribution is lessons learned, deployment experience, operational measurement, or insight from real-world health data are best suited here. Papers whose contribution is methodological novelty are better suited to the regular track. Review both calls, and contact the track chairs if you would like help deciding.

Can I submit a short industry case study?

Yes. You may submit a 2–4-page extended abstract or a full paper of up to 8 pages, including references. The short format is intended for focused case studies, deployments, lessons learned, emerging applications, and industry problem statements.
