



No study reported more than three of five.

What a systematic review of 71 published papers found about five features of digital health platforms in breast cancer and oncology.

Author: MyBC Health

Contact: Mike Haniff, PhD, MD, MSHI, MSCRA

Summary 1.2, review not yet published

Date: 2026-10-08


Summary

The team building MyBC Health carried out a systematic review of digital health platforms used in the care of adults with breast cancer and other cancers. The review asked which of five features each paper reported. The five are interoperability with the electronic health record (EHR, the hospital's record system) or through FHIR (the data standard for exchanging records); objective wearable or sensor data; toxicity surveillance with escalation to the care team; AI or clinical decision support; and data security and privacy. From 379 records, 71 papers were included: 42 reports of primary studies and 29 reviews. Toxicity surveillance with escalation was coded in 61 of the 71, data security in 25, objective wearable or sensor data in 19, EHR or FHIR interoperability in 11, and AI or clinical decision support in 7. No primary study reported more than three of the five features, and none reported all five. As run, the review was a single reviewer's mapping of PubMed abstracts, not registered in advance; it has not been peer reviewed or published. This summary gives its figures, how they were produced, and its limitations. It does not show that MyBC Health, or any platform, is effective or safe.

What was reviewed

The search ran five feature-specific queries on PubMed/MEDLINE on July 10, 2026, each covering January 2014 to the search date. The five queries matched 6,286 records in total (hits summed across the queries, not unique records). The 80 highest-ranked records of each query were exported, 400 in all; the manuscript reports 379 unique records, and the files do not record how the 400 and the 379 reconcile. Consensus semantic retrieval, ClinicalTrials.gov, and PubMed's related-article retrieval for the benchmark reviews and the highest-relevance included studies supplemented the search; no record-level provenance says what they contributed. The window starts in January 2014. HL7 published the first FHIR draft standard for trial use in February 2014, and SMART on FHIR, which runs apps on that standard, was first demonstrated the same month; those standards are what make the interoperability feature assessable.

A study was eligible if it described a digital health platform, mobile application, web portal or connected-device system used in the care of adults (18 years or older) with breast cancer, or a broader oncology population applicable to breast cancer. The system also had to implement or evaluate at least one of the five features concretely. Randomized trials, non-randomized studies, cohorts, feasibility studies, cross-sectional and mixed-methods evaluations, application feature analyses and existing reviews were eligible; reviews were tabulated separately from primary studies. Only English-language records were eligible. Excluded were protocol-only papers without results or a feature description, editorials and commentaries without primary data, genomics or imaging decision tools with no patient-facing platform, and systems that implemented none of the five features concretely. Eligibility was judged from each record's abstract and its PubMed or PubMed Central metadata; full texts were not retrieved.



Study selection, from the review's PRISMA 2020 flow: PubMed/MEDLINE, five feature searches, the 80 highest-ranked records of each, January 2014 to July 10, 2026. Eligibility was assessed from abstracts and indexed metadata, not from retrieved full texts. The record excluded at citation verification was a non-English article. Consensus (five queries) and ClinicalTrials.gov (three queries) were used to corroborate feature-level claims and describe the trial landscape; they are outside the 379.

The five features

- 1. EHR or FHIR interoperability: the platform exchanges data with an electronic health record, through FHIR, SMART on FHIR (which runs apps on FHIR) or another documented integration.
- 2. Objective wearable or sensor data: the platform captures data from a wearable or sensor. That includes any wearable-derived estimate of functional status, such as an estimate from a movement sensor of ECOG performance status (the 0 to 5 scale clinicians use for how a patient is managing day to day).
- 3. Toxicity surveillance and escalation: the platform collects symptoms or toxicity from the patient away from the clinic and escalates or alerts the care team.
- 4. AI or clinical decision support: the platform has an AI or decision-support feature. The review set out to record governance as well (validation, drift checks, a human in the loop, non-diagnostic framing). In practice the feature was scored on the presence of the component, and governance detail was noted where the record reported it.
- 5. Data security and privacy: the record reports security or privacy measures, or discusses security or privacy as a design consideration or a risk.

A feature was scored as present when the record's title, abstract and indexed metadata described it concretely. Depth and quality of implementation were not graded, and risk of bias was not assessed. An internal appraisal found codes that do not meet that bar. One security code is recorded as implied, and some toxicity codes belong to records whose abstracts describe no escalation (a lighting intervention and a nutrition and fitness app among them). The counts below are counts of codes as recorded.

Results

FEATURE	SHARE OF 71	PAPERS	PERCENT	PRIMARY (42)	REVIEWS (29)
Toxicity surveillance and escalation		61 of 71	86%	37	24
Data security and privacy		25 of 71	35%	10	15
Objective wearable or sensor data		19 of 71	27%	8	11
EHR or FHIR interoperability		11 of 71	15%	6	5
AI or clinical decision support		7 of 71	10%	4	3

Papers coded for each feature, of the 71 included, with the split between primary studies and reviews. Percentages are rounded to the nearest whole number. A paper can be coded for more than one feature, so the counts add to more than 71.

None of the 42 primary studies reported more than three of the five features.

The review's platform table has 19 rows from the primary studies; grouped by the system each describes, they cover 16 systems, four of the rows being reports of one trial. No row, and no system after grouping, combined interoperability, wearable-derived functional status, toxicity escalation, AI and documented security. The rows fall into recognizable groups: electronic patient-reported outcome systems with alerting and, often, a security description; wearable-and-app combinations that capture objective activity data; and a few EHR-integrated systems built on FHIR or mCODE (an oncology data standard). The table does not hold every primary study: 23 of the 42 are outside it, and the review does not state the rule that selected the 19. The ceiling of three features holds for each of the 42 primary studies taken on its own, not only for the rows in the table. A platform studied in several papers could have more features across them than any one paper reports.

AI or clinical decision support was the rarest feature, in 7 of the 71 papers (10%). Of those seven, the extraction notes validation for two, records no drift monitoring or human-oversight design for any, and is blank for two; a blank field does not show that the abstract reported nothing. One paper outside the seven describes a validated alert-suppression algorithm with nurse review and is not coded for this feature, so the seven under-count governance detail. Reviews and primary studies were both concentrated on symptom surveillance. Reviews were coded for security (15 of 29, 52%) and wearables (11 of 29, 38%) more often than primary studies were (10 of 42, 24%, and 8 of 42, 19%). Interoperability and AI were uncommon in both.

How to read the numbers

- The counts are papers, of 71, coded as reporting a feature. A paper could be coded for several features, so the five counts add to 123, more than 71.
- The percentages are each count divided by 71, rounded to the nearest whole number: 61 of 71 is 86%, and 7 of 71 is 10%.
- The 19 rows of the review's platform table are 19 reports of primary studies, each with the system it describes: nine rows carry a platform's proper name and seven a descriptive label. Grouped by system, the rows describe

16 systems, because four rows report one trial. The review does not state the rule that selected the 19 rows, and 23 of the 42 primary studies are outside the table.

- A feature was scored present or absent from the abstract and indexed metadata. A platform that has a feature but did not describe it there is counted as not reporting it, so the counts can under-count; a feature described in an abstract but not delivered counts the other way.
- Reviews (29) were scored alongside primary studies (42), and the table shows the split. The 71 are papers, not independent platforms or trials: reviews overlap each other and the primary papers, and several papers report one trial. The percentages describe papers, not platforms, patients or the market.
- The review did not pool effectiveness results, and it did not assess risk of bias. It reports which features papers described. Nothing here measures whether any platform works, and nothing here shows that MyBC Health, or a platform with all five features, is more effective or safer.
- MyBC Health is not one of the 71 papers. The review's platform table carries it as a twentieth row, from a developer fact sheet rather than a published study, with all five features marked. That row is the developer's description, not evidence.
- The figures in this summary come from the review's manuscript, its evidence map and its search log. The review's classifications in the evidence map are used as they stand, including the ones the appraisal questions. This summary recomputed the percentages, grouped the table rows by system, and describes the method from the review's files rather than from the manuscript's wording.

Limitations

The review was not registered in advance. One reviewer screened and extracted the papers; the check that followed resolved every identifier against CrossRef and PubMed, and did not repeat the screening or extraction judgments. The team building MyBC Health carried it out and has a commercial interest in the result; the first pass of the evidence map was produced with an AI literature tool and checked by the author. The review lists MyBC once, as a row of its platform table from a developer fact sheet. That row is not evidence that MyBC has these features.

Screening and extraction judgments were not repeated independently by a second reviewer. The search used one database, PubMed/MEDLINE, and kept only the 80 highest-ranked records of each query, so eligible records ranked lower, or indexed elsewhere, were never seen; a planned OpenAlex sweep could not be completed. Eligibility and features were judged from abstracts and indexed metadata, not from full texts, so the counts can err in either direction. Risk of bias was not assessed, and no synthesis of effectiveness was performed. Only English-language records were included. An internal appraisal of the manuscript found nine included records whose recorded populations need checking against the criteria, three of them clear mismatches: a COVID-19 study, a study of patients after atrial fibrillation ablation, and a pediatric study. It found six study-design labels that disagree with PubMed's publication types and two more that disagree with the papers' titles, and three records that are protocol or design papers, so the 42 and the 29 may move. It found fifteen records whose population, platform and sample-size fields are generic and whose outcome field is a sentence from the abstract; all fifteen count toward the 61. The author has yet to adjudicate these, and the counts above may change when that is done. The review has not been peer reviewed or published, and this summary will be revised when it is.

What MyBC Health is building in each of the five areas

This section describes design intent. It is not evidence, and it is not part of the review. MyBC Health is in development and is not yet available to patients. A feature in the current build is described in the present tense; a designed feature that is not yet in the build is described as being built. Statements about availability, EHR connections and regulatory status are the company's own, not findings of the review.

1 Toxicity surveillance and escalation

A patient logs a symptom from the phone or the watch, and the entry is time-stamped when it is submitted. The CareDay screen has quick buttons to log a symptom or message the care team. MyBC records symptom answers in PRO-CTCAE, the patient-reported measure of symptomatic side effects that oncology trials use, so results can reach the chart as data. Per-patient alert thresholds and escalation to the care team are being built. Whether and when a clinician reads a report depends on the clinic; the app is being built to tell the patient when reports are not being monitored, until a named reader is assigned and outside the clinic's coverage hours.

2 Data security and privacy

One identity, one audit trail and one permission model across every surface is the design. Provider access and record changes are written to an audit log, and emergency (break-glass) access is recorded as a distinct event for escalation. Withdrawing permission updates the consent record and queues the clean-up of shared data. A patient chooses what is shared, by category of data, and who else sees the record. MyBC does not sell patient data and does not use identifiable health data to train general-purpose AI models.

3 Objective wearable or sensor data

Steps, heart rate and sleep come in from Apple Health on iPhone and from Health Connect on Android, when the patient grants access. Blood oxygen, heart-rate variability, sleeping wrist temperature and sleep stages are read from Apple Health when a watch records them. W-BCOG estimates ECOG performance status from a 14-day rolling window of wrist and patient-reported signals. The signals and their weights are steps (35%), non-sedentary hours (25%), fatigue (20%), sleep efficiency (10%), resting heart rate (5%) and the patient's own rating (5%); when an input is missing, the others are reweighted. The default scorer is rule-based and deterministic (version rule-based-ts-1.0.0); the code also allows an optional machine-learning scorer, which is not used by default and falls back to the rules. The scorer needs at least seven days in the window with 60% or more of the expected data, and otherwise produces no score. W-BCOG is an investigational estimate, not a validated substitute for a clinician's ECOG assessment, and has not been cleared by any regulator.

4 EHR or FHIR interoperability

MyBC is being built on SMART on FHIR R4 (the current FHIR release) to read and write Patient, Observation, MedicationRequest, Condition and Encounter, with USCDI v3 (the U.S. core data set) as the coverage target. Today its server defines a CDS Hooks catalog, the standard for showing advice inside a clinician's record system, with services on patient-view, medication-prescribe and order-sign, and a SMART launch that requests Patient and Condition reads. There is no live EHR connection today. Live write-back is not implemented: the transport, credentials and endpoint configuration are still to be built, and it also waits on executed business associate agreements.

5 AI or clinical decision support

The assistant answers from the patient's own record for questions such as what medication they are on or when the next appointment is. For questions such as what a lab test measures, it draws on general sources, including guidelines and drug information. Three layers are designed to sit between a question and an answer. An intent router is in place. Identifier scrubbing today covers six identifier types in the web portal (Social Security number, medical record number, phone, email, date of birth and street address), with a server-side field blocklist; coverage of all 18 identifier types in the HIPAA Safe Harbor list is being built. NVIDIA NeMo Guardrails is wired in for hallucination detection, topical-drift and prompt-injection defense; it is off by default while it runs in a logging-only mode. A TypeScript output guard screens responses for prescribing, diagnosing and medication-change language and offers to flag such a question for the care team; it is one layer, not a guarantee. For clinicians, a treatment-plan recommendation shows its NCCN (National Comprehensive Cancer Network) guideline citation and evidence category where supplied, and one that needs review is approved, declined or modified, with the rule and version shown for audit. MyBC is not designed to prescribe, diagnose or change a medication.

W-BCOG is investigational. No EHR connection is live. MyBC Health has not been cleared or approved by any regulator. MyBC is not a substitute for professional medical advice, diagnosis or treatment.

Contact

MyBC Health

mybchealth.com

Mike Haniff, PhD, MD, MSHI, MSCRA

Clinics, researchers and partners:

partners@mybchealth.com

Summary 1.2, 2026-10-08

General questions: support@mybchealth.com

This summary reports an unpublished systematic review by the team building MyBC Health. It cites no third-party study by name. When the review is published, a later version of this summary will carry its citation.