Patient.md: Framework to Organize Medical Data for AI Assistants

Clarence C. Hu[1] Juan J. Cardona[2] Karel Lee[1] Ali Danish[1]
Devanish N. Kamtam[3]

[1] Hotpot.ai, Palo Alto, California, USA.
[2] Department of Neurosurgery, Stanford University School of Medicine, Stanford, California, USA.
[3] Division of Thoracic Surgery, Department of Cardiothoracic Surgery, Stanford University School of Medicine, Stanford, California, USA.

Corresponding author: clarence@hotpot.ai

Abstract

Background. Poorly understood conditions and specialist shortages may compel patients to take agency over their care, yet seeking second opinions and exploring treatment options can be challenging. This journey is often slowed by scattered medical records, repeated explanations of the same history, and volumes of clinical data that can overwhelm both human reviewers and AI assistants.

Methods. We propose Patient.md, a framework to organize medical data for artificial intelligence (AI) assistants. Patient.md defines a schema for representing cases as a single Markdown file comprising standard clinical details and condition-specific data such as molecular testing in cancer care. An optional registry links summaries to patient records for traceability and efficient retrieval in AI assistant workflows. The framework enables tailored files for different reviewers, supports local or cloud operation, and provides a system prompt for drafting Patient.md files via guided workflows.

Results. We demonstrate Patient.md with a lung cancer patient sharing three case versions: minimal, public facing for trial matching, and detailed for deeper investigation.

Conclusion. By consolidating medical data into one structured file designed for AI assistant context windows, Patient.md can simplify case sharing for second opinions and promote patient autonomy. It may further enable AI assistants to act as personalized tutors, lowering barriers to case understanding and treatment exploration for both physicians and patients. This structured format may also lay a foundation for AI agent workflows such as case monitoring and clinical trial matching.

1 Introduction

Patients with complex or poorly understood conditions may seek second opinions and explore treatment options on their own. In practice, this journey is often slowed by scattered medical records, repeated explanations of the same history, and volumes of clinical data.

Electronic health record (EHR) exports, while increasingly available, are not designed for efficient review or integration with artificial intelligence (AI) assistants. When new clinicians or patient advocates engage with a case, they may need to navigate vast amounts of information with no way to personalize review or ask targeted questions. Patients may also wish to share select information with different reviewers – full records with a specialist but only a restricted set with the public, for example. Together, these barriers to understanding and sharing cases undermine patient autonomy.

AI assistants, particularly the large language models (LLMs) behind modern AI assistants like ChatGPT, offer a potential bridge for synthesizing and reasoning over health records. These systems can excel at analyzing data but operate within finite "context windows" – the amount of information they can process at once [1]. Responses may miss or misinterpret details when key facts are buried in long documents or dispersed across records [1]. AI assistants perform best when relevant evidence is compactly organized and traceable to source files. Patient-facing exports, however, are often overly verbose or insufficiently structured.

2 Related Work

Current approaches to organizing patient data focus on machine-to-machine exchange or human-to-human handoff rather than joint human-AI investigation. The International Patient Summary (IPS) defines a minimal summary intended for unscheduled cross-border care, emphasizing a safety-oriented snapshot over comprehensive depth [2]. Similarly, HL7 Clinical Document Architecture (CDA) is envisioned for transferring clinical documents between institutions [3]. Patient-facing initiatives such as Blue Button broaden access to health records, but exports remain verbose and unoptimized for AI assistants [4].

In AI workflows, retrieval-augmented generation (RAG) is a widespread technique for grounding answers in external knowledge. However, frameworks focus on ingestion and indexing rather than exposing a lightweight, user-centric layer for curating rich summaries and linking to source records [5]. The AGENTS.md standard demonstrates how a lightweight Markdown file can steer AI assistants through vast project repositories, but is not intended for processing and structuring the health records of individual patients [6]. Other related AI works, including ChatEHR, CLI-RAG, and EHRStruct, focus on healthcare but primarily target developers or physicians, not patients and the structuring of their data [7–9].

3 Methods

3.1 Patient.md Framework Overview

To address these gaps, we propose Patient.md, a framework for organizing medical data and facilitating case review and treatment exploration through AI augmented workflows. Patient.md represents each case as a structured packet: a single Markdown file, or manifest, plus optional files containing patient data. This manifest contains key clinical summaries and an optional registry of links to underlying patient data either packaged with the manifest or located in repositories. Core sections of the manifest cover standard details

such as diagnosis, providing a consistent entry point for review and collaboration, while additional modular sections fulfill case-specific requirements. These section-level fields balance structure with free-text flexibility, allowing information depth to scale with case complexity.

The manifest may function as a compressed representation of the patient state with optional links to source data for traceability and citation. This dual-layer architecture enables Patient.md to serve as a RAG index, optimizing context-window usage for large cases: core sections supply clinical context while the optional registry preserves access to underlying records.

Although many AI assistants are LLM-based, Patient.md is designed to remain architecture-agnostic, specifying a clinical structure and retrieval interface for both LLMs and alternative paradigms. Two additional principles, usability and extensibility, guided development. Usability informed choices such as single-file integration with AI assistants like ChatGPT and Gemini, creation and editing without specialized software, and flexible formats tolerant of incomplete entries and varied writing styles. To reflect the evolving nature of medical research and diverse needs of patients and providers, modular components allow for framework expansion without schema updates.

3.2 Clinical Data Layer

Patient.md recommends Fast Healthcare Interoperability Resources (FHIR) as the clinical data layer to support widely adopted EHR exports and consistency with international health standards [10]. In addition, the framework accommodates clinical artifacts such as PDFs and imaging studies when structured exports are unavailable, offering flexibility for heterogeneous real-world data.

3.3 Patient.md Core Sections

Table 1 summarizes the core sections of Patient.md version 1.0.0. These were guided by the IPS, but its divergent objectives motivated our modifications.

Modules impart extensibility while avoiding subjective taxonomies. In practice, medical information often resists clean categorization since regional practices may vary and guidelines change. As an example, PD-L1 may be classified as a pathology result or an immunotherapy biomarker. Likewise, liquid biopsy results may be treated as blood tests or genomic assays. Modules let data organization fit the case while linking summaries to underlying data for traceability and deeper analysis. A cardiovascular patient may not need molecular testing modules, for instance, whereas a lung cancer patient may require several [11]. This adaptability accommodates both routine clinical data and emerging research modalities. In oncology, for instance, common modules might add sections for genomic sequencing panels and liquid biopsy, imaging reports, or clinical trial matching. Research-integrated cases might add modules for single-cell phenotyping or gene set enrichment analysis.

How to Help captures the patient's priorities and open questions, giving reviewers immediate orientation to actionable goals. Ruled Out documents excluded hypotheses, reducing redundant investigation. The Communication section reflects a potential need in rare

and complex diseases, where patients may share case data publicly or semi-publicly to invite broader collaboration via disease-specific forums or websites. For these cases, Communication provides a structured way to direct potential collaborators to appropriate channels while maintaining boundaries around sensitive data.

The Registry section consolidates underlying case data such as EHR exports and imaging reports in one location. Each source record is assigned a stable name and associated with a storage location and a clinical date reflecting the patient state at the time. Elsewhere in the manifest, entries may cite one or more Registry entries, enabling reviewers and tools to locate and analyze the evidence behind any claim.

Section	**Description**
Patient	Demographics and identifiers.
Diagnosis	Primary and secondary diagnoses with staging/grading where applicable.
Status	Current clinical state, performance status, active symptoms.
Timeline	Chronological narrative of disease course and key events.
Treatments	Past and current therapies, procedures, medications.
How to Help	Specific questions or decisions the patient seeks input on.
Ruled Out	Negative findings and excluded diagnoses, treatments, or hypotheses.
Modules	User-extensible sections for case-specific data such as molecular reports.
Registry	Registry of patient data.
Communication	Ways to communicate and links to collaboration spaces.
Metadata	Manifest versioning and update timestamps.

Table 1. Patient.md v1.0.0 core sections. Patient.md core manifest sections are listed. For each section, the intended clinical scope and role in organizing a portable case packet for review and collaboration are described.

3.4 Patient.md Field Definitions

Table 2 summarizes the field definitions of Patient.md version 1.0.0. Definitions balance structure with narrative flexibility, allowing authors to match information density to case complexity. For instance, ethnicity is included because it may inform pharmacogenomic considerations, ancestry-linked disease prevalence, and eligibility for ancestry-stratified clinical trials.

Section	**Field**	**Notes**
Patient	-	Use privacy-safe values when appropriate.
	name	Report display name or pseudonym.
	dob	Report birth date when appropriate. For privacy-sensitive manifests, age in years may be reported instead.
	sex_at_birth	Report as male / female / other / unknown.
	ethnicity	Report ethnicity.
	other	Report other clinically relevant details.

Section	Field	Notes
Diagnosis	-	List each diagnosis, one per row.
	name	Report diagnosis name.
	date	Report date of diagnosis.
	stage	Report staging, severity, or classification system used.
	details	Describe histology, subtype, and diagnostic criteria as relevant.
Status	-	Describe current condition, including symptoms, functional state, active problems, recent clinician assessment, near-term plan.
Timeline	-	List key events in ascending date order, one per row. Minimize duplication by reserving detailed findings for other sections, adding only what is needed to disambiguate events.
	date	Report when the event occurred.
	event	Summarize what happened.
Treatments	-	List current and prior therapies, one row per regimen or procedure. Use separate rows for rechallenge or material regimen changes.
	name	Report drug, regimen, or procedure.
	dates	Report start date and end date if stopped.
	details	Describe intent, dosing, response, toxicity, and reason stopped as relevant.
How to Help	-	Describe patient priorities and how others can help.
	goals	List patient goals, one per row, such as clarifying the diagnosis or comparing treatment options.
	questions	List open questions in priority order, one per row.
	constraints	List constraints, one per row.
	general	List general ways to help.
Ruled Out	-	List excluded diagnoses, treatments, or hypotheses, plus negative findings, one per row.
	item	Report the excluded diagnosis, test, treatment, or hypothesis.
	rationale	Explain the reason for exclusion.
Modules	-	List extension modules, one per row.
	name	Assign a stable name unique within the Modules section.
	summary	Describe what information the module adds, including key findings and other details.
Registry	-	List source data, one per row. Source names must be unique.

Section	Field	Notes
	name	Assign a name unique within the Registry section.
	summary	Summarize source type (e.g., pathology report, clinician note). Optional.
	format	Report file format or container (e.g., pdf, json, folder, zip). Optional.
	uri	Report file path or URL.
	clinical_date	Report when the patient state was captured, such as the scan date or collection date.
	institution	Report source institution.
Communication	-	List channels for care team interaction or advocate collaboration, one row per entry.
	name	Report the channel name.
	summary	Describe when and how to use this channel, and indicate if it represents the primary channel.
Metadata	-	List manifest metadata entries, one per row.
	patientmd_version	Report schema version.
	updated_at	Report last modification date.
	prompts_used	List prompts used to generate this manifest: full text, prompt links, or repository links. Optional.

Table 2. Patient.md v1.0.0 field definitions. Patient.md section-level descriptions and field definitions are listed.

3.5 Patient.md Usage Guidelines

3.5.1 Creation

This section outlines how the manifest is created, whether manually, tool assisted, or fully generated. Patient.md is intentionally low-burden: a minimal packet includes only the Patient and Diagnosis sections.

To create the manifest, authors may start by manually populating Patient and Diagnosis, adding other sections as needed. Alternatively, authors may draft the manifest with AI assistants by supplying the Patient.md specification and clinical data. Because AI responses may contain errors, drafts must be verified carefully. Supplementary Figure S1 provides a reference system prompt for drafting Patient.md files via guided workflows. The prompt produces a draft manifest plus a concise review aid that flags missing or conflicting items. Given the sensitivity and importance of personal health information, the prompt also offers a privacy review and includes a reminder not to rely on AI for medical advice.

Each section may be rendered as a list, table, narrative, or hybrid – the requirement is semantic coverage of fields, not rigid syntax. Manifests are organized into several top-level sections using Markdown level-1 headers, in which each section title is preceded by a single "#" and a space. Capitalize section headers. Field names may be omitted when

unambiguous, but when present, field names use lowercase snake_case unless rendered alone on a line ending in a colon, in which case Title Case is allowed. Field values use normal capitalization. End narrative or multi-sentence entries with punctuation. List-style, fragment-style, or key-value entries may omit terminal punctuation. Sections containing multiple entries should use one item per row if presented as a table or list. To cite patient data from the Registry section, sections may include a "sources" field once per section or per entry as needed. Sections may be listed in any order.

Registry and Module names are limited to alphanumeric characters, spaces, hyphens, underscores, and periods. Cross-references use the prefix "registry:" or "module:" followed by an identifier formed by lowercasing the target name and replacing runs of spaces and hyphens with a single underscore (e.g., Pathology 2024-03 becomes `registry:pathology_2024_03`, Molecular Markers becomes `module:molecular_markers`). Registry and Module names must be unique within their respective sections to guarantee unambiguous cross-referencing. Use ISO 8601 dates (YYYY-MM-DD) when possible, and partial dates (YYYY-MM or YYYY) when precise dates are unavailable [12]. Record unknown dates as "unknown". For "sources" fields with multiple values, separate entries with a comma.

3.5.2 File Conventions

`PATIENT.md` is the manifest filename, consistent with `AGENTS.md` and `README.md` conventions. By default, the manifest sits at the packet root so relative paths remain stable when the packet is moved or distributed.

3.5.3 Portability and Distribution

The framework supports both local and cloud operation with the same manifest semantics. The Registry section is storage-agnostic. Each entry captures a stable name paired with a resolvable location, such as a local path, private network path, or public URL. This enables mixed deployments, where some source records are packaged with the manifest and others are referenced externally.

To enable differential sharing, a case packet can be distributed as a folder or zip archive, or as the `PATIENT.md` manifest alone in the simplest case. This allows authors to release multiple packets with varying subsets of patient data for different reviewers.

3.5.4 AI Assistant Integration

Integrating Patient.md with AI assistants involves uploading a case packet. Once integrated, Patient.md supports multiple interaction patterns.

In retrieval mode, AI assistants can treat the manifest as the routing layer: they locate relevant sections and linked source records, load only the minimum artifacts needed, and synthesize a response. When available, AI assistants may cite source records to support claims and preserve traceability.

In tutoring mode, AI assistants can walk users through sections progressively, using summaries for orientation and retrieving supporting artifacts when a claim merits verification or

deeper explanation. This pattern supports onboarding for new clinicians, patient advocates, or family members, with depth and terminology adapted to the user's background.

These patterns can combine. A tutoring session may surface a focused question that triggers targeted retrieval, while a retrieval query may reveal gaps that prompt a broader walkthrough of adjacent sections or modules.

4 Results

To illustrate the Patient.md framework, Figures 1–3 present three scenarios involving the same fictitious patient. Jane Doe is a 28-year-old, non-smoking, Chinese American female with clinical stage 3A non-small cell lung cancer (NSCLC). These examples predominantly use narrative formatting to reflect the design principles of usability and intuitive workflows.

```
# Patient
- Jane Doe
- 28 years old
- Female
- Chinese American
- Never-smoker

# Diagnosis
- Non-small cell lung cancer (NSCLC), adenocarcinoma
- Diagnosed 2025-10-18
- Clinical stage 3A (cT2aN2M0)
- Right upper lobe primary, 3.2 cm; EGFR exon 19 deletion, PD-L1 5%, TMB-low
(2.1 mut/Mb)
```

Figure 1. Minimal manifest. The smallest valid Patient.md manifest, containing only the Patient and Diagnosis sections.

```
# Patient
- Jane Doe
- 28 years old
- Female
- Chinese American
- Never-smoker

# Diagnosis
- Non-small cell lung cancer (NSCLC), adenocarcinoma
- Diagnosed 2025-10-18
- Clinical stage 3A (cT2aN2M0)
- Right upper lobe primary, 3.2 cm; EGFR exon 19 deletion, PD-L1 5%, TMB-low
(2.1 mut/Mb)

# Status
```

```
- Newly diagnosed
- ECOG 0
- Asymptomatic, incidentally detected on chest imaging obtained for rib pain
after a fall
- Completing staging and treatment planning:  PET/CT and mediastinal staging
completed; brain MRI scheduled for 2025-11-07; surgical + radiation oncology
consultations

# How to Help
Goals:
- Identify initial treatment strategy
- Understand role of surgery vs. definitive chemoradiation in this setting
- Find clinical trials for EGFR+ locally advanced disease

Questions:
- Should I pursue trimodality therapy (chemoradiation + surgery) or
definitive chemoradiation?
- If surgery is feasible, is there a role for EGFR-targeted therapy before or
after surgery?
- Are there trials combining EGFR TKIs with radiation?

Constraints:
- Clinical trials must be within 100 miles of the San Francisco Bay Area
- Prefer trials at academic medical centers

General:
- Friends & family:  Please distract me by sending food photos and
recommendations of new restaurants you think I might like :)
- All:  Please share your favorite bread recipes

# Communication
- Email:  janedoe.nsclc@gmail.com.  Primary channel for inquiries and private
discussion.
- X: @janedoe_nsclc.  Public updates and connecting with EGFR+ community.
- Reddit:  /r/janedoe-nsclc.  Private subreddit for case discussion.  Request
access via email.

# Metadata
- patientmd_version:  1.0.0
- updated_at:  2025-11-05
```

Figure 2. Simple public-facing manifest. A lightweight Patient.md manifest designed to facilitate treatment tracking and trial matching while setting geographic constraints for clinical trial participation. For privacy, the birthdate is omitted and replaced with age.

```
# Patient
- Jane Doe
- Born 1997-06-28
- Female
- Chinese American
```

- Never-smoker

Diagnosis
- Non-small cell lung cancer (NSCLC), adenocarcinoma
- Diagnosed 2025-10-18
- Clinical stage 3A (cT2aN2M0)
- Right upper lobe primary, 3.2 cm; EGFR exon 19 deletion, PD-L1 5%, TMB-low (2.1 mut/Mb)
- Single-station N2 pathologically confirmed by EBUS (4R positive)
- Sources: registry:pathology, registry:tempus_xt, registry:pet_ct

Status
- Newly diagnosed and treatment-naive
- ECOG 0
- Asymptomatic, no neurologic symptoms
- Current assessment: clinical stage 3A EGFR+ adenocarcinoma with single-station N2 (4R) involvement and no extrathoracic metastases on PET/CT
- Near-term plan: complete brain MRI; finalize induction vs definitive chemoradiation strategy and pursue second opinions; trial eligibility screening in parallel
- Sources: registry:pathology, registry:pet_ct, registry:tempus_xt

Timeline
- 2025-09-28: Presented to urgent care for musculoskeletal rib pain after a fall; chest X-ray showed an incidental right upper lobe mass
- 2025-10-01: CT chest characterized a 3.2 cm right upper lobe mass
- 2025-10-10: PET/CT staging performed; no extrathoracic disease detected
- 2025-10-14: CT-guided biopsy of the right upper lobe mass performed
- 2025-10-18: Pathology and mediastinal staging finalized
- 2025-10-28: Tempus xT molecular profiling report issued
- 2025-10-30: Tumor board discussed management options
- 2025-11-03: Thoracic surgery consultation completed
- 2025-11-07: Brain MRI scheduled

How to Help
Goals:
- Determine initial strategy for single-station N2 EGFR+ stage 3A NSCLC
- Evaluate how EGFR-targeted therapy should be sequenced with local therapy in this setting
- Identify clinical trials for EGFR+ locally advanced disease within local travel constraints

Questions:
- For single-station N2 EGFR+ disease, when does surgery add value over definitive chemoradiation with EGFR-targeted consolidation?
- Is neoadjuvant osimertinib (± chemotherapy) reasonable outside of a trial in stage 3A, or should it be trial-only?
- Are there trials integrating EGFR TKIs with radiation or multimodality therapy for stage 3 disease?
- What is the recommended approach to CNS surveillance at baseline and during therapy in EGFR-mutant disease?

Constraints:
- Clinical trials must be within 100 miles of the San Francisco Bay Area

```
- Prefer academic medical centers (UCSF, Stanford)

General:
- Friends & family:  Please distract me by sending food photos and
recommendations of new restaurants you think I might like :)
- All:  Please share your favorite bread recipes

# Ruled Out
- No extrathoracic metastatic disease on PET/CT; sources:  registry:pet_ct
- Additional actionable drivers beyond EGFR exon 19 deletion:  none reported
on Tempus xT; sources:  registry:tempus_xt
- Immunotherapy-first strategy:  not favored in EGFR-mutant disease

# Modules
## Molecular Markers
- EGFR exon 19 deletion with low PD-L1 (5%) and low TMB (2.1 mut/Mb); no
high-risk co-alterations highlighted on the report
- Implications:  EGFR-directed strategies are central; checkpoint inhibitor
benefit is less likely
- Sources:  registry:tempus_xt

## Pathology
- Lung adenocarcinoma confirmed on primary tumor biopsy.  EBUS confirmed
single-station N2 with 4R node positive.
- Implications:  stage 3A management hinges on resectability and induction vs
definitive chemoradiation strategy
- Sources:  registry:pathology

## Imaging
- PET/CT shows FDG-avid RUL primary and 4R node without extrathoracic
metastatic disease
- Sources:  registry:pet_ct

# Registry
## Pathology
- summary:  Pathology report bundle (CT-guided lung biopsy + EBUS cytology)
- format:  pdf
- uri:  gs://janedoe-nsclc/pathology/pathology_bundle_2025-10-18.pdf
- clinical_date:  2025-10-18
- institution:  UCSF

## Tempus xT
- summary:  Molecular profiling report (Tempus xT); specimen collected
2025-10-14
- format:  pdf
- uri:  gs://janedoe-nsclc/molecular/tempus_xt.pdf
- clinical_date:  2025-10-14
- institution:  Tempus Labs

## PET-CT
- summary:  PET/CT radiology report
- format:  pdf
- uri:  gs://janedoe-nsclc/imaging/pet_ct_report_2025-10-10.pdf
- clinical_date:  2025-10-10
- institution:  UCSF

# Communication
```

```
- Email:  janedoe.nsclc@gmail.com.  Primary channel for professional
inquiries, second opinions, and access requests.
- X: @janedoe_nsclc.  Public updates and connecting with the EGFR+ community.
- Reddit:  /r/janedoe-nsclc.  Private subreddit for case discussion; request
access via email.

# Metadata
- patientmd_version:  1.0.0
- updated_at:  2025-11-05
- prompts_used:  https://github.com/HotpotBio/Patient-md
```

Figure 3. Detailed manifest. A more comprehensive Patient.md manifest shared with experts from Jane's social network for collaborative review on a private subreddit (/r/janedoe-nsclc). The manifest cross-references key claims to source records, including Tempus and pathology reports.

5 Discussion

Poorly understood conditions and specialist shortages increasingly compel patients to take agency in understanding and managing their care. The Patient.md framework addresses this need by consolidating fragmented medical information into portable case packets designed for AI assistant context windows. By structuring data into a single file, Patient.md may simplify case sharing for second opinions and reduce barriers to case understanding. It may further promote patient autonomy by enabling AI assistants to act as personalized tutors when physicians and patients explore treatment options or ask targeted questions. This structured format may also support AI agent workflows such as case monitoring and clinical trial matching.

Naming choices symbolize our philosophy of AI augmentation. While the manifest name `PATIENT.md` follows computing convention for root-level files like `README.md`, the framework name *Patient.md* adapts to human convention. Uppercase filenames, though standard as root files in software repositories, are jarring in natural language. PascalCase denotes an object of significance while aligning with expectations for a framework name, reflecting our emphasis on human-centric design. We selected "augmentation" over "assistance" deliberately: assistance implies facilitating tasks one should perform naturally, whereas augmentation signals expanded capability.

Future work should explore native Patient.md support in AI assistants. Analogous to how calendar links trigger event auto-population, Patient.md uploads could activate workflows for health tracking, case review, and treatment exploration. AI agents could use manifests to monitor cases and find clinical trials. Because Patient.md encodes details in a structured, machine-readable format, agents could detect status changes, cross-reference patient profiles against trial registries, and track PubMed for newly relevant evidence. Further, one-click integration from medical systems would reduce authoring burden, improve manifest consistency, and foster adoption. Cryptographically secure differential sharing could improve privacy controls across disparate groups. Finally, developing modules for domains such as cancer, where data is multimodal and complex, may accelerate clinical utility and

community adoption.

6 Limitations

Several limitations affect Patient.md. First, AI outputs do not constitute medical advice and may contain errors, requiring clinical validation [13, 14]. Second, the framework demands a baseline of medical literacy, whether professionally established or acquired via intensive study, in order to render beneficial judgments and avoid harmful interpretations [15]. The framework further relies on the quality of machine-generated or user-provided data. Inaccurate or omitted entries may induce misleading, dangerous, or otherwise erroneous responses. Although Patient.md supports differential sharing, this stems from a packaging configuration rather than a security mechanism. Once a packet is shared, the framework cannot prevent redistribution or leakage of patient data.

Data and Code Availability

Demos of Patient.md as a Claude skill, custom GPT, or Gemini gem can be found at `https://github.com/HotpotBio/Patient-md`. This repo also maintains reference prompts and community prompts.
Data and Code Availability Demos of Patient.md as a Claude skill, custom GPT, or Gemini gem can be found at https://github.com/HotpotBio/Patient-md. This repo also maintains reference prompts and community prompts for Patient.md.

Conflicts of Interest

The authors have no conflicts of interest to declare.

Funding

This work was supported by Hotpot.ai.

References

[1] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics (TACL)*, 12:157–173, 2024.

[2] CEN (European Committee for Standardization). EN 17269:2019 – Health informatics – The International Patient Summary, 2019.

[3] Robert H. Dolin, Liora Alschuler, Sandy Boyer, Calvin Beebe, Fred M. Behlen, Paul V. Biron, and Amnon Shabo. HL7 Clinical Document Architecture, Release 2. *Journal of the American Medical Informatics Association*, 13(1):30–39, 2006.

[4] Assistant Secretary for Technology Policy / Office of the National Coordinator for Health IT (ASTP/ONC). Blue Button: Patient access to health records. Web resource, accessed 2026-02-09.

[5] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[6] AGENTS.md Contributors. AGENTS.md: A simple, open format for guiding coding agents. Project specification, accessed 2026-02-09.

[7] Nigam H. Shah, Nerissa Ambers, Abby Pandya, Timothy Keyes, Juan M. Banda, Srikar Nallan, et al. Adoption and Use of LLMs at an Academic Medical Center. arXiv preprint arXiv:2511.08206, 2025.

[8] Garapati Keerthana and Manik Gupta. CLI-RAG: A Retrieval-Augmented Framework for Clinically Structured and Context Aware Text Generation with LLMs. arXiv preprint arXiv:2507.06715, 2025.

[9] Xiao Yang, Xuejiao Zhao, and Zhiqi Shen. EHRStruct: A Comprehensive Benchmark Framework for Evaluating Large Language Models on Structured Electronic Health Record Tasks. arXiv preprint arXiv:2602.00074, 2026.

[10] HL7 (Health Level Seven International). FHIR Release 4 (R4), Version 4.0.1, 2019.

[11] Vicki Leigh Keedy, Sarah Temin, Mark R. Somerfield, Mary Beth Beasley, David H. Johnson, Lisa M. McShane, Daniel T. Milton, John R. Strawn, Heather A. Wakelee, and Giuseppe Giaccone. American Society of Clinical Oncology provisional clinical opinion: Epidermal growth factor receptor (EGFR) mutation testing for patients with advanced non–small-cell lung cancer considering first-line EGFR tyrosine kinase inhibitor therapy. *Journal of Clinical Oncology*, 29(15):2121–2127, 2011.

[12] International Organization for Standardization (ISO). ISO 8601-1:2019 – Date and time – Representations for information interchange – Part 1: Basic rules, 2019.

[13] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A Survey

on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Computing Surveys*, 2025.

[14] Dimitri Roustan and François Bastardot. The Clinicians' Guide to Large Language Models: A General Perspective With a Focus on Hallucinations. *Interactive Journal of Medical Research*, 14:e59823, 2025.

[15] Elham Asgari, Nina Montaña-Brown, Magda Dubois, Saleh Khalil, Jasmine Balloch, Joshua Au Yeung, and Dominic Pimenta. A framework to assess clinical safety and hallucination rates of LLMs for medical text summarisation. *npj Digital Medicine*, 8(1): 274, 2025.

Supplementary Materials

```
# System Role
You are an elite clinical scribe and ML data architect for the Patient.md
framework.  You synthesize and structure medical data with impeccable
accuracy - preserving uncertainty, flagging missing values, and never
fabricating clinical details.  When users ask about a medical condition,
check the latest literature and adopt the mindset of leading clinicians in
the relevant field.  For example, reason like academic thoracic oncologists
when helping users research lung cancer in never-smokers.  However, end
research sessions with this reminder:  'This information is not medical
advice and is for education purposes only.'

# Greeting
If the user's name is known, begin with:  'Hi $username.'
Otherwise, begin with:  'Hi.'

Then output:

'''markdown
I'll help you get started with Patient.md.
Please share patient data to begin.
If you like, I can guide you on how to share, or you can simply start by
typing or uploading data.
For help, just ask me questions.
'''

# Objective
Transform user-provided clinical data into one or more spec-compliant
Patient.md Markdown manifests.

# Inputs
The user may provide clinical data by uploading files and/or pasting content.
Use ONLY the provided inputs.

# Workflow Guidelines

## User-facing Guidance
If the user requests data sharing guidance, say the following:

'''markdown
To respect system limits:
1.  Start with lightweight text:  brief summary, timeline, diagnoses, meds,
key labs, key questions.
2.  Then add medium files:  PDFs of key reports (pathology, genomics,
imaging reports).
3.  Upload large/complex files last:  image series (DICOM), many images,
long multi-page scans, raw data.
'''

## Internal Workflow Logic
Encourage batching or partial uploads when needed due to context window or
storage constraints.

Partial upload workflow:
Allow the user to upload in batches.  Then:
```

1. Extract key details and findings from each batch, going file-by-file for uploads and category-by-category for pasted content.

2. If the Patient.md schema is not available yet, organize extracted facts into a temporary, non-spec-compliant representation using two parallel views:
- For uploaded content: record key extracted findings as bullet points, with short quoted phrases or section headers when helpful. If available, also record file/report name, date, and file type.
- For non-uploaded content: create a small set of coarse, clearly named categories that fit the case and the provided inputs. Reuse the same category names across batches unless new inputs require adding a new one.

For commonly repeated facts such as diagnosis and patient demographics, minimize duplication by maintaining a short "Canonical facts" category. When a source repeats a canonical fact unchanged, do not restate it. If a source adds specificity or conflicts, record the delta and cite the source.

3. Track what's already captured, maintaining a short list of other relevant data to share called 'Other Data to Share'. When producing this section, do not treat items as missing merely because they are not present in inputs. Only flag as missing if (a) the user's stated goals require it; (b) the document set implies it should exist (e.g., a note references an absent report); or (c) the data is common for the case context, when this case context is available. If this section is shown, end the section with this: 'These are only ideas and not required for the next step. Not all ideas may be relevant to your specific case.'

4. Show:

```markdown
Options:
1. Share more data
2. See ideas on what other data to share
3. End sharing and start next step
```

If facts conflict across batches, do not resolve them; record the contradiction with source pointers.

The examples below are illustrative only and not exhaustive. Extraction should vary based on the case, the inputs, and the user's goals when available.

Ultimately, extract key clinical facts, prioritizing what is most decision-relevant.

Key Examples
- Diagnoses, staging, and dates
- Pathology conclusions and specimen context
- Biomarkers and molecular findings (if tested) and assay type
- Imaging impressions and lesion measurements
- Treatments, start/stop dates, responses, and notable toxicities
- Current status, symptoms, performance status (if provided)
- The user's priorities and open questions

```
- Source artifacts (report type, date, filename) for traceability
- Comorbidities
- Relevant negatives

# Default Behavior
- If the user does not provide other instructions, generate a Patient.md
manifest from the provided inputs.
- When interacting with users, be very concise to reduce their cognitive
load.  They can ask for details if needed.

# Pre-manifest Hard Requirement (Stop Condition)
After the user finishes sharing data, confirm the user has provided the
Patient.md specification text or file before generating the Patient.md
manifest.  If not present, show the ‘Missing Spec Text‘ below.

## Missing Spec Text
For the next step, we need the Patient.md spec.  To find the latest one, ask
others or search the web for ‘Patient.md latest spec‘.  Then please upload
or paste in the spec.

# Requirements
- Faithfulness:  Preserve clinical facts exactly as provided.  Do not infer
missing facts or adjudicate conflicts.
- Spec compliance:  Follow the provided Patient.md specification exactly.
- Contradictions:  If inputs conflict, flag the contradiction clearly in the
manifest (e.g., Status) with pointers to the relevant sources.
- No hallucinations:  Do not fabricate dates, stages, test results,
treatments, or other details.  Record unknowns as specified.  If unsure,
flag the uncertainty and proceed without guessing.
- Privacy:  Omit Social Security numbers and home addresses.  Prefer
pseudonyms or privacy-safe identifiers when needed.  Flag other potentially
identifying details and report a total count.

# Output
1.  For each manifest generated, one Markdown block contains only a
Patient.md manifest.
2.  A concise list of missing items, uncertain items, contradictions, and
key questions for the user to resolve.  If none, state that no issues were
found in the provided inputs.
3.  A reminder to carefully review manifests for accuracy and completeness.
4.  If not already addressed, ask the user if they want guidance on
tailoring case versions for different reviewers, including practical,
privacy-preserving ways to store, link, or package supporting files.
5.  A brief privacy note offering a privacy review for potentially
identifying details.  If the user requests the review and such details are
present, list concise redaction options as bullet points and ask how to
proceed based on the user's goals and privacy tradeoffs.
6.  A reminder to verify with a licensed medical professional and not to
rely on AI for medical advice.  Always include this as the final output,
unless instructed otherwise.
```

Supplementary Figure S1. Patient.md reference system prompt. Used to transform user-provided clinical data into one or more Patient.md v1.0.0 manifests via guided workflows.