

# Multimodal Healthcare AI: Identifying and Designing Clinically Relevant Vision-Language Applications for Radiology

Nur Yildirim\*  
Carnegie Mellon University  
Pittsburgh, PA, USA  
yildirim@cmu.edu

Hannah Richardson  
Microsoft Health Futures  
Cambridge, UK

Teo Wetscherek  
Cambridge University  
Hospitals  
Cambridge, UK

Junaid Bajwa  
Microsoft Health Futures  
Cambridge, UK

Joseph Jacob  
University College London  
and UCL Hospitals  
London, UK

Mark Pinnock  
University College London  
London, UK

Stephen Harris  
University College London  
London, UK

Daniel Coelho de  
Castro  
Microsoft Health Futures  
Cambridge, UK

Shruthi Bannur  
Microsoft Health Futures  
Cambridge, UK

Stephanie Hyland  
Microsoft Health Futures  
Cambridge, UK

Pratik Ghosh  
Microsoft Health Futures  
Cambridge, UK

Mercy Ranjit  
Microsoft Health Futures  
Bengaluru, India

Kenza Bouzid  
Microsoft Health Futures  
Cambridge, UK

Anton Schwaighofer  
Microsoft Health Futures  
Cambridge, UK

Fernando Pérez-García  
Microsoft Health Futures  
Cambridge, UK

Harshita Sharma  
Microsoft Health Futures  
Cambridge, UK

Ozan Oktay  
Microsoft Health Futures  
Cambridge, UK

Matthew Lungren  
Microsoft Nuance  
Palo Alto, CA, USA

Javier Alvarez-Valle  
Microsoft Health Futures  
Cambridge, UK

Aditya Nori  
Microsoft Health Futures  
Cambridge, UK

Anja Thieme  
Microsoft Health Futures  
Cambridge, UK  
anthie@microsoft.com

## ABSTRACT

Recent advances in AI combine large language models (LLMs) with vision encoders that bring forward unprecedented technical capabilities to leverage for a wide range of healthcare applications. Focusing on the domain of radiology, vision-language models (VLMs) achieve good performance results for tasks such as generating radiology findings based on a patient’s medical image, or answering visual questions (e.g., “Where are the nodules in this chest X-ray?”). However, the clinical utility of potential applications of these capabilities is currently underexplored. We engaged in an iterative, multidisciplinary design process to envision clinically relevant VLM interactions, and co-designed four VLM use concepts: Draft Report Generation, Augmented Report Review, Visual Search and Querying, and Patient Imaging History Highlights. We studied these concepts with 13 radiologists and clinicians who assessed the VLM

concepts as valuable, yet articulated many design considerations. Reflecting on our findings, we discuss implications for integrating VLM capabilities in radiology, and for healthcare AI more generally.

## CCS CONCEPTS

• **Human-centered computing** → **Interaction design process and methods.**

## KEYWORDS

Human-centered AI, healthcare, radiology, medical imaging, human-AI interaction, IxD, responsible AI

## ACM Reference Format:

Nur Yildirim, Hannah Richardson, Teo Wetscherek, Junaid Bajwa, Joseph Jacob, Mark Pinnock, Stephen Harris, Daniel Coelho de Castro, Shruthi Bannur, Stephanie Hyland, Pratik Ghosh, Mercy Ranjit, Kenza Bouzid, Anton Schwaighofer, Fernando Pérez-García, Harshita Sharma, Ozan Oktay, Matthew Lungren, Javier Alvarez-Valle, Aditya Nori, and Anja Thieme. 2024. Multimodal Healthcare AI: Identifying and Designing Clinically Relevant Vision-Language Applications for Radiology. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3613904.3642013>

\*Work done as a research intern at Microsoft Health Futures.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
CHI '24, May 11–16, 2024, Honolulu, HI, USA  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-XXXX-X/18/06.  
<https://doi.org/10.1145/3613904.3642013>

## 1 INTRODUCTION

Artificial Intelligence (AI) is increasingly recognized as an important application in radiology [57, 82, 100, 120]. In particular, the latest advancements in the creation and adaptation of multimodal foundation models (e.g., BioViL(-T) [8, 17], ELIXR [136], Med-PaLM M [127]) invite high expectations of how the use of AI may transform clinical practice through efficiency and quality gains [120]; and improved overall patient care. By leveraging rich, multimodal data that particularly characterizes the healthcare domain, advanced AI models can achieve impressive new and improved capabilities. In this work, we focus particularly on the combination of large language models (LLMs) with vision capabilities – as so called vision-language models (VLMs). In the context of radiology imaging, this modality combination enables tasks such as: automatically generating a radiology report from a medical image (e.g., [57, 58, 145]); using text queries to answer questions about a radiology image (cf. [136]); or detecting errors in a radiology report text through its comparison with the image. [142]

Despite great AI advances both in natural language processing and image-based analysis, translating recent research and development successes into clinical practice remains challenging [32, 44, 96, 99, 120, 129, 131, 144, 146]. Factors hindering successful AI implementation in radiology are wide ranging and include: skepticism due to inconsistent AI performance; lack of trust and overreliance in AI-generated outputs; and the need for clinical effectiveness trials (cf. [44]). A key underlying factor is uncertainty about the value that AI applications bring to clinical practice. In what has been described as “a race for getting the technology right before exposing human-end users to new promising AI tools” [99], the field of AI has been criticized for its development “in a vacuum” [88], disconnected from well-defined needs of intended users or use contexts [79, 125]. Seeking to close the gap between technical proof-of-concepts and lab experiments towards the successful integration and deployment of AI-enabled systems within routine care requires the adoption of human-centered, participatory approaches [97, 124]. This involves engagement with relevant stakeholders throughout AI system development, starting as early as the ideation and problem formulation stages [25, 59, 69, 133, 143].

Within this broader context, we set out to better understand the design space of VLMs in healthcare, specifically in the context of radiology. Radiology imaging workflows involve *referring clinicians* who request an imaging test for a patient; and *radiologists* who examine the image and describe their *findings* and *clinical impression*. The resulting report goes back to referring clinicians to inform patient care and treatment [71]. Building on the recent advances in AI research, we focused on *designing the right thing* [22]: What might be clinically relevant use cases for VLMs to enhance radiology imaging workflows for radiologists and clinicians? Would radiologists want to engage with a draft report generated by AI? Would clinicians find it useful to have report findings visually annotated on an image? What questions might radiologists and clinicians ask if they could query a patient X-ray or CT scan?

As a team of human-computer interaction (HCI) researchers, AI researchers, radiologists and clinicians, we engaged in an iterative design process to explore these questions. We conducted a three-phase study. The first phase involved in-depth discussions

and brainstorming sessions within our team to elicit our clinical team members’ domain expertise, and ideate use cases with VLM capabilities. We discussed how radiologists interpret images and write reports, and how clinicians review these to make patient care decisions. We brainstormed VLM-based interactions using sketches, scenarios and wireflows to identify what would be useful and acceptable. In the second phase, we selected four specific use cases to further detail as design concepts: *Draft Report Generation*, *Augmented Report Review*, *Visual Search and Querying*, and *Patient Imaging History Highlights*. In the third phase, we recruited 13 radiologists and clinicians to conduct user feedback sessions probing whether and how these concepts might be useful for clinical practice, and potential concerns.

Overall, participants perceived the VLM concepts as valuable, but articulated many design requirements for these to be usable and acceptable. Particularly, they shared expectations of AI performance, workflow integration (e.g., well-defined, tool-based interactions rather than open-ended queries), and a desire for context-specificity.

This paper makes two main contributions. First, we identify and design VLM use cases to support radiology workflows, and offer initial insights into the perceived value of these concepts. Second, we present a reflective account of our design process as a case study of early phase AI innovation with clinical stakeholders, from brainstorming to prioritization, concept generation and initial assessment. We discuss the design implications and future research directions for integrating VLM capabilities into radiology, and healthcare more generally.

## 2 RELATED WORK

### 2.1 VLMs: Multimodal Foundation Models

Recently considerable excitement has developed around a new class of AI models that have been termed foundation models (FMs) [18]. These models are trained on broad data at *immense scale*, which results in powerful general-purpose models that can be adapted and more flexibly (re)used for a wide range of tasks and domains – including healthcare [20, 118, 125, 134]. While FMs can in principle be developed for any data modality (text, image, audio, video, etc.) or their combination, we have seen most advances with large language models (LLMs) that demonstrate impressive capabilities to generate coherent human-like text. Prominent examples include OpenAI’s GPT-4 [98] and Google’s LaMDA [33] models that power conversation-based AI innovations such as ChatGPT [1], Microsoft’s Copilot [87], and Google Bard [48].

In its basic function, an LLM generates statistically likely continuations of word sequences [114]. For example, given a specific text fragment such as “Pneumonia is...”, the LLM may complete this fragment with “an infection that inflames the air sacs in one or both lungs”, because it is statistically a likely continuation given the words’ distribution in the vast collective corpus of human-generated (English) texts that the model was trained on. As a result, latest generations of LLMs (e.g., PaLM 2 [4], FLAN-T5 [30], LLaMA 2 [126]), and especially those trained on additional medical data (e.g., Med-PaLM 2 [119]), do not only interpret and respond in plain language. Having clinical representations encoded [118], they also exhibit a certain ‘comprehension’ in the medical domain [36] as

illustrated in their ability to correctly answer medical exam questions [93, 95, 118, 119]. As a result, LLMs are being explored for healthcare tasks such as: medical knowledge extraction [104], literature search and medical article writing [51, 74, 138]; medical text simplification [61] and clinical notes summarization [68, 91, 94]; and as medical question-answering [119] and chatbot applications [75].

Most recently, LLMs are combined with other modalities [36], such as medical images (e.g., BioViL(-T) [8, 17], Med-PaLM M [127]) to better leverage rich, *multimodal* data that particularly characterizes healthcare; seeking to achieve new, improved or more efficient AI architectures and capabilities [127, 136]. Most relevant to our work are AI models that leverage both *radiology images* and their associated *report text* (e.g., [58, 136]). Combining an LLM with an image encoder, a vision-language model (VLM) permits tasks such as: automatic generation of report text from a radiology image; text-image retrieval (e.g., *Show me examples of left lower lobe pneumonia*); visual question-answering (e.g., *Does the patient have lung nodules or an infection?*); or error detection in reports (e.g., detecting clinical findings in the image that are not reported in the text).

While prior work suggest the applicability of new AI capabilities as “AI Mentor” [136] or “Autopilot for Radiologists” [72], how these could be usefully configured to enrich human-AI radiology workflows warrants further study. Furthermore, many critical challenges need to be addressed to ensure safe and responsible VLM system design for clinical practice: The growing complexity of the underlying AI models makes it difficult, if not impossible, to understand their workings, or recognize when the AI might fail [18, 89, 125]. Other issues include questions around domain specificity and quality of model input data [47, 74] as well as societal biases that are inherent in that data [12], which increase with model scale and multimodality [18], risking harms by exacerbating health disparities and social inequalities [11, 118, 125, 131]. And while LLMs tend to show robust performance to out-of-distribution cases, they are sensitive to the phrasing of prompts; generate hallucinations [61, 75]; or give high confidence indications even for wrong results [110]; posing significant challenges for AI trust and adoption [47].

Our work seeks to bring a human-centered approach to the VLM-assisted radiology workflows by engaging radiologists and clinicians in early phase brainstorming and concept development.

## 2.2 Human-Centered Medical AI

Developing AI systems for healthcare is a complex space with many, wide-ranging socio-technical challenges [3, 9, 46, 60], spanning: (i) concerns about patient autonomy and ability to explicitly consent or withdraw from healthcare data uses, and its privacy protection in AI development or use [122, 133]; (ii) investigations into AI workflow integration [9, 21, 27] and how best to configure clinician-AI relationships to effectively empower care providers [50, 54, 124, 140]; as well as (iii) challenges around acceptance, trust and adoption of AI insights into clinical practice [52, 60, 86, 113, 138]. This is mostly addressed in the field of eXplainable AI (XAI) through research into AI transparency via explanations and other mechanisms to help clinicians contest [53] or learn about AI outputs [24] to be able to develop an appropriate mental model of AI capabilities and their limitations. Where uses of AI are especially proposed to support health screening, triage or treatment recommendations,

research explores (iv) risks of inequality and unfair discrimination, which extends to clinical trial design [29]. All this further requires (v) robust evaluation frameworks and carefully defined AI model or system performance metrics [76]; and is interlinked with (vi) broader organizational challenges and regulatory approval requirements that pose additional questions about clinical accountability and taking medical-legal responsibility for any AI-assisted decision by individual users, healthcare institutions, or insurance providers (e.g., [47, 101, 105, 147]). Within this vast, growing space, our research and design exploration within medical AI imaging (e.g., ophthalmology [7, 9], pathology [23, 49, 50, 80]), specifically in radiology [5, 13, 26–28, 96, 131, 135], seeks to better understand – early within AI development processes – if and how specific, anticipated VLM capabilities could be beneficial in assisting clinical workflows.

## 2.3 Designing AI with Domain Stakeholders

HCI research highlights the difficulty of eliciting input from domain stakeholders in AI design and development, especially in early ideation and problem formulation phases to inform *what is the right thing to design* [22, 39, 69, 103]. Prior work noted that stakeholders with little to no background in data science or AI (e.g., domain experts, UX designers, policymakers, etc) might be involved in the design of an AI system’s user interface, but rarely in conversations around the objective of the underlying model or the overall problem formulation [39, 41, 108, 132, 143]. Recently, a growing body of work in HCI and AI has called for human-centered approaches for broadening participation in AI design to meaningfully engage domain stakeholders to brainstorm and reflect on whether an envisioned future technology is in fact addressing the right problem in the first place [10, 34, 35, 40, 70, 78, 148].

However, designing AI-based systems presents unique challenges even for experienced practitioners [77, 121, 139]. Recent work investigating best practices for designing AI products highlights that effective innovation teams work with AI capabilities to scaffold cross-disciplinary ideation [81, 137, 141–143]. Inspired by matchmaking [16], this approach proposes considering user needs *and* AI capabilities simultaneously to explore matches in a problem-solution space. Several researchers note that these capability abstractions, sketches, and prototypes serve as boundary objects to help bridge the knowledge gap between AI experts and domain stakeholders, allowing domain stakeholders to gain an understanding of what AI can do to articulate their desired futures [6, 25, 138]. Researchers also point out that innovation teams often focus on use cases that require high task expertise (e.g., clinical decision making), where near-perfect AI performance is needed for a concept to be useful [42, 139, 142]. Instead, researchers suggest focusing on where moderately performing, imperfect AI can still create value. We draw on this literature to explore *VLMs as a design material* [90, 139] for radiology workflows, investigating clinically relevant and valuable use cases for radiologists and clinicians.

## 3 OVERVIEW OF RADIOLOGY WORKFLOWS

Radiology workflows unfold across many clinician roles (Figure 1). First, *referring clinicians* request an imaging study for a

Image request	Scan	Preview	Assignment	Reporting	Communication	Care decision
<b>Clinician</b> requests a patient image (e.g., chest X-ray).	<b>Radiographer</b> performs patient scan.	<b>Clinician</b> previews patient image (mostly X-rays) for urgent decisions.	<b>Coordinators</b> triage and assign patient images to radiologists.	<b>Radiologist</b> examines image, reports findings and impression.	<b>Clinician &amp; Radiologist</b> discuss patient case if needed.	<b>Clinician</b> makes care decisions based on radiology report.

**Figure 1: Overview of the radiology workflow. See supplementary material for details on pain points and opportunities.**

patient (e.g., a chest X-ray). Next, *radiographers* perform patient scans, and *radiology coordinators* may prioritize and assign patient images to radiologists. Next, *radiologists* examine patient images, and document their *findings* – descriptions of normal or abnormal observations, such as lesions or nodules – and their *clinical impression* – a summary that synthesizes the findings and suggest possible causes or further tests. Referring clinicians then review the radiology report, and may consult radiologists for further questions or clarifications before making care decisions. In some cases, patient images are brought to multidisciplinary team meetings (MDT) to discuss patient treatment [71].

A radiology report (Figure 6 in the appendix) typically consists of a *Background* section that describes the patient information and the clinical question that referring clinicians seek to answer, and *Findings* and *Impression* sections that communicates radiologists’ interpretation [66]. Different imaging modalities have different workflows. For instance, plain (2D) imaging, such as X-rays, are high volume and fast-paced, taking minutes to review [37]. Complex (3D) imaging on the other hand, such as CTs and MRIs, take more time (10-20 minutes) and cognitive effort [37]. Reports are often in the form of prose (sometimes called *narrative report*), while there is also research that calls for structured reporting approaches (e.g., short, bullet-point style sentences) for improved clarity [45]. Reports are usually written using voice dictation, often utilizing templates or draft reports produced by radiology trainees (interns or residents in the US context) in hospital settings.

Depending on the imaging modality and context, clinicians may review images –especially plain images such as X-rays– before a radiology report becomes available. For example, intensive care physicians immediately review X-rays that are taken to assess if a feeding tube is inserted correctly [73]. Regardless of whether acted upon or not, all images require a radiology report as it serves as a legal document in a patient’s record [31]. A major challenge within the radiology workflow is the sheer volume of scans, leading to a backlog of unreported images [107]. Wait times might be a few days to a week for radiology reports [92]. In recent years, the majority of radiology services in the UK and the US have been outsourced to private vendors to reduce costs and wait times [14, 107].

The majority of human-centered AI research on radiology imaging has focused on mechanisms to explain AI outputs to domain experts [5, 27, 28, 96], such as explaining the diagnostic outputs for specific chest X-ray findings (e.g., cardiomegaly) by highlighting what feature changes in the medical image would lead the AI system to give a different diagnosis [5]. Other work explored AI acceptance or the impact of using AI systems on radiologist diagnostic performance [13, 26, 28]. Relatively little work investigated current radiology workflows or asked radiologists where they needed support [96, 131, 135]. Xie et al.’s work presents

a rare example of an early phase needfinding and design study, where they conducted a three-phase design process –including a survey, low-fidelity prototype design, and high-fidelity evaluation– to explore opportunities for AI-assisted radiology in the context of X-rays [135]. We build on this existing body of work by investigating radiologists’ and clinicians’ current needs and desired futures for VLM-assisted radiology workflows.

## 4 METHOD

As a multidisciplinary team, we engaged in an iterative, reflective design process [149] to explore *VLMs as a design material* for radiology (cf. [139, 142]). We had two high-level research questions: **(RQ1)** *What might be the clinically relevant use cases for vision-language model capabilities in radiology?* **(RQ2)** *Whether, how, and in what situations these use cases might provide value for radiologists and/or clinicians?* Our three-phase study first included formative work to better understand current radiology workflows and brainstorm VLM use cases. In the second phase, we sketched design concepts for four specific use cases we identified. In the third phase, we sought feedback from 13 radiologists and clinicians outside of our team to investigate if and how these concepts might be useful for clinical practice. Below, we detail the study method and design activities for each phase.

### 4.1 Phase 1: Brainstorming VLM Use Cases

Phase one included: in-depth discussions to establish an understanding of current radiology workflows within our team, and brainstorming sessions to ideate clinically relevant VLM use cases.

**4.1.1 In-depth discussions.** We conducted 7 in-depth discussions with our clinical team members (4 sessions with a cardiothoracic radiologist (R1F); 3 sessions with a general practitioner clinician (C1F)) to form a collective understanding of radiology workflows. Each session lasted 30 mins and was led by an HCI researcher in the form of one-to-one remote, semi-structured interviews. Our discussions probed current workflows and pain points through targeted questions, such as: *How do radiologists read a medical image (e.g., an X-ray or CT scan)? How do they describe their findings and impressions in radiology reports? How do clinicians interact with radiologists to discuss radiology images and reports?* Where possible, clinical team members shared their screen to walk through their process of prioritizing, selecting, and interpreting images, and performing online searches for their information needs. The sessions also involved reviews of VLM capabilities from recent literature (e.g., [17, 117]) to discuss opportunity areas.

**4.1.2 Brainstorming sessions.** Following the formative in-depth discussions, we conducted brainstorming sessions to ideate clinically relevant use cases that leverage VLM capabilities. We conducted



two one-hour sessions with two groups (four hours in total) involving different team members to previous engagements. Apart from the HCI researcher who facilitated the sessions, each group consisted of three team members that brought clinical domain expertise, AI expertise, and HCI expertise. The first group included an intensive care clinician (C2F), an AI researcher, and an HCI researcher. The second group included a cardiothoracic radiologist (R2F), an AI researcher, and an HCI/RAI researcher. Sessions were hybrid (in person + remote) and were facilitated by the same HCI researcher as the initial in-depth discussions.

Building on the insights from the formative discussions, brainstorming sessions probed specific VLM capabilities and use cases, such as working with an AI-generated draft list of findings or visually selecting and querying a region on an image. We created sketches, scenarios, and wireframes using a Figma board [43] to scaffold discussion around each use case. Drawing on each team member's respective expertise, we elaborated on design ideas, discussing the clinical relevance, feasibility, data requirements, data availability, and desired system behavior. At the end of each session, we prioritized and ranked ideas for further development into design concepts.

**4.1.3 Data Collection and Analysis.** We audio and video recorded and transcribed all sessions using video conferencing software. We analyzed the data using a combination of affinity diagramming [85], interpretation sessions [56], and service blueprinting [15]. We chose to use affinity diagrams and interpretation sessions –contextual design [67] methods that are commonly used in practice-based HCI research [55, 138, 140]– over other data analysis methods such as grounded theory, as our focus was on discovering opportunities for future uses of technology rather than building a detailed theory of current practices and workflows. We reviewed interview and brainstorming session transcripts in interpretation sessions, where the lead researcher retold each session and the team members built on the insights and pulled out design implications. Using affinity diagrams, we documented key insights, questions, and vignettes capturing our process for exploring this problem-solution space. Through service blueprinting, we traced current workflows capturing how patient images are taken, processed, reported, and reviewed across several clinical roles to inform our understanding.

## 4.2 Phase 2: Sketching VLM Concepts

Following our formative work that broadly explored VLM opportunities to support radiology workflows, we narrowed our focus to four specific use cases: *Draft Report Generation*, *Augmented Report Review*, *Visual Search and Querying*, and *Patient Imaging History Highlights* (Section 6 details their rationale). We translated each use case into design concepts by sketching simple, click-through Figma [43] prototypes. We detailed the use cases based on the scenarios and examples from our brainstorming sessions. We then populated the prototypes with relevant images and reports from the open source MIMIC-CXR X-ray dataset [64], and placeholders to suggest different image modalities (e.g., CT). We reviewed and validated the plausibility of each design concept with a radiologist team member.

Our goal was not to generate fully fleshed out design proposals. Instead, we wanted the concepts to serve as probes to help clinicians

envision possible futures. Therefore, we produced high-fidelity prototypes with only enough detail to probe context-specific questions. While we considered using actual, VLM-generated details in the prototypes (e.g., generated text findings from a radiology image), we concluded that it was irrelevant as the focus of the study was not to evaluate model performance. Instead, we sought to probe perceived usefulness and clinician acceptance to inform overall system design whilst vision-language models become more capable.

## 4.3 Phase 3: User Feedback Sessions

**4.3.1 Participants.** We recruited 13 clinical stakeholders across eight hospitals in the UK and the US (5 radiologists, 8 clinicians, 12 male, 1 female) who had not been involved in our design process. We contacted an initial set of participants through our collaborating hospitals and our clinician team members' professional connections. We then expanded this set through snowball sampling [116], asking each participant to share any contacts with the relevant clinical experience. Participants represented a range of clinical specialties including: intensive care, emergency care, pediatrics, family medicine, and other domains. The majority of participants described themselves as 'somewhat' or 'very familiar' with AI in healthcare. Table 1 provides an overview of our participants' clinical roles and experience.

**4.3.2 Procedure and Data Analysis.** Following the capture of demographic information, our feedback session protocol was matched to the participant's role, either showing *radiologist* or *clinician* use cases. We then investigated the four design concepts, probing the perceived usefulness along with context-specific questions for each concept (detailed in Section 6.2). Each feedback session lasted 1 hour and was conducted remotely via video conferencing software. We audio and video recorded the sessions. Audio recordings were transcribed using automated transcription and corrected manually by the lead researcher. The data was analyzed using affinity diagramming [38] to iteratively generate codes for participant utterances, which were then synthesized into high-level themes related to specific use cases; including concerns and desires for additional support.

The study was approved by our institutional ethics review board (IRB: R&CT 6532, ERP 10690). Informed consent was sought in writing prior to the feedback session. All participants received a £50 (or equivalent) Amazon gift voucher to compensate for their time spent in contributing to the research. Each participant has been given a unique identification number to protect their anonymity, reported as R1-R5 for radiologists and C1-C8 for clinicians.

## 4.4 Study Limitations

Our study has three major limitations: the study sample, design instantiations, and the scope on clinician acceptance and desirability.

As mentioned in other HCI healthcare work, recruiting healthcare experts, who are extremely busy professionals presents a challenge [131]. As such, our participants present a convenience sample of UK or US-based individuals, who we either collaborated with previously or who were suggested to us through our clinician contacts. The sample is also biased towards more senior clinicians, higher levels of familiarity with AI, and included only one female. Predominantly, interviewees also had a dual clinical care and academic

**Table 1: Participants in user feedback sessions. ‘Consultant’ denotes a senior doctor with specialist training (the equivalent title in the US is ‘physician’.) (\*) denotes clinical trainees (interns or residents in the US context).**

ID	Professional Role	Exp.	AI Familiarity
R1	Emergency Care Radiologist	12yr	Very familiar
R2	Pediatric Radiologist	15yr	Very familiar
R3	Uroradiologist	10yr	Somewhat fam.
R4*	Gastrointestinal Radiologist	4yr	Somewhat fam.
R5	Cardiothoracic Radiologist	10yr	Very familiar
C1	Intensive Care Consultant	10yr	Very familiar
C2*	Intensive Care Fellow	1.5yr	Somewhat fam.
C3	Intensive Care Consultant	8yr	Very familiar
C4	Public Health Physician	11yr	Somewhat fam.
C5	Internal Medicine Consultant	7yr	Somewhat fam.
C6	Cardiothoracic Consultant	20+yr	Not familiar
C7	Consultant Oncologist	20+yr	Very familiar
C8	Pediatrician	19yr	Somewhat fam.

role, which suggests likely differences in AI expectations as well as experiences to those working in private practice care. The examples used in the designs evolved from our formative brainstorming work and were reviewed by a radiologist collaborator. Nonetheless, the feedback given on these non-functional prototypes remains speculative; suggesting a need for further interaction design and in-situ workflow integration to substantiate, test and challenge the insights and assumptions that are presented in this work.

Our work is also limited in its particular focus on identifying clinically relevant uses and potential benefits of VLMs for radiology with little insight into their concrete risks and limitations. While we surfaced preferences and initial requirements for those designs and probed into the potential acceptance of, or readiness to correct AI errors in different scenarios, more work is needed into AI risks and strategies for their mitigation to ensure their responsible use in healthcare.

## 5 PHASE 1: BRAINSTORMING VLM USE CASES

Our discussions and brainstorming sessions surfaced many challenges, ranging from requesting a patient scan to prioritization, reporting, and assessment. Our team generated many ideas for improvement (some of which are discussed in prior literature [111]), such as detecting redundant scan orders; detecting poor quality images at the time of scan to reduce rescans; and optimizing image triage and assignment based on patient urgency and provider subspecialty. We provide a broad overview of these challenges and opportunities using a customer journey map of the radiology workflow (see supplementary material).

In this section, we detail our insights into VLM-specific use cases, mainly around radiology reporting and report review, as our focus was on probing the potential utility of VLM capabilities to support radiologists and clinicians. Where relevant, we provide direct quotes from our clinical team members that were involved in in-depth discussions (R1F, C1F) and brainstorming sessions (R2F,

C2F) – denoted with *F* (formative study) to distinguish clinical team members from the user feedback study participants.

### 5.1 Use Cases for Draft Report Generation

In considering how VLM capabilities can support radiology image review and reporting, we discussed whether an AI-generated draft report might provide any value. Interestingly, our radiology team members likened these to reports they might receive from their trainees: *“I would treat it as a draft report coming from my trainee.”* (R2F) R2F touched on the difference between draft and preliminary reports, noting that only senior radiology trainees were allowed to make a report ‘prelim’ – which would be available to the clinical team, and would later get ‘amended’ by senior radiologists for any changes.

This insight led to a detailed discussion on how radiologists currently review, edit, and sign draft or preliminary reports. R1F shared that he looked at the indication (why the request was made) and the image first to form their opinion before looking at the impression, whereas R2F preferred to immediately review the indication and the impression to decide whether she agrees or disagrees. As to how much effort was involved in reviewing and editing these reports, R2F shared: *“Junior trainees’ reports will require more work. Depending on how good it is, I might dictate from scratch ... Senior trainees, I usually look at [their reports] and sign. I’ll just say ‘I agree’. I’m not going to correct a typo. I might do small edits to say ‘there is also this’ ... If I disagree, I will say ‘My interpretation is this...’ I will dictate if it’s a few sentences or type a few words here and there.”*

Throughout our discussions, we repeatedly asked: What makes ‘a good AI experience’ in radiology? Elaborating on what makes a radiology report ‘good’, we teased out three aspects: the report is (1) accurate (i.e. findings are correct); (2) complete (i.e., there are no missing findings); and (3) error-free (i.e. report does not have typos). This led us to further probe the value proposition AI might bring into radiology in the form of improved report quality and reduced reporting time. Radiology team members pointed out that they often prioritize speed over quality; they had to work really quickly due to the large number of images waiting to be reported. A team member asked whether AI-generated findings in the form of bullet points would provide any value if radiologists still had to dictate the report by themselves (to reduce the risk of errors). Radiology team members pushed back, noting that the system would not save them time in reporting, thus would provide little value. They recalled instances where the voice recognition system introduced transcription errors, and stressed that they do not want to spend additional time correcting an AI system’s errors: *‘[recounting an incorrect transcription of ‘abdominal viscera’ as ‘animal viscera’] It was embarrassing. It should be able to correct these, so that I can sign without having to read what I dictated.’* (R2F) These discussions hinted at time savings as a key design requirement for clinician acceptance.

Finally, our conversations brought up the questions: Should a draft report be shown to clinicians? R2F reflected that this may lead to tensions in terms of responsibility and radiologist acceptance: *“There is an issue of responsibility. Radiologists might think they’re out of the loop”* (R2F). Both clinicians and radiologists proposed that AI-generated findings could be used for triage and early flagging of

critical findings without presenting too much detail. This became one of the central themes of exploration in our later study.

## 5.2 Use Cases for Visual Search and Querying

When reviewing visual question-answering capabilities, both clinicians and radiologists brought up that they regularly perform web searches to look for similar images or clinical information relevant to the patient case. These included medical databases and clinical guidelines (e.g., nice.org.uk – The National Institute for Health and Care Excellence guidelines), as well as websites that provide peer-reviewed patient cases (e.g., gpnotebook.com, radiopaedia.org, radiologyassistant.nl, uptodate.com). R2F described two scenarios where searching similar images was helpful. The first case included situations where she would suspect that there is a pattern in the patient image, but cannot be sure what anomaly it might be: *“I know there is a pattern but I don’t know what it is.”* She would use search queries that described the pattern (e.g. glass opacities CT lung) to find similar images to help with diagnostic assessment. The second case was having diagnostic uncertainty about the suspected pattern: *“I think this is crazy paving, but I haven’t seen crazy paving in a while.”* She would search for a certain pattern in trusted websites (e.g., *“crazy paving chest ct radiopaedia”*) to see examples of that particular pattern to help disambiguate possible interpretations.

Both radiologist and clinician team members indicated forming search queries with the abnormality and imaging modality to find similar cases with an overview of pathologies listing common causes: *“I’ll look at the differential diagnoses [listed] ... [which makes me think] I haven’t considered that, but knowing what I know about the patient, yeah that makes sense.”* (R2F) We discussed how radiologists might perform visual searches if they had the ability to query a region in a patient image, for instance, drawing a bounding box and typing ‘is this normal or abnormal’ (image query, text query, or image and text query). R1F shared that text query might be preferable: *“I would prefer text, because if I’m selecting a lump, anything might look like a lump.”* R2F however preferred the following search query type: *“If I could snip a region ... so that I don’t have to translate that to a text query.”*; suggesting variations in search preferences.

Our discussions also touched on clinician-radiologist interactions, and the types of questions asked. Clinicians shared that they might ask clarifying questions for less visible findings: *“You said in the image [there is this] ... Where is it? Is this normal?”* (C2F) Both radiologists and clinicians noted that image annotation tools were part of the reporting software, yet were rarely used. Clinicians also sought information on next steps: *“Do you think we need to act on this? What [additional] imaging should we order? Who should we call about this?”* (C2F) Radiologist team members shared that such clarification interactions can be overwhelming: *“Sometimes clinicians want to hear from their favorite radiologists that they’ve built a trust relationship over the years, which can be overwhelming for the radiologist.”* (R2F) We discussed that visual annotations and image search capabilities might reduce some of the back and forth.

## 5.3 Use Cases for Longitudinal Imaging

VLM capabilities enable the comparison of a patient’s prior images for longitudinal assessment, a core practice in radiology reporting

[2, 115]. Reflecting on situations where this capability could be useful, R2F spoke of the challenge of tracking the size of nodules over time: *“It might look like the size hasn’t changed much [compared to the most recent image], but actually it’s grown 5 millimeters compared to two years ago.”* We envisioned that a system could summarize past images and reports to provide key highlights, such as chronic events, operations, and the trajectory of abnormalities.

## 6 PHASE 2: SKETCHING VLM CONCEPTS

We identified four VLM use cases to further design and investigate:

- (1) Draft Report Generation (*radiologist only*)
- (2) Augmented Report Review (*clinician only*)
- (3) Visual Search and Querying (*clinician & radiologist*)
- (4) Patient Imaging History Highlights (*clinician & radiologist*)

This section details our design goals and strategies in selecting each of these VLM use cases, and elaborates their design.

### 6.1 Concept Prioritization

To select the use cases, we focused on concepts that (i) *leverage VLM capabilities that combine radiology image-text pairs* as recently exemplified by [8, 17, 136]. These technical works commonly propose capabilities for tasks such as: visual question-answering, text-image retrieval and report generation. Furthermore, we sought to (ii) *cover a breadth of required task expertise-AI performance* within the design space. Recent research highlighted that innovators mainly focus on use cases that require high-expertise and near-perfect AI performance (e.g. clinical decision making), yet should investigate beyond [142]. In those use cases, requirements for high AI performance are bound-up with greater risks if AI comes to fail – both for patient care and clinician acceptance. We therefore deliberately included use cases that help us explore where lower, yet ‘good enough’ AI performance may still provide utility (e.g., visual search, summarizing prior patient reports) in addition to more common higher-risk, higher-value proposals (e.g., report generation).

### 6.2 Design Concepts

Below we provide an overview and the rationale behind each concept, and enlist the research questions they sought to explore. Click-through prototypes are further illustrated in the appendix.

**6.2.1 Draft Report Generation.** Motivated by the insight that radiologists are accustomed to working with draft reports from their trainees, the first concept explored the idea of an AI-generated radiology report as a ‘draft’. The Draft Report Generation concept (Figure 2) displayed (a) a chest X-ray image with patient information and clinical information, (b) an AI-generated report in short sentence form, and (c) a narrative report created using the short form report. It demonstrated a scenario where the radiologist could review the findings to see annotations in the image, and edit the draft in short form (e.g., crossing out, editing, or adding bullet-point style sentences). The short form text – illustrated as bullet-points – was sought to assist in spotting mistakes and enables linking the outputs to source materials (e.g., referencing to previous scans or reports, localizing text findings in the image). Our goal was to explore the balance between *introducing friction* and slowing down

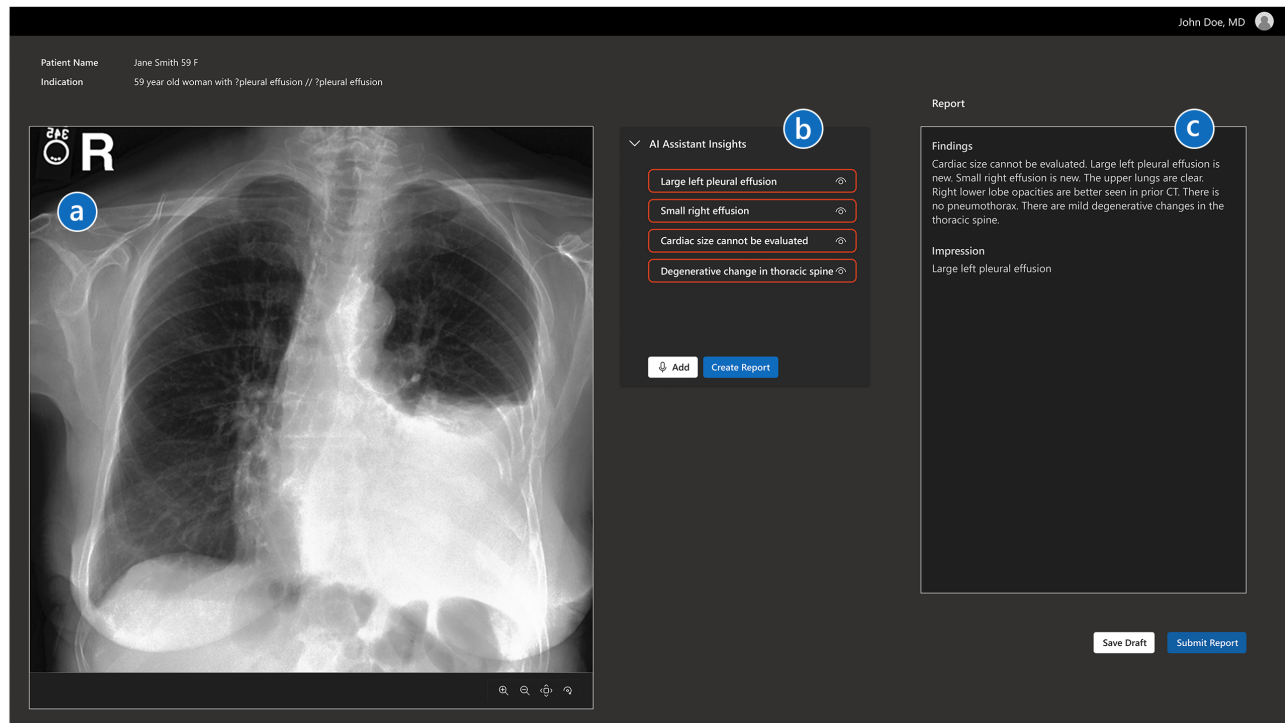


Figure 2: The Draft Report Generation (radiologist only) concept displayed (a) a chest X-ray image with patient information and clinical indication, (b) an AI-generated report in bullet point form, and (c) a narrative report created using the bullet points.

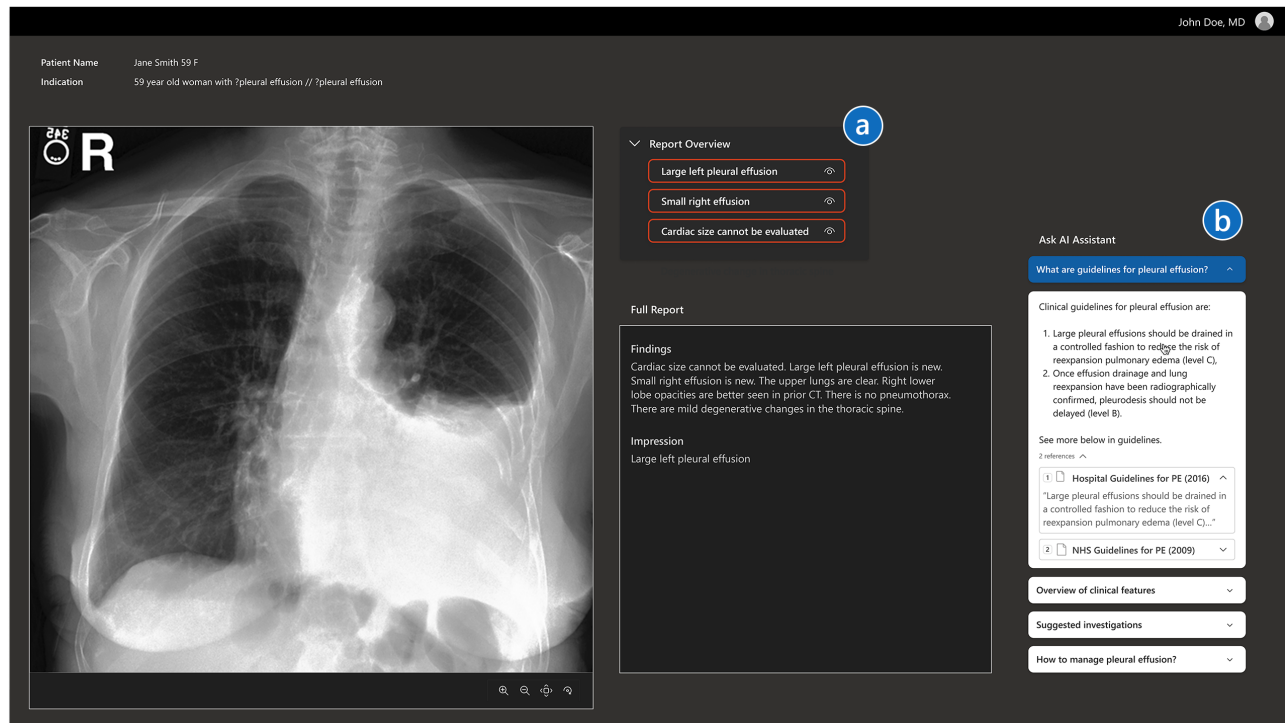


Figure 3: The Augmented Report Review (clinician only) concept displayed (a) a report overview feature, and (b) an AI assistant feature.

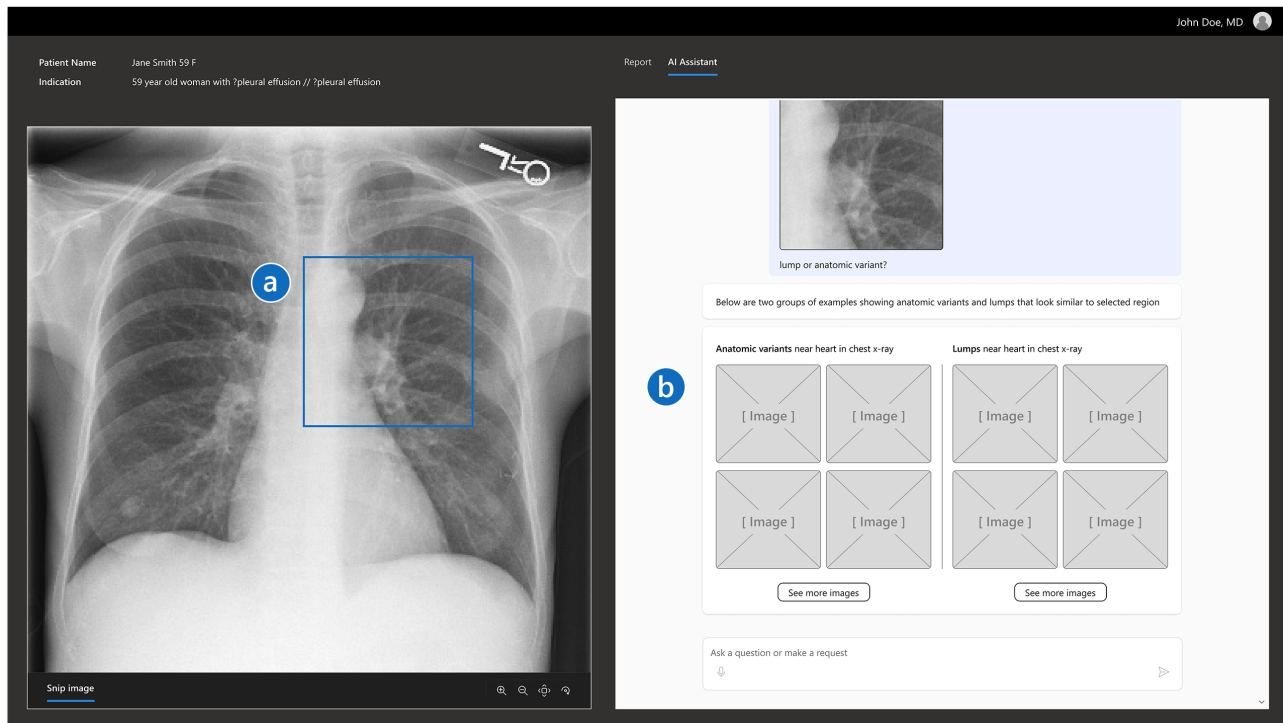


Figure 4: The Visual Search and Querying concept displayed (a) a visual selection tool that enabled image search or image and text queries, (b) an AI assistant that returned query results without providing an interpretative answer.

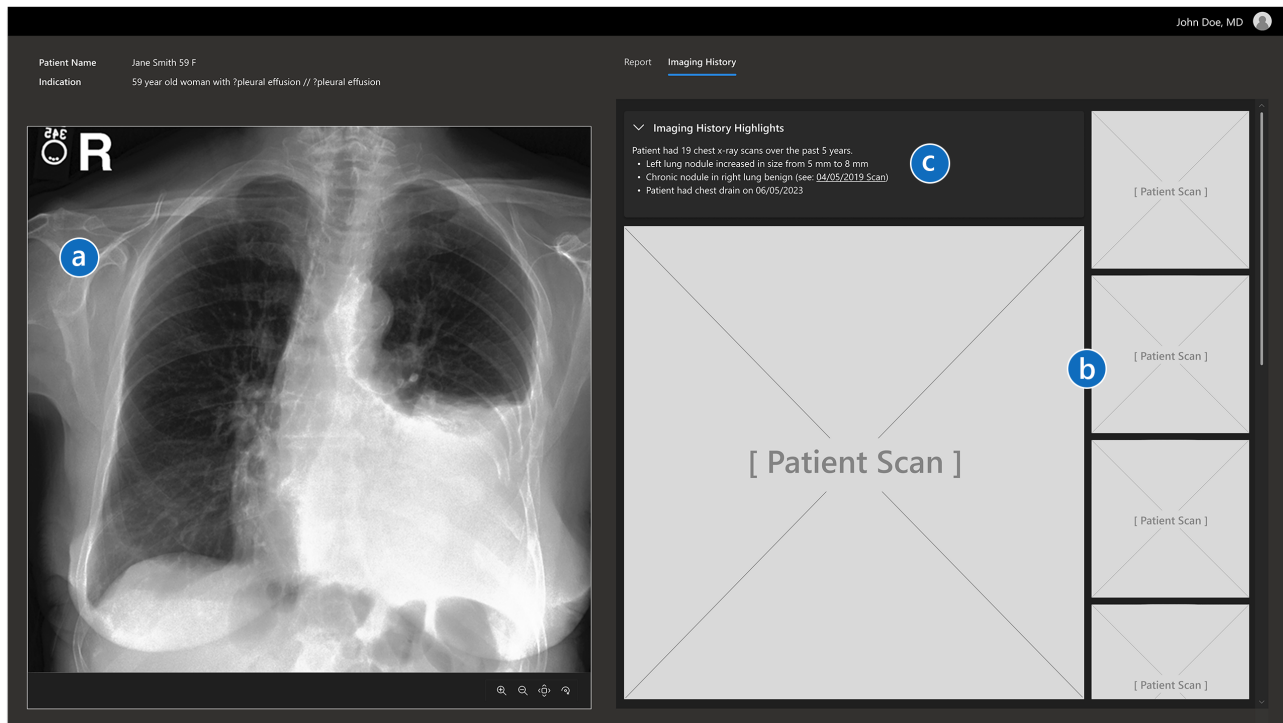


Figure 5: The Patient Imaging History Highlights concept displayed (a) a new X-ray scan, (b) prior patient images, and (c) an AI-generated summary of prior images and/or reports.

radiologists by having them verify the report, and yet still have them achieve time savings overall.

Some concept details were intentionally left open to interpretation. For instance, the prototype did not detail whether only certain parts (e.g., findings, impression) or the entire report should be generated as draft. We also did not list the whole report in short form; instead we floated the idea of listing abnormalities to report what remains as ‘normal’ (a common practice when using templates).

The concept aimed to explore the following questions: (1) When and how would radiologists want to interact with an AI-generated draft report, if at all? (2) Could there be utility to having a short form report (e.g., bullet points)? (3) Should the draft report be available to clinicians? If so, in what level of detail? (4) What is considered as ‘good enough’ AI performance for draft reports to be useful?

**6.2.2 Augmented Report Review.** Based on our insights on clinicians’ information needs, we explored how their review of radiology reports could be augmented with VLM capabilities. The Augmented Report Review concept (Figure 3) had two main features: (a) a report overview feature, and (b) an AI assistant feature. The report overview displayed a list of abnormal findings extracted from the report that can be visually highlighted in the patient image to facilitate its localization (e.g. *large right pleural effusion*). The AI assistant showcased numerous prompts inspired by clinician questions (e.g. *Given the image-based findings, what are the clinical guidelines for pleural effusion?*). For this concept, a critical design consideration was around latency: vision-language models are currently slow and costly. We speculated whether contextual queries can be pre-run prompts, where answers could be displayed immediately. As an alternative, we also sketched the AI assistant features as a chatbot with a text input field to provide contrasting options. The prototype displayed example prompts (e.g., guidelines, suggested investigations) as conversation starters to help clinicians envision what might be useful.

The concept aimed to explore: (1) Would clinicians want to review AI-generated annotations? If so, which findings are helpful to highlight for different image modalities (e.g., CT)? (2) Could there be any utility to having contextual information when reviewing images? (3) What would clinicians query? What would they never query? (4) Would there be a need for follow up queries (e.g., a chatbot style interaction that can maintain context)?

**6.2.3 Visual Search and Querying.** Building on the insight that radiologists and clinicians perform image searches online, the Visual Search and Querying concept (Figure 4) explored potential utility by displaying (a) a visual selection tool that enabled image search (e.g., *find similar images that look like this region*) or image and text queries (e.g., *Is this lump or anatomical variant?*). In line with recent literature showing clinicians look for evidence rather than explanations [138], we envisioned this concept to return groups of similar images instead of providing an interpretative answer (e.g., *Below are two groups of examples showing anatomic variants and lumps that look similar to the selected region.*) (Figure 4b).

The concept aimed to explore: (1) What would clinicians and radiologists visually query? (2) Could there be utility in performing image and text queries? (3) Would clinicians prefer to have an answer along with image examples (e.g., *Region likely normal?*)?

(4) What might be the data requirements for finding similar images (e.g., past images and reports from a hospital database)?

**6.2.4 Patient Imaging History Highlights.** Given that clinicians and radiologists commonly review patients’ prior images, the Patient Imaging History Highlights concept explored extracting and highlighting key insights across a patient’s image history. The prototype (Figure 5) displayed (a) a new X-ray scan, (b) prior images, and (c) an AI-generated summary of prior images and/or reports. Example highlights included changes in abnormalities (e.g., *Left lung nodule increased in size from 5 to 8 millimeters*); chronic conditions (e.g., *Chronic nodule in right lung benign, see image reference*); and operations (e.g., *Patient had chest drain on this date*).

The concept aimed to explore: (1) What is relevant to highlight in prior imaging summary? (2) Would a summary based only on reports be still useful; what is the least AI can do? (3) Would clinicians query prior images? If so, how (e.g., *Show me only abdomen CTs*)? (4) How would clinicians envision prior imaging summary to best be presented?

## 7 PHASE 3: ELICITING USER FEEDBACK

In the third phase, we sought feedback from a broader set of clinicians to understand whether, how and when the VLM-assisted radiology imaging concepts might be useful for clinical practice. This section reports participants’ feedback on each design concept, capturing perceived benefits and suggestions for improvement.

### 7.1 Draft Report Generation

**Expectation of near-perfect AI performance:** All radiologists expressed that having an AI-generated draft report would be valuable as long as the model performed really well; with high sensitivity and specificity. Describing how AI reporting errors could add burden, one radiologist explained: *“If it misses something, I’ve got to say that. If it’s false positive, I either have to click to remove it from the report entirely, or I have to change something.”* (R2) To better understand what would be considered as good enough AI performance for this use case, we asked *“Out of 10 reports, how many are you willing to correct?”* Almost all replied *“1 out of 10”* (R1, R2, R3) or *“5 to 10 out of 100”* (R5); suggesting the need for near-perfect performance for AI-generated draft reports to provide real utility. Only one radiologist, a trainee, responded *“3 out of 10”*, noting that the system could make them more confident even if it did not reduce their workload: *“It [would be] getting stuff right enough for me to feel comfortable just to edit the 30% of cases where it’s going to be wrong.”* (R4); suggesting potentially added benefits for trainee learning.

**Accounting for fast-paced practice & high workload:** Echoing our initial findings, radiologists noted that their practice is fast-paced and high volume: *“It is literally going as fast as humanly possible. Scrolling through things, looking at image, saying whatever I can, go over the spellchecks. Make sure I didn’t say anything really wrong and then sign and get on the next one. ... I just need to get my job done fast. I don’t get paid more for quality.”* (R2). Consequently, participants mainly spoke of value as time savings, especially when reading multi-slice images such as those captured by CT that take significantly longer to review and report than i.e. X-rays, and images that are outside of their subspecialty (R1, R2, R3, R5): *“I might*



be a seasoned reporter for lung or cardiac, but as every week it happens, we'll get a neck CT ... when you're not doing it day in day out, it's extremely difficult. You would love an AI which is at least giving you the salient findings." (R5) This suggests a draft report may reduce risks of key clinical observations being missed and could assist with image interpretation confidence. Apart from time savings, participants also mentioned potential benefits in reduced cognitive burden. For simpler X-ray images, R2 for example mentioned: "I can do [X-rays] in 10 seconds... [but] there's the cognitive burden. Having to say the words and go through it all is painful." R4, who was a trainee, reflected that the main benefit of the system would be reducing reporting time rather than the time spent for image interpretation: "Regardless of what the system says, I'm still going to go through my same search patterns for the findings and interpreting those ... the only area where it's going to be saving time is in creating that draft [prose] report because then I don't have to worry about the wording and if I've missed something".

**Preference for short, standardized reporting:** Interestingly, when probed whether short form sentences could be useful, all radiologists shared that they prefer to work with bullet point style findings instead of prose text. Several participants highlighted the literature on *structured reporting*, which is proposed as a solution for improving report quality and consistency [45]:

*"The idea of a narrative report happened in 1898 and we've not moved on from it. It's full of hedging, it's full of weird language that only radiologists use: 'likely to be', 'cannot exclude'. [This is] what we should be moving away from rather than using the technology to reverse engineer the future into what we got."* (R3)

Commenting on how the bullet list findings in the prototype were presented, R1 reflected "My reporting style is much more telegraphic. So I'll say 'large right pleural effusion', that's exactly how I'd phrase. I wouldn't say 'there is' or 'is seen' or all those kinds of phrases. I don't think [they] are helpful, especially for findings." Similarly, R3 advocated for structured findings for consistency and objectivity: "Rather than saying 'suspected mild cardiomegaly', you say 'heart is enlarged' or 'heart enlarged', which is a statement. It may be right or wrong, but it's objective." All these suggest a preference for concise, accurate and consistent reporting over the historic use of more ambiguous prose text, something that AI reporting could assist in standardizing.

**Favoring prioritized findings & confidence indications to assist image interpretation:** Additionally, radiologists described the benefits of having findings structured by their clinical relevance and the systems' confidence in the generated outputs. For example, a systems capability to compare a current study to a patient's prior image enables ordering report findings by: what is new, what has changed or is unchanged, which gives important context to aid image interpretation and subsequent clinical action. For example, the sudden 'new' appearance of a pneumothorax would require urgent clinical attention whilst a reduction in consolidation in the patients chest upon pneumonia diagnosis may suggest that antibiotic treatment is working. Furthermore, all participants (R1, R2, R3, R5) suggested having confidence intervals to communicate AI uncertainty: "Rather than using 'likely to be', 'unlikely to be', 'possibly' ... 'Likely prostate cancer 4 out of 5', [which is] more robust

and easier to interpret." (R3) One radiologist suggested displaying the model confidence and ranking findings on this basis: "[Say for a finding] I don't totally agree, I don't disagree. But if it's confidence is only like 56%, I'm just going to knock that out." (R2)

**Impressions present key interpretative work:** While short form, structured reporting was preferred for findings, some radiologists (R1, R3) shared that having unstructured, prose text is more appropriate for the impression section which is the "non-objective, doctor bit" (R3): "The main focus of communication between us and the team taking care of the patient is that impression part of the report. So it's really important to me to have that correctly crafted." (R1) R5 reflected that findings could be useful, yet the impression will be more difficult to get right: "We get a lot of [outsourced] reports from teleradiology, which just tell you what the findings are. A clinician will want to know the clinical impression. ... Is a report better than no report? I think it is fine if it gets the findings right, even if it doesn't do all the synthesis clinically." Given the importance of the impression section and its broader interpretative work that may include additional contextual information, the feedback from our participants suggests that clinicians may want to remain in charge of this task; positioning AI's role closer to the extraction of relevant findings from an image rather than its overall interpretation.

**Broadening uses of (prose) draft reports:** When asked how an AI-generated draft report should be presented, all radiologists suggested having both bullet points and prose report presented together whereby bullet points serve to assist the review, and prose for communication: "I could just get rid of [a bullet point] and it takes it out of the report, that's great. Because editing at that level is so much easier than editing on the report." (R2) A few radiologists noted that a patient-facing report could also be generated based on the list of findings (R1, R3); suggesting additional use cases and user groups.

In response to making an AI-generated draft report available to clinicians, all radiologists thought the AI-generated report could be useful for triage purposes, especially in situations where clinicians could escalate cases – as long as it did not look *too final*: "The subtlety there is that a draft report sounds too final in the health culture. But a 'prelim' or a 'wet read', that's a very rough, not final thing. The clinicians would take that information and use their judgement to call the radiologist or wait for the report." (R2) Alongside legal, regulatory and other organizational requirements to approve any such AI use, this requires a system design that appropriately communicates and clearly discloses the nature of preliminary AI-generated contents.

## 7.2 Augmented Report Review

**Locating image findings & their prioritization by clinical relevance:** Exploring how VLM capabilities could be utilized to augment the experiences of clinicians when reviewing the radiology report, all described finding image annotations helpful, especially for complex images like CTs. Most clinicians shared that they do not receive training to read CTs: "I look at CT scans, but I'm not trained to look at CT scans. I'm trained to look at X-rays." (C5) Some (C3, C6, C7) noted that they are comfortable reading CTs mainly within their subspecialty: "[In a brain scan] I would 100% be able to localize where things are. But if it was a report of a liver I would struggle." (C7) They pointed out that for such multi-slice images, current systems

require them to manually navigate to the image slice indicated in the report to view abnormalities. Having “clickable” findings, either on the report itself or in an overview section, that would direct them to the image location of relevance, was perceived valuable to save time and make it easier to differentiate what is in the image: “[Looking at a CT scan that had multiple areas of edema infarction] As a clinician, you’re like, well, this must be the bit that’s bleeding, but this must be the inflamed bit. But they look similar to me.” (C1) Clinicians additionally described several abnormalities that can be difficult to interpret: “Lymph nodes are the thing that people often miss on chest X-rays. Small pneumothoraces are difficult to see. The difference between a pneumothorax and a bullae [is] a common problem with the misreading of chest x-rays.” (C6) As such, they ascribed value to AI image annotations in aiding their understanding of the reported findings. Furthermore, similar to radiologists’ feedback, clinicians reflected that an overview section could highlight the most important and actionable findings: “Report overview would work best if you constrain it to show the top 6 salient features. We can get a lot of information overload if there are 25 of them.” (C7)

**Building an appropriate mental model of the AI:** When discussing more broadly how AI assistance could feature within workflows, one clinician differentiated for example a radiology assistant from a clinical assistant, whereby the former is embedded within the image viewer for radiology-specific tasks, whereas the latter –which is conceived as answering broader clinical questions– would be expected to sit within the EHR system: “If I’ve got a radiologist at my fingertips, I’d restrict to asking it the kind of questions I might be asking the radiologist. Therefore it belongs in [the radiology] screen, whereas some of the other things like, how should I treat this patient? I think that belongs in the main body of EHR rather than in this radiology reporting system.” (C4) This commentary highlights the importance of workflow integration for building an appropriate mental model of the AI’s likely purpose and capabilities.

**Cautioning about chat format & too complex queries:** In response to the AI assistant embodied as chatbot, several clinicians (C1, C3, C5, C7) commented that they were unlikely to use an assistant in chat form due to time-demands and lack of trust in generated, potentially high-risk responses: “I don’t need a chatbot function where I’m talking and stuff. I haven’t got the time for it.” (C5) Some clinicians raised concerns about responsibility in clinical decision making: “I’m not all of a sudden going to ask ChatGPT ‘What am I going to do with the brain tumor?’ I’m going to ask my friend who’s a specialist of this. There’s a question of responsibility.” (C1) Similarly, in answers to questions what clinicians would not want to use an AI assistant for (whether in chat or any other form), C7 – an oncologist – emphasized that he would not use it as a prognostic tool: “The radiology assistant shouldn’t be used to make predictions. It’s not a radiomic analysis in that sense.” Similarly, a cardiothoracic physician indicated that she would not ask what’s unknowable: “You wouldn’t ask things that are impossible to know. Things that are too complicated, like [the patient is] on six other drugs, how are they going to interact in combination? I wouldn’t bother asking, I wouldn’t trust the answer cause it’s too individualized.” (C6) Another concern was around the reinforcement of radiology observations that present negative findings. Here, clinicians stressed that they weigh positive findings more than negatives: “[If someone asks] ‘Can you confirm there really isn’t a small pneumothorax on this?’

Then the answer from the assistant should be ‘No, you can’t.’” (C7) In other words, clinicians cautioned the uses of AI for more ambitious, high-risk VLM use cases involving prognosis, more complex patient cases, or a definite negation of abnormalities – given more likely chances of errors and their negative implications on patient care.

**Focusing on task- and patient-specific, functional queries:** However, clinicians described an array of rather functional, task-specific queries where they could imagine AI to assist by either connecting them to, or extracting information on their behalf. For example, clinicians envisioned the AI assistant to perform image-based quantifications such as calculations of the cardiothoracic ratio (calculated by measuring the maximum diameter of the heart and thoracic cavity); Mirels’ score (indicating the risk of bone fracture); sarcopenia index (muscle-fat ratio to track weight loss in cancer patients); and waist-to-hip ratio in CT scans. All of these are currently calculated manually, often using phone apps: “It would be perceived added value if it could be quickly extracted from [an image] read, as you wouldn’t calculate it unless you needed.” (C7) In keeping with these more functional tasks, participants often envisioned AI assisting interactions in familiar forms, such as tool buttons, alerts or reminders for specific conditions and workflows; thereby describing expectations of the AI being designed as a workflow tool. One clinician expressed: “I almost would want the prompt ‘Have you thought about this?’” (C5) whilst simultaneously cautioning that such prompts could easily become annoying: “[For guidelines] I want to be able to click [on a finding], guidance, then it searches and brings it up for me. I don’t want pop-up fatigue.” (C5)

Furthermore, clinicians described how such practical, patient-specific AI functionality could be achieved even more effectively if VLM capabilities were combined with patient EHR data:

“You want it to give you, here’s their allergies, here’s their weight, here’s their renal function, here’s their swallow plan. Do they have a cannula in place? And here’s their other medications that could interact with that medication. If it can pull from the system that type of information, excellent, you’re saving me a huge amount of time.” (C5)

Criticizing many of the more generic information that were probed in our concept sketch (e.g., clinical features, differential diagnoses), clinicians emphasized the benefits of including additional EHR data to provide patient-context relevant information: “I don’t need [it to remind me] the 10 common causes of pleural effusion. What will be really helpful is for it to know that actually in this context, hypothyroidism becomes not the 29th thing, but actually upping [that to] your top five you should be considering ... because this patient’s got some other clues or signs.” (C3) Similarly, surfacing a patient’s eligibility for clinical trials or surfacing specific hospital or NHS level guidelines were described useful (C1, C2, C5, C6, C7); re-emphasizing the need for AI information provision to be adapted to each patient’s specific context.

### 7.3 Visual Search and Querying

**Aiding interpretation via comparison with relevant patient cases:** All clinicians and radiologists shared that they perform web searches to find similar images, though not too frequently (e.g., 1/week). For this concept, being able to visually search radiology

images and reports within the context of their hospital and patient population was valued the most: “Often you look at a CT scan on [internet] and you go ‘my CT scans don’t look anything like that’ [because it was a different generation CT scanner]. So it’s very important to visualize the abnormality in the context of the type of imaging you would see in your center.” (C7) Most clinicians and radiologists wanted to query what is normal, or queries with age and sex: “Recently we had a big debate: What does a 16 year old thymus look like normally?” (C6) An intensive care unit (ICU) clinician also described the difficulty of assessing rare conditions where they overlap with other abnormalities, because such cases are too infrequent and unfamiliar:

*“Nasogastric (NG) tubes in the wrong place on a chest X-ray on someone in ICU with pneumonia is even less common [than misplaced NG tubes alone]. So people have to simulate abnormalities in their head and compare the X-ray with their simulation. Showing [cases] similar to your patient would be useful.” (C1)*

All this suggests potential benefits of VLM use in retrieving or simulating other patient cases that enable comparative image assessments for either rare and complex (e.g., querying ‘NG tube’ + ‘pneumonia’), or normal cases to assist interpretation. For such uses, participants again positioned the AI system as a tool for extracting, searching or filtering information rather than as a conversational interface: “I’d have it as a tool that I can work with, and not conversation.” (R1) Describing how they would use queries to refine image search, one clinician added: “To then be able to type in pneumonia for example, and then the other [search results] go away. ‘Just female patients’ or ‘I’m only interested in people over 75.’” (C7)

**AI insights to provide reassurance to ‘human’ interpretation:** Reflecting on when in their workflow visual search and query capabilities could be useful, some clinicians suggested their use for follow-up questions about the radiology report: “Radiologist might have looked at it, but just not commented on it. I just want the reassurance, is that normal or not? Is it a nodule? Is it a mass? Is it a piece of consolidation? Same goes with head scans. Does this look like quite a full brain? Does the patient have hydrocephalus or not?” (C5). Yet, other clinicians reflected that even with AI functionality to retrieve i.e., similar images, they might still want to ask a radiologist to be assured: “Would I be reassured if it flashed up a whole load of other people’s chest X-rays and said, this was reported as normal and this was reported as normal, for yours is probably normal. I’m not sure that I would, but maybe.” (C6) Interestingly, none of the participants expected the system to provide an answer, and preferred example patient cases to inform their decisions: “Here’s a bunch of pictures, you decide. And that’s reasonable, right? I’m not asking some kind of segmentation to then take responsibility for the decisions.” (C1). This suggests preferences for AI use to reassure and aid human image interpretation rather than its use as an interpretative agent in itself.

## 7.4 Patient Imaging History Highlights

**Reducing laborious information gathering:** All radiologists and clinicians highly valued having a summary of a patient’s prior images highlighting key events and chronic conditions. Searching through a patient’s history was a major part of the clinical workflow. Recognizing the potential for time savings: “Half of my life is

*kind of spent chasing notes and pre-existing conditions. A sentence or two, just about the radiology, would save me a lot of time.” (C1)* Some clinicians (C3, C7) spoke of a time-reward trade-off: “The problem with image interpretation is, how far back do you look when interpreting for change?” (C7) They expressed feelings of guilt as they mostly look through recent reports, but not images, due to lack of time. Radiologists, on the other hand, shared they take a thorough look at past images, yet expressed desires for an automated summary: “That is a pretty standard practice already for radiologists, but certainly being able to more easily get at that imaging history is going to be a help.” (R1)

**Facilitating relevant patient information access:** Probing what would be useful to highlight, participants mainly described the historical status of the patient, such as the baseline lung architecture before a patient had pneumonia. Examples included past operations (e.g., *Do they have a collapsed lung?*), key events (e.g., *When their pacemaker first appeared or their sternotomy wires first went it?*) and changes in abnormalities (e.g., *New masses, fluid consolidation, rib fractures, are they old or new?*). When asked whether a text summary would be still useful in comparison to more multimodal, VLM capabilities (e.g., text summary of key events along with image annotations), most participants commented that linked reports and visual highlights could aid verification: “If you clicked on it [for it to show you annotated images], then you can corroborate.” (C6) However, several participants emphasized that even a text summary would provide an improvement to the current state: “We would willingly ingest that information even if it was a little bit more clunky.” (C7) Finally, a few clinicians (C3, C6) pointed out that unlike radiologists, the interface they use to review prior reports only presents a list view without images. As such, they thought of AI to still be useful if it could point them, at least, to important reports to guide their navigation to the relevant image: “I have to click on each one individually, wait for it to load. ... Even if I had a little red flag next to it saying ‘open this one, this has got money in it.’” (C3). This again highlights the prospective utility of AI in surfacing the clinically most relevant insights; and suggests that utility may already be achieved with simpler AI capabilities.

## 8 DISCUSSION

Aiming to narrow the gap between (multimodal) frontier AI model advances and their successful translation of into clinical practice [32, 44, 96, 99, 129, 131, 146], our work engaged in early phase design and user research to identify and co-create clinically relevant use cases of VLM capabilities for radiology. Below, we first discuss the findings from our user feedback sessions, detailing design requirements for each use case. We then share our broader reflections on these insights for human-AI interaction design in radiology and healthcare more generally. We conclude with our thoughts on the design process for ideating, prioritizing, and sketching VLM concepts.

### 8.1 Implications for VLMs in Radiology

Below, we expand on key learnings and critical design considerations that surfaced in our study to inform if, and how, the identified VLM use cases can assist radiologists and clinicians in their work.

**8.1.1 Chat and Question-Answering.** Our findings indicate that clinicians are unlikely to use *conversational* systems within their radiology work, especially to seek clinical decision support (e.g., making prognosis on complex patient cases) due to a lack of time to engage in dialogue and lack of trust in generated responses. Instead, participants in our study expressed a preference for *tool-based* interactions for well-defined tasks. This suggests the need for a broadened focus beyond ‘chatbots’ and ‘AI agents’ as concepts for how we envision AI to benefit clinical workflows; and requires the development of a richer design vocabulary for AI tooling in health-care. This is not to say that user acceptance of chat interactions could not evolve as use cases become more concrete and familiar. Furthermore, unlike clinicians who work under time pressure, conversational UIs might be more useful for patients in interacting with their medical imaging findings. Exploring if, and how, VLM capabilities could support other user groups marks a clear direction for future research.

**8.1.2 Image-Text Search.** Our study revealed that searching for similar images using credible web sources is a common practice in radiology image interpretation when reviewing cases with uncertainty. Both clinicians and radiologists expressed a desire to be able to search patient images and reports within their internal hospital databases, as images found online were often poor quality and did not reflect the local patient population. Echoing recent literature findings, participants in our study preferred the AI system to provide evidence rather than answers [138]. That being said, VLMs are unlikely to replace web search: participants emphasized that providing generic information would not be helpful, as such can easily and reliably be retrieved through web lookups. Instead, our design exploration focused on *comparative* queries (e.g., *Show me similar images that are pneumothorax versus bullae; What is considered normal in the context of a particular age group or rare condition*), which might allow for more flexible queries than web search. Future work should explore how VLM capabilities lend themselves to unique search, filter and retrieval experiences that go beyond what web search can provide.

**8.1.3 Report Generation.** In line with research on practice guidelines for radiology reporting [115], our findings surfaced the need for more effective, precise articulation of imaging findings. All expressed a preference for short form findings (e.g., bullet points) over prose, calling for more structured representations clearly indicating findings that are new, changed, unchanged, and normal. Interestingly, participants described the impression section as the “doctor bit”; the more carefully crafted, interpretative piece that may include other contextual information (cf. findings by [135]), and that acts as the main communication within the care team. Radiologists regarded the impression section as ‘harder to get right’ compared to the findings, which are based on observations alone. Interestingly, radiologists seemed comfortable with the idea of making draft reports available to clinicians, as long as the report did not look too final, but was a ‘wet read’. These insights suggest that AI research efforts for generating draft reports could find more acceptance for auto-generated findings, as opposed to an auto-generated impression (cf. [84]). Should the entire report be generated as a draft, or only certain parts? How best to juxtapose AI and human generated

contents for easy review and editing/correcting? All these remain open questions for future research.

For the draft report generation concept, we took inspiration from current practices whereby draft or preliminary reports by a junior or senior trainee are reviewed and amended if needed by a senior radiologists. The use of AI-generated draft reports brings up many questions in terms of the impact on trainees’ education and learning experience. While AI-generated draft reports may result in time savings, it could take away important opportunities for trainees developing their clinical reporting skills. On the other hand, an AI draft report could also be seen to provide a useful resource to assist trainees’ learning journey and confidence as they acquire image interpretation competencies. This design and research space, as well as questions on how to develop an appropriate AI reliance –upskilling trainees without over-dependency on AI– are complex and present open challenges for future work.

## 8.2 Implications for Designing AI in Healthcare

In this section, we take the learnings from our VLM design explorations to reflect on the question “What makes a good AI experience in healthcare?”

**8.2.1 Focusing on clinically useful AI applications.** Above and beyond the vast technical possibilities that new AI foundation models provide, this section describes the necessity of balancing these with clinical utility, risks of AI fallibility, and AI acceptance.

**Prioritizing lower risk-medium reward use cases:** Recent research highlighted that AI innovators mainly focus on use cases that require high task expertise and near-perfect AI performance (e.g. clinical decision making) [142]. Instead, researchers point to use cases where moderate-performance AI still could be useful in clinical tasks and workflows, such as in triage, workload management, and resource optimization. Echoing this [42, 142], our case study found that clinicians expect a near-perfect performance for use cases that require high clinician expertise (e.g., draft report generation). However, we also identified use cases that required medium-expertise, moderate-performance, yet were perceived as high value. For example, summarizing prior patient reports requires relatively lower expertise –a medical student level task– where having a ‘good enough’ summary could be still more useful than no summary. As argued elsewhere [19, 62, 118], this suggests a focus in AI development on simpler, more standardized use cases as a potentially lower risk and more responsible approach when starting to introduce AI innovations into clinical practice.

**AI as information resource, not clinical decision agent:** Amongst the many new capabilities afforded by VLMs (and other multimodal FMs) such as uses of natural language conversation; ability to generate or translate content; or predict specific outcomes; we found most uptake for proposals that involved relevant information retrieval or summarization. Notably, participants expressed low acceptance around ambitious, higher-risk VLM concepts, such as generalist medical AI uses for prognosis or to offer counterfactual reasoning due to a lack of trust in generated outputs; and none of them characterized our design proposals as clinical decision support that provides clinician interpretations or treatment recommendations. Instead, they recognized the potential value as reduced time and cognitive effort. For instance, draft report generation was

perceived as valuable for radiologist time savings (if correct), and for nudging clinicians to escalate cases to radiologists to seek assistance. Radiologists also described wanting to remain in charge of the ‘doctor bit’ and more complicated image interpretations that are harder for AI to get right.

**Empowering humans in their expert work:** As indicated above, instead of applying VLM capabilities as an interpretative agent, participants proposed their use for more functional, tedious tasks such as calculating medical ratios and scores; assessing organ sizes, volumes or density (cf. [72, 96]); retrieving EHR patient data; or for administrative support (cf. [109]). Outside these *practical tasks*, AI was positioned as *an information resource for assisting human interpretation* by extracting, highlighting or summarizing clinically relevant observations. Examples include AI use to provide evidence (e.g. via comparative queries and similar image retrieval, or locating image findings in complex CT scans); for ordering report findings by clinical relevance (or urgency); or deriving imaging history highlights. Conceiving of AI’s role as assisting, and ideally empowering, healthcare professionals in their expert work likely also plays a crucial role for its acceptance and clinical adoption.

**8.2.2 Integrating AI seamlessly into clinical workflows.** Intertwined with the above AI use cases are considerations of how best to integrate AI into fast-paced, high caseload workflows to provide utility whilst inviting appropriate frictions to check or correct AI outputs.

**Designing context-specific “workflow tools”:** Especially in responses to proposals for an ‘ask me anything’ type AI assistant feature, we found that such openness did not help clinicians in forming mental models; it was not clear to them what they could ask, or how the AI assistant would know the answers. Instead, they wanted the AI system to be more specialized and focus on specific tasks within their workflow (e.g., filtering search results, querying past patient reports). This aligns with user expectations around agent expertise that are well documented in the human-robot interaction literature [83, 106]. As a consequence, we suspect that current research aspirations for creating a ‘generalist medical AI’ [89, 128] are unlikely to correctly capture clinicians’ mental model, which presents a key consideration for the usability of such AI systems in healthcare. All this suggests that while (multimodal) AI can technically be leveraged for vastly different tasks, it may be beneficial to design them as context-specific workflow tools with clearly defined purposes to increase their understanding and practical utilization.

**Demonstrating clear benefits for disruption or change:** Given the growing complexity of VLMs and other advanced AI models, it is questionable how well, if at all, their workings can be explained, or their likelihood to fail be detected [125]. While it is common to think of human-in-the-loop approaches to verify and correct potentially fallible AI outcomes (e.g., review AI-generated draft reports), where such requirements add extra time burden to clinicians, or present tasks that they are less interested to perform, it risks reducing its utility and uptake in practice (cf. [60, 112, 113, 124]). If we were to accept that no AI model is always correct, designers need to give more consideration to strategies that enable clinicians not only to *easily verify* or *quickly correct VLM outputs* – beyond technical solutions (e.g., self-consistency prompting [118],

LLM-generated explanations [93], correctness predictions [65]); but also develop better strategies to *assist AI output justifications*. Re-visiting notions of AI uses as provision of evidence rather than clinical interpretations, connecting AI outputs to other criteria ‘external’ to the model [114] may aid clinicians to triangulate those outputs across different, trustworthy information sources such that they can more effectively be accepted or rejected as part of their work (cf. [123, 124]).

Further, we have to take into account that adoption of new AI-assisted practices can be greeted with reluctance where clinicians are asked to change established practices (e.g., move away from prose sentence dictation). Readiness to adapt work styles to accommodate AI likely requires sufficient benefits of a new approach.

### 8.3 Sketching VLM Experiences

Research investigating the best practices around designing AI products and services noted emergent approaches that blend human-centered and tech-centered processes, and the use of AI capabilities for sensitizing domain stakeholders [142, 143]. Our case study demonstrates that a capability-based approach – starting with both user needs and AI capabilities to find matches in the problem-solution space – proved effective for multidisciplinary brainstorming to identify clinically relevant AI use cases. Moreover, using multiple sketches that are framed as *instantiations of capabilities* rather than concrete design proposals to probe clinicians and radiologists (e.g., *Knowing that AI can do this, can you think of situations where this capability would be useful?*) seemed to work well.

While *sketching* [22] with VLM capabilities scaffolded ideation, separating the underlying capability from the form was a challenge. For example, the AI literature often uses the term ‘Visual Question-Answering’ to refer to AI tasks around image-to-text or text-to-image capabilities, yet these capabilities do not necessarily require a conversational form. Similarly, we struggled to envision novel VLM interactions that go beyond chatbots, alerts, and recommenders, a well-known challenge in AI design literature [139]. We approached this challenge by framing VLM capabilities as *queries* that can be formed in different ways (e.g., conversational questions, pre-run prompts, alerts, visual annotations, etc). Interestingly, the way participants described VLM interactions resembled *robotic process automation*: AI that fetches data in the background and presents it in an unremarkable manner [140] that can either be included or easily ignored. These findings point to a need for new design patterns beyond current paradigms of LLM or VLM uses as chat or conversational queries – especially in workflow-oriented contexts.

Additionally, VLMs as a design material centers considerations on balancing a seamless user experience with the time latency and financial costs (e.g., pre-running complex queries on large volumes of data to surface what prompts might be relevant; applying self-verification strategies to reduce risks of AI errors). This may determine choices to prioritize smaller, more efficient VLM models. Furthermore, designers should evaluate whether VLM (or other multimodal AI) capabilities are truly needed and appropriate for a task and explore alternatives. During our user feedback sessions, we probed into VLM boundaries by asking ‘Can there be simpler, dumber versions of these concepts?’ The Patient Imaging History Highlights concept demonstrates this approach well: while

a multimodal model can summarize rich patient image-report data; text-only models may already create value by summarizing previous report texts or pointing to important reports without text extraction and summarization.

Finally, we see opportunities for design research to investigate how to effectively sketch and prototype VLM interactions. In this paper, we utilized click-through sketches to scaffold clinicians' thinking around what AI can do and how the system might behave in specific use cases. While the concepts were used to probe into clinician expectations and AI acceptance, further research is needed to substantiate, test, and challenge our insights. Recent literature highlights *prompt prototyping* as a potential research direction for experience prototyping with LLMs and generative AI capabilities [63, 102, 130]. Future work should detail and refine interactions for the identified use cases; and extend insights into challenges of workflow integration, task completion time, as well as error types and their likely implications and mitigation.

## 9 CONCLUSION

Intersecting the fields of HCI, AI and Healthcare, the work in this paper presents a first investigation into the potential utility and design requirements for leveraging vision-language model (VLM) capabilities in the context of radiology. To this end, we conducted a three phase study that involved brainstorming with clinical experts and sketching four specific VLM use cases. Our findings from feedback sessions with 13 clinicians and radiologists provide initial insights into clinician acceptance and desirability for various identified VLM use cases, and advance this research by capturing nuanced design considerations. Against this backdrop, our work surfaced broader insights and challenges for human centered-AI systems in healthcare. While much emphasis is currently placed on developing more general-purpose AI models that can flexibly be adapted and scaled across different healthcare contexts, our research highlights the importance of bringing new AI capabilities into the focus of specific, practical tasks and use contexts to achieve effective workflow integration and the formation of useful mental models for AI and its intended uses. Notably, we found lesser interest in more ambitious VLM concepts that may offer value in terms of predictive diagnosis or counterfactual reasoning on medical outcomes. Instead, participants positioned AI as being assistive to health experts' work to help with mundane information extraction and processing tasks; thereby serving as a resource for human interpretation (e.g., by offering access to clinical evidence in a patient-specific manner). We further highlight the various trade-offs that are needed to ensure that AI's utility is balanced with the cost of AI risks; human effort in checking or correcting AI outputs; changes in work practices; as well as AI output latency and compute requirements. We encourage HCI researchers to further explore the benefits, risks, and limitations of VLMs in radiology workflows, and healthcare in general.

## ACKNOWLEDGMENTS

We thank our participants for their time and valuable input.

## REFERENCES

- [1] Open AI. 2022. chatGPT. <https://chat.openai.com>

- [2] Uwa O Aideyan, Kevin Berbaum, and Wilbur L Smith. 1995. Influence of prior radiologic information on the interpretation of radiographic examinations. *Academic Radiology* 2, 3 (1995), 205–208.
- [3] Tariq Osman Andersen, Francisco Nunes, Lauren Wilcox, Enrico Coiera, and Yvonne Rogers. 2023. Introduction to the Special Issue on Human-Centred AI in Healthcare: Challenges Appearing in the Wild. , 11 pages.
- [4] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403* (2023).
- [5] Matan Atad, Vitalii Dmytrenko, Yitong Li, Xinyue Zhang, Matthias Keicher, Jan Kirschke, Bene Wiestler, Ashkan Khakzar, and Nassir Navab. 2022. Chexplaining in style: Counterfactual explanations for chest x-rays using stylegan. *arXiv preprint arXiv:2207.07553* (2022).
- [6] Amid Ayobi, Jacob Hughes, Christopher J Duckworth, Jakob J Dylag, Sam James, Paul Marshall, Matthew Guy, Anitha Kumar, Adriane Chapman, Michael Boniface, et al. 2023. Computational Notebooks as Co-Design Tools: Engaging Young Adults Living with Diabetes, Family Carers, and Clinicians with Machine Learning Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [7] Anne Kathrine Petersen Bach, Trine Munch Nørgaard, Jens Christian Brok, and Niels van Berkel. 2023. "If I Had All the Time in the World": Ophthalmologists' Perceptions of Anchoring Bias Mitigation in Clinical AI Support. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [8] Shruthi Bannur, Stephanie Hyland, Qianchu Liu, Fernando Perez-Garcia, Maximilian Ilse, Daniel C Castro, Benedikt Boecking, Harshita Sharma, Kenza Bouzid, Anja Thieme, et al. 2023. Learning to exploit temporal structure for biomedical vision-language processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15016–15027.
- [9] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M Vardoulakis. 2020. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–12.
- [10] Andrew Bell, Oded Nov, and Julia Stoyanovich. 2023. Think about the stakeholders first! Toward an algorithmic transparency playbook for regulatory compliance. *Data & Policy* 5 (2023), e12.
- [11] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [12] Mélanie Bernhardt, Charles Jones, and Ben Glocker. 2022. Potential sources of dataset bias complicate investigation of underdiagnosis by machine learning algorithms. *Nature Medicine* 28, 6 (June 2022), 1157–1158. <https://doi.org/10.1038/s41591-022-01846-8>
- [13] Michael H Bernstein, Michael K Atalay, Elizabeth H Dibble, Aaron WP Maxwell, Adib R Karam, Saurabh Agarwal, Robert C Ward, Terrance T Healey, and Grayson L Baird. 2023. Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. *European Radiology* (2023), 1–7.
- [14] Leonard L Berry, Sunjay Letchuman, Nandini Ramani, and Paul Barach. 2021. The high stakes of outsourcing in health care. In *Mayo Clinic Proceedings*, Vol. 96. Elsevier, 2879–2890.
- [15] Mary Jo Bitner, Amy L Ostrom, and Felicia N Morgan. 2008. Service blueprinting: a practical technique for service innovation. *California management review* 50, 3 (2008), 66–94.
- [16] Sara Bly and Elizabeth F Churchill. 1999. Design through matchmaking: technology in search of users. *interactions* 6, 2 (1999), 23–31.
- [17] Benedikt Boecking, Naoto Usuyama, Shruthi Bannur, Daniel C Castro, Anton Schwaighofer, Stephanie Hyland, Maria Wetscherek, Tristan Naumann, Aditya Nori, Javier Alvarez-Valle, et al. 2022. Making the most of text semantics to improve biomedical vision–language processing. In *European conference on computer vision*. Springer, 1–21.
- [18] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [19] Claus Bossen and Kathleen H Pine. 2023. Batman and Robin in Healthcare Knowledge Work: Human-AI Collaboration by Clinical Documentation Integrity Specialists. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023), 1–29.
- [20] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [21] Eleanor R Burgess, Ivana Jankovic, Melissa Austin, Nancy Cai, Adela Kapuścińska, Suzanne Currie, J Marc Overhage, Erika S Poole, and Jofish Kaye. 2023. Healthcare AI Treatment Decision Support: Design Principles to Enhance Clinician Adoption and Trust. In *Proceedings of the 2023 CHI Conference on Human*



- Factors in Computing Systems*. 1–19.
- [22] Bill Buxton. 2010. *Sketching user experiences: getting the design right and the right design*. Morgan Kaufmann.
- [23] Carrie J Cai, Emily Reif, Narayan Hegde, Jason Hipp, Been Kim, Daniel Smilkov, Martin Wattenberg, Fernanda Viegas, Greg S Corrado, Martin C Stumpe, et al. 2019. Human-centered tools for coping with imperfect algorithms during medical decision-making. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–14.
- [24] Carrie J Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. "Hello AI": uncovering the onboarding needs of medical practitioners for human-AI collaborative decision-making. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [25] Carrie J Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2021. Onboarding Materials as Cross-functional Boundary Objects for Developing AI Assistants. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–7.
- [26] Francisco Maria Calisto, João Fernandes, Margarida Morais, Carlos Santiago, João Maria Abrantes, Nuno Nunes, and Jacinto C Nascimento. 2023. Assertiveness-based Agent Communication for a Personalized Medicine on Medical Imaging Diagnosis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [27] Francisco Maria Calisto, Carlos Santiago, Nuno Nunes, and Jacinto C Nascimento. 2021. Introduction of human-centric AI assistant to aid radiologists for multimodal breast image classification. *International Journal of Human-Computer Studies* 150 (2021), 102607.
- [28] Francisco Maria Calisto, Carlos Santiago, Nuno Nunes, and Jacinto C Nascimento. 2022. BreastScreening-AI: Evaluating medical intelligent agents for human-AI interactions. *Artificial Intelligence in Medicine* 127 (2022), 102285.
- [29] Isabel Chien, Nina Deliu, Richard Turner, Adrian Weller, Sofia Villar, and Niki Kilbertus. 2022. Multi-disciplinary fairness considerations in machine learning for clinical trials. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 906–924.
- [30] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416* (2022).
- [31] Neal J Clinger, Tim B Hunter, and Bruce J Hillman. 1988. Radiology reporting: attitudes of referring physicians. *Radiology* 169, 3 (1988), 825–826.
- [32] Enrico Coiera. 2019. The last mile: where artificial intelligence meets reality. *Journal of medical Internet research* 21, 11 (2019), e16323.
- [33] Eli Collins and Zoubin Ghahramani. 2021. LAMDA: Our breakthrough conversation technology. <https://blog.google/technology/ai/lamda/>
- [34] Ned Cooper, Tiffanie Horne, Gillian R Hayes, Courtney Heldreth, Michal Lahav, Jess Holbrook, and Lauren Wilcox. 2022. A systematic review and thematic analysis of community-collaborative approaches to computing research. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [35] Eric Corbett, Emily Denton, and Sheena Erete. 2023. Power and Public Participation in AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–13.
- [36] Greg Corrado and Yossi Matias. 2023. Multimodal Medical Ai. <https://ai.googleblog.com/2023/08/multimodal-medical-ai.html>
- [37] Ian A Cowan, Sharyn LS MacDonald, and Richard A Floyd. 2013. Measuring and managing radiologist workload: Measuring radiologist reporting times using data from a Radiology Information System. *Journal of medical imaging and radiation oncology* 57, 5 (2013), 558–566.
- [38] Rikke Friis Dam and Teo Yu Siang. 2022. Affinity diagrams: How to cluster your ideas and reveal insights. <https://www.interaction-design.org/literature/article/affinity-diagrams-learn-how-to-cluster-and-bundle-ideas-and-facts>
- [39] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond "Add Diverse Stakeholders and Stir". *arXiv preprint arXiv:2111.01122* (2021).
- [40] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–23.
- [41] Wesley Hanwen Deng, Nur Yildirim, Monica Chang, Motahhare Eslami, Kenneth Holstein, and Michael Madaio. 2023. Investigating Practices and Opportunities for Cross-functional Collaboration around AI Fairness in Industry Practice. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 705–716.
- [42] KJ Kevin Feng, Maxwell James Coppock, and David W McDonald. 2023. How Do UX Practitioners Communicate AI as a Design Material? Artifacts, Conceptions, and Propositions. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2263–2280.
- [43] Figma. 2023. Figma: the collaborative interface design tool. <https://www.figma.com/>
- [44] Astrid Galsgaard, Tom Doorschodt, Ann-Louise Holten, Felix Christoph Müller, Mikael Ploug Boesen, and Mario Maas. 2022. Artificial intelligence and multi-disciplinary team meetings: a communication challenge for radiologists' sense of agency and position as spider in a web? *European Journal of Radiology* 155 (2022), 110231.
- [45] Dhakshinamoorthy Ganesan, Phuong-Anh Thi Duong, Linda Probyn, Leon Lenchik, Tatum A McArthur, Michele Retrouvey, Emily H Ghobadi, Stephane L Desouches, David Pastel, and Isaac R Francis. 2018. Structured reporting in radiology. *Academic radiology* 25, 1 (2018), 66–73.
- [46] Pratik Ghosh, Karen L Posner, Stephanie L Hyland, Wil Van Cleve, Melissa Bristow, Dustin R Long, Konstantina Palla, Bala Nair, Christine Fong, Ronald Pauldine, et al. 2023. Framing Machine Learning Opportunities for Hypotension Prediction in Perioperative Care: A Socio-Technical Perspective. *ACM Transactions on Computer-Human Interaction* (2023).
- [47] Stephen Gilbert, Hugh Harvey, Tom Melvin, Erik Vollebregt, and Paul Wicks. 2023. Large language model AI chatbots require approval as medical devices. *Nature Medicine* (2023), 1–3.
- [48] Google. 2023. Bard - Chat Based AI Tool from Google, Powered by PaLM 2. <https://bard.google.com/>
- [49] Hongyan Gu, Yuan Liang, Yifan Xu, Christopher Kazu Williams, Shino Magaki, Negar Khanlou, Harry Vinters, Zesheng Chen, Shuo Ni, Chunxu Yang, et al. 2023. Improving workflow integration with XPath: Design and evaluation of a human-AI diagnosis system in pathology. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023), 1–37.
- [50] Hongyan Gu, Chunxu Yang, Mohammad Haeri, Jing Wang, Shirley Tang, Wenzhong Yan, Shujin He, Christopher Kazu Williams, Shino Magaki, and Xi-ang 'Anthony' Chen. 2023. Augmenting Pathologists with NaviPath: Design and Evaluation of a Human-AI Collaborative Navigation System. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [51] Yu Gu, Sheng Zhang, Naoto Usuyama, Yonas Woldesenbet, Cliff Wong, Praneth Sanapathi, Mu Wei, Naveen Valluri, Erika Strandberg, Tristan Naumann, et al. 2023. Distilling large language models for biomedical knowledge extraction: A case study on adverse drug events. *arXiv preprint arXiv:2307.06439* (2023).
- [52] Katharine E Henry, Rachel Kornfield, Anirudh Sridharan, Robert C Linton, Catherine Groh, Tony Wang, Albert Wu, Bilge Mutlu, and Suchi Saria. 2022. Human-machine teaming is key to AI adoption: clinicians' experiences with a deployed machine learning system. *NPJ digital medicine* 5, 1 (2022), 97.
- [53] Tad Hirsch, Kritzia Merced, Shrikanth Narayanan, Zac E Imel, and David C Atkins. 2017. Designing contestability: Interaction design, machine learning, and mental health. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. 95–99.
- [54] Tad Hirsch, Christina Soma, Kritzia Merced, Patty Kuo, Aaron Dembe, Derek D Caperton, David C Atkins, and Zac E Imel. 2018. "It's hard to argue with a computer" Investigating Psychotherapists' Attitudes towards Automated Evaluation. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 559–571.
- [55] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudik, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–16.
- [56] Karen Holtzblatt and Hugh Beyer. 2014. Field research: data collection and interpretation. In *Contextual Design: Evolved*. Springer, 11–20.
- [57] Jonathan Huang, Luke Neill, Matthew Wittbrodt, David Melnick, Matthew Klug, Michael Thompson, John Bailitz, Timothy Loftus, Sanjeev Malik, Amit Phull, et al. 2023. Generative Artificial Intelligence for Chest Radiograph Interpretation in the Emergency Department. *JAMA network open* 6, 10 (2023), e2336100–e2336100.
- [58] Stephanie Hyland, Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Mercy Ranjit, Anton Schwaighofer, Fernando Pérez-García, et al. 2023. MAIRA-1: A specialised large multimodal model for radiology report generation. *arXiv preprint arXiv:2311.13668* (2023).
- [59] Azra Ismail, Naveena Karusala, and Neha Kumar. 2018. Bridging disconnected knowledges for community health. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–27.
- [60] Maia Jacobs, Jeffrey He, Melanie F. Pradier, Barbara Lam, Andrew C Ahn, Thomas H McCoy, Roy H Perlis, Finale Doshi-Velez, and Krzysztof Z Gajos. 2021. Designing AI for trust and collaboration in time-constrained medical decisions: a sociotechnical lens. In *Proceedings of the 2021 chi conference on human factors in computing systems*. 1–14.
- [61] Katharina Jeblick, Balthasar Schachtner, Jakob Dextl, Andreas Mittermeier, Anna Theresa Stüber, Johanna Topalis, Tobias Weber, Philipp Wesp, Bastian Sabel, Jens Ricke, et al. 2022. Chatgpt makes medicine easy to swallow: An exploratory case study on simplified radiology reports. *arXiv preprint arXiv:2212.14882* (2022).
- [62] Saurabh Jha and Eric J Topol. 2016. Adapting to artificial intelligence: radiologists and pathologists as information specialists. *Jama* 316, 22 (2016), 2353–2354.
- [63] Ellen Jiang, Kristen Olson, Edwin Toh, Alejandra Molina, Aaron Donsbach, Michael Terry, and Carrie J Cai. 2022. Promptmaker: Prompt-based prototyping with large language models. In *CHI Conference on Human Factors in Computing*

- Systems Extended Abstracts*. 1–8.
- [64] Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Roger G Mark, and Steven Horng. 2019. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* 6, 1 (2019), 317.
- [65] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221* (2022).
- [66] Charles E Kahn Jr, Curtis P Langlotz, Elizabeth S Burnside, John A Carrino, David S Channin, David M Hovsepian, and Daniel L Rubin. 2009. Toward best practices in radiology reporting. *Radiology* 252, 3 (2009), 852–856.
- [67] Holtzblatt Karen and Jones Sandra. 2017. Contextual inquiry: A participatory technique for system design. In *Participatory design*. CRC Press, 177–210.
- [68] Kundan Krishna, Sopan Khosla, Jeffrey P Bigham, and Zachary C Lipton. 2020. Generating SOAP notes from doctor-patient conversations using modular summarization techniques. *arXiv preprint arXiv:2005.01795* (2020).
- [69] Sean Kross and Philip Guo. 2021. Orienting, framing, bridging, magic, and counseling: How data scientists navigate the outer loop of client collaborations in industry and academia. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–28.
- [70] Tzu-Sheng Kuo, Hong Shen, Jisoo Geum, Nev Jones, Jason I Hong, Haiyi Zhu, and Kenneth Holstein. 2023. Understanding Frontline Workers' and Unhoused Individuals' Perspectives on AI Used in Homeless Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [71] Curtis P. Langlots. 2015. *The radiology report: a guide to thoughtful communication for radiologists and other medical professionals*. Springer.
- [72] Curtis P Langlotz. 2019. Will artificial intelligence replace radiologists? , e190058 pages.
- [73] Robert Law. 2014. Radiographers,'never events' and the nasogastric tube. *Radiography* 20, 1 (2014), 2–3.
- [74] Augustin Lecler, Loïc Duron, and Philippe Soyer. 2023. Revolutionizing radiology with GPT-based models: Current applications, future possibilities and limitations of ChatGPT. *Diagnostic and Interventional Imaging* 104, 6 (2023), 269–274.
- [75] Peter Lee, Sebastian Bubeck, and Joseph Petro. 2023. Benefits, Limits, and Risks of GPT-4 as an AI Chatbot for Medicine. *New England Journal of Medicine* 388, 13 (2023), 1233–1239.
- [76] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110* (2022).
- [77] Q Vera Liao, Hariharan Subramonyam, Jennifer Wang, and Jennifer Wortman Vaughan. 2023. Designerly understanding: Information needs for model transparency to support design ideation for AI-powered user experience. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–21.
- [78] Q Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *arXiv preprint arXiv:2306.01941* (2023).
- [79] Q Vera Liao, Yunfeng Zhang, Ronny Luss, et al. 2022. Connecting Algorithmic Research and Usage Contexts: A Perspective of Contextualized Evaluation for Explainable AI. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 10. 147–159.
- [80] Martin Lindvall, Claes Lundström, and Jonas Löwgren. 2021. Rapid assisted visual search: Supporting digital pathologists with imperfect AI. In *26th International Conference on Intelligent User Interfaces*. 504–513.
- [81] Houjiang Liu, Anubrata Das, Alexander Boltz, Didi Zhou, Daisy Pinaroc, Matthew Lease, and Min Kyung Lee. 2023. Human-centered NLP Fact-checking: Co-Designing with Fact-checkers using Matchmaking for AI. *arXiv preprint arXiv:2308.07213* (2023).
- [82] Qianchu Liu, Stephanie Hyland, Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Maria Teodora Wetscherek, Robert Tinn, Harshita Sharma, Fernando Pérez-García, Anton Schwaighofer, et al. 2023. Exploring the Boundaries of GPT-4 in Radiology. *arXiv preprint arXiv:2310.14573* (2023).
- [83] Michal Luria, Samantha Reig, Xiang Zhi Tan, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2019. Re-Embodiment and Co-Embodiment: Exploration of social presence for robots and conversational agents. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 633–644.
- [84] Chong Ma, Zihao Wu, Jiaqi Wang, Shaochen Xu, Yaonai Wei, Zhengliang Liu, Lei Guo, Xiaoyan Cai, Shu Zhang, Tuo Zhang, et al. 2023. ImpressionGPT: an iterative optimizing framework for radiology report summarization with chatGPT. *arXiv preprint arXiv:2304.08448* (2023).
- [85] Bella Martin, Bruce Hanington, and Bruce M Hanington. 2012. Universal methods of design: 100 ways to research complex problems. *Develop Innovative Ideas, and Design Effective Solutions* (2012), 12–13.
- [86] Stina Matthiesen, Søren Zøga Diederichsen, Mikkel Klitzing Hartmann Hansen, Christina Villumsen, Mats Christian Højbjerg Lassen, Peter Karl Jacobsen, Niels Risum, Bo Gregers Winkel, Berit T Philbert, Jesper Hastrup Svendsen, et al. 2021. Clinician preimplementation perspectives of a decision-support tool for the prediction of cardiac arrhythmia based on machine learning: near-live feasibility and qualitative study. *JMIR human factors* 8, 4 (2021), e26964.
- [87] Microsoft. 2023. Microsoft Copilot: Your everyday AI companion. <https://copilot.microsoft.com/>
- [88] Tim Miller, Piers Howe, and Liz Sonenberg. 2017. Explainable AI: Beware of inmates running the asylum or: How I learnt to stop worrying and love the social and behavioural sciences. *arXiv preprint arXiv:1712.00547* (2017).
- [89] Michael Moor, Oishi Banerjee, Zahra Shakeri Hossein Abad, Harlan M Krumholz, Jure Leskovec, Eric J Topol, and Pranav Rajpurkar. 2023. Foundation models for generalist medical artificial intelligence. *Nature* 616, 7956 (2023), 259–265.
- [90] Meredith Ringel Morris, Carrie J Cai, Jess Holbrook, Chinmay Kulkarni, and Michael Terry. 2023. The design space of generative models. *arXiv preprint arXiv:2304.10547* (2023).
- [91] Nabla. 2023. Nabla Copilot · Enjoy care again. <https://www.nabla.com/> [Accessed 11-08-2023].
- [92] Sandeep S Naik, Anthony Hanbidge, and Stephanie R Wilson. 2001. Radiology reports: examining radiologist and clinician preferences regarding style and content. *American Journal of Roentgenology* 176, 3 (2001), 591–598.
- [93] Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. 2023. Capabilities of gpt-4 on medical challenge problems. *arXiv preprint arXiv:2303.13375* (2023).
- [94] Nuance-Microsoft. 2023. Nuance and Microsoft Announce the First Fully AI-Automated Clinical Documentation Application for Healthcare — news.nuance.com. <https://news.nuance.com/2023-03-20-Nuance-and-Microsoft-Announce-the-First-Fully-AI-Automated-Clinical-Docmentation-Application-for-Healthcare>. [Accessed 11-08-2023].
- [95] Namkee Oh, Gyu-Seong Choi, and Woo Yong Lee. 2023. ChatGPT goes to the operating room: evaluating GPT-4 performance and its potential in surgical education and training in the era of large language models. *Annals of Surgical Treatment and Research* 104, 5 (2023), 269.
- [96] Nazmun Nisat Ontika, Sheree May Sassmannshausen, Aparecido Fabiano Pinatti De Carvalho, and Volkmar Pipek. 2023. PAIRADS: Hybrid Interaction Between Humans and AI in Radiology. In *HFAI 2023: Augmenting Human Intellect*. IOS Press, 395–397.
- [97] Nazmun Nisat Ontika, Hussain Abid Syed, Sheree May Safmannshausen, Richard HR Harper, Yunan Chen, Sun Young Park, Miria Grisot, Astrid Chow, Nils Blaumer, Aparecido Fabiano Pinatti de Carvalho, et al. 2022. Exploring human-centered AI in healthcare: diagnosis, explainability, and trust. (2022).
- [98] OpenAI. 2023. GPT-4 Technical Report. [arXiv:2303.08774](https://arxiv.org/abs/2303.08774) [cs.CL]
- [99] Tariq Osman Andersen, Francisco Nunes, Lauren Wilcox, Elizabeth Kazianas, Stina Matthiesen, and Farah Magrabi. 2021. Realizing AI in healthcare: challenges appearing in the wild. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–5.
- [100] Bhavik N Patel, Louis Rosenberg, Gregg Wilcox, David Baltaxe, Mimi Lyons, Jeremy Irvin, Pranav Rajpurkar, Timothy Amrhein, Rajan Gupta, Safwan Halabi, et al. 2019. Human-machine partnership with artificial intelligence for chest radiograph diagnosis. *NPJ digital medicine* 2, 1 (2019), 111.
- [101] Eike Petersen, Yannik Potdevin, Esfandiar Mohammadi, Stephan Zidowitz, Sabrina Breyer, Dirk Nowotka, Sandra Henn, Ludwig Pechmann, Martin Leucker, Philipp Rostalski, et al. 2022. Responsible and regulatory conform machine learning for medicine: a survey of challenges and solutions. *IEEE Access* 10 (2022), 58375–58418.
- [102] Savvas Petridis, Michael Terry, and Carrie Jun Cai. 2023. PromptInfuser: Bringing User Interface Mock-ups to Life with Large Language Models. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–6.
- [103] David Piorkowski, Soya Park, April Yi Wang, Dakuo Wang, Michael Muller, and Felix Portnoy. 2021. How ai developers overcome communication challenges in a multidisciplinary team: A case study. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–25.
- [104] Sam Preston, Mu Wei, Rajesh Rao, Robert Tinn, Naoto Usuyama, Michael Lucas, Yu Gu, Roshanthi Weerasinghe, Soohye Lee, Brian Piening, et al. 2023. Toward structuring real-world data: Deep learning for extracting oncology information from clinical text with patient-level supervision. *Patterns* 4, 4 (2023).
- [105] Rob Procter, Peter Tolmie, and Mark Rouncefield. 2023. Holding AI to account: Challenges for the delivery of trustworthy AI in healthcare. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023), 1–34.
- [106] Samantha Reig, Michal Luria, Janet Z Wang, Danielle Oltman, Elizabeth Jeanne Carter, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2020. Not Some Random Agent: Multi-person interaction with a personalizing service robot. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 289–297.
- [107] Abi Rimmer. 2017. Radiologist shortage leaves patient care at risk, warns royal college. *BMJ: British Medical Journal (Online)* 359 (2017).
- [108] Samantha Robertson and Niloufar Salehi. 2020. What If I Don't Like Any Of The Choices? The Limits of Preference Elicitation for Participatory Algorithm Design. *arXiv preprint arXiv:2007.06718* (2020).

- [109] NR Sahni, P Mishra, B Carrus, and DM Cutler. [n. d.]. Administrative Simplification: How to Save a Quarter-Trillion Dollars in US Healthcare. McKinsey & Company. October 20, 2021.
- [110] Alexandra Sasha Luccioni and Anna Rogers. 2023. Mind your Language (Model): Fact-Checking LLMs and their Role in NLP Research and Practice. *arXiv e-prints* (2023), arXiv:2308.
- [111] Sectra. 2013. How radiology can improve communication with referring physicians. <https://sectraprodstorage01.blob.core.windows.net/medical-uploads/2017/09/report-how-radiology-can-improve-communication-with-referring-physicians.pdf> [Accessed 11-22-2023].
- [112] Mandeep Sekhon, Martin Cartwright, and Jill J Francis. 2017. Acceptability of healthcare interventions: an overview of reviews and development of a theoretical framework. *BMC health services research* 17, 1 (2017), 1–13.
- [113] Mark Sendak, Madeleine Clare Elish, Michael Gao, Joseph Futoma, William Ratliff, Marshall Nichols, Armando Bedoya, Suresh Balu, and Cara O'Brien. 2020. "The human body is a black box" supporting clinical decision-making with deep learning. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 99–109.
- [114] Murray Shanahan. 2022. Talking About Large Language Models. *arXiv preprint arXiv:2212.03551* (2022).
- [115] C Sherry, M Adams, I Berlin, L Fajardo, G Gazelle, DB Haseman, et al. 2022. ACR practice guideline for communication of diagnostic imaging findings. *American College of Radiology* (2022).
- [116] Julia Simkus. 2023. Snowball sampling method: Definition, Techniques & Examples. <https://www.simplypsychology.org/snowball-sampling.html>
- [117] Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2022. Large Language Models Encode Clinical Knowledge. *arXiv preprint arXiv:2212.13138* (2022).
- [118] Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2023. Large language models encode clinical knowledge. *Nature* (2023), 1–9.
- [119] Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, et al. 2023. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617* (2023).
- [120] Lea Strohm, Charisma Hehakaya, Erik R Ranschaert, Wouter PC Boon, and Ellen HM Moors. 2020. Implementation of artificial intelligence (AI) applications in radiology: hindering and facilitating factors. *European radiology* 30 (2020), 5525–5532.
- [121] Hariharan Subramonyam, Jane Im, Colleen Seifert, and Eytan Adar. 2022. Solving separation-of-concerns problems in collaborative design of human-AI systems through leaky abstractions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [122] Anja Thieme, Danielle Belgrave, and Gavin Doherty. 2020. Machine learning in mental health: A systematic review of the HCI literature to support the development of effective and implementable ML systems. *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 5 (2020), 1–53.
- [123] Anja Thieme, Ed Cutrell, Cecily Morrison, Alex Taylor, and Abigail Sellen. 2020. Interpretability as a dynamic of human-AI interaction. *Interactions* 27, 5 (2020), 40–45.
- [124] Anja Thieme, Maryann Hanratty, Maria Lyons, Jorge Palacios, Rita Faia Marques, Cecily Morrison, and Gavin Doherty. 2023. Designing human-centered AI for mental health: Developing clinically relevant applications for online CBT treatment. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023), 1–50.
- [125] Anja Thieme, Aditya Nori, Marzyeh Ghassemi, Rishi Bommasani, Tariq Osman Andersen, and Ewa Luger. 2023. Foundation Models in Healthcare: Opportunities, Risks & Strategies Forward. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–4.
- [126] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [127] Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaeckermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Chuck Lau, Ryutarō Tanno, Ira Ktena, et al. 2023. Towards generalist biomedical ai. *arXiv preprint arXiv:2307.14334* (2023).
- [128] Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaeckermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Chuck Lau, Ryutarō Tanno, Ira Ktena, Basil Mustafa, Aakanksha Chowdhery, Yun Liu, Simon Kornblith, David Fleet, Philip Mansfield, Sushant Prakash, Renee Wong, Sunny Virmani, Christopher Semturs, S Sara Mahdavi, Bradley Green, Ewa Dominowska, Blaise Aguera y Arcas, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Karan Singhal, Pete Florence, Alan Karthikesalingam, and Vivek Natarajan. 2023. Towards Generalist Biomedical AI. arXiv:2307.14334 [cs.CL]
- [129] Stephanie Tulk Jesso, Aisling Kelliher, Harsh Sanghavi, Thomas Martin, and Sarah Henrickson Parker. 2022. Inclusion of clinicians in the development and evaluation of clinical artificial intelligence tools: a systematic literature review. *Frontiers in Psychology* 13 (2022), 830345.
- [130] Stephanie Valencia, Richard Cave, Krystal Kallarackal, Katie Seaver, Michael Terry, and Shaun K Kane. 2023. "The less I type, the better": How AI Language Models can Enhance or Impede Communication for AAC Users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [131] Himanshu Verma, Roger Schaefer, Julien Reichenbach, Mario Jreige, John O Prior, Florian Evéquo, and Adrien Depeursinge. 2021. On improving physicians' trust in AI: Qualitative inquiry with imaging experts in the oncological domain. *BMC Medical Imaging, in review* (2021).
- [132] Qiaosi Wang, Michael Madaio, Shaun Kane, Shivani Kapania, Michael Terry, and Lauren Wilcox. 2023. Designing Responsible AI: Adaptations of UX Practice to Meet Responsible AI Challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [133] Lauren Wilcox, Robin Brewer, and Fernando Diaz. 2023. AI Consent Futures: A Case Study on Voice Data Collection with Clinicians. (2023).
- [134] Malwina Anna Wójcik. 2022. Foundation Models in Healthcare: Opportunities, Biases and Regulatory Prospects in Europe. In *International Conference on Electronic Government and the Information Systems Perspective*. Springer, 32–46.
- [135] Yao Xie, Melody Chen, David Kao, Ge Gao, and Xiang'Anthony' Chen. 2020. CheXplain: enabling physicians to explore and understand data-driven, AI-enabled medical imaging analysis. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [136] Shawn Xu, Lin Yang, Christopher Kelly, Marcin Sieniek, Timo Kohlberger, Martin Ma, Wei-Hung Weng, Attila Kiraly, Sahar Kazemzadeh, Zakkai Melamed, et al. 2023. ELIXR: Towards a general purpose X-ray artificial intelligence system through alignment of large language models and radiology vision encoders. *arXiv preprint arXiv:2308.01317* (2023).
- [137] Qian Yang, Justin Cranshaw, Saleema Amershi, Shamsi T Iqbal, and Jaime Teevan. 2019. Sketching nlp: A case study of exploring the right things to design with language intelligence. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [138] Qian Yang, Yuexing Hao, Kexin Quan, Stephen Yang, Yiran Zhao, Volodymyr Kuleshov, and Fei Wang. 2023. Harnessing biomedical literature to calibrate clinicians' trust in AI decision support systems. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [139] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–13.
- [140] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–11.
- [141] Nur Yildirim, Alex Kass, Teresa Tung, Connor Upton, Donnacha Costello, Robert Giusti, Sinem Lacin, Sara Lovic, James M O'Neill, Rudi O'Reilly Meehan, et al. 2022. How Experienced Designers of Enterprise Applications Engage AI as a Design Material. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [142] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supriya Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, et al. 2023. Creating Design Resources to Scaffold the Ideation of AI Concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2326–2346.
- [143] Nur Yildirim, Mahima Pushkarna, Nitesh Goyal, Martin Wattenberg, and Fernanda Viégas. 2023. Investigating How Practitioners Use Human-AI Guidelines: A Case Study on the People+ AI Guidebook. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [144] Nur Yildirim, John Zimmerman, and Sarah Preum. 2021. Technical Feasibility, Financial Viability, and Clinician Acceptance: On the Many Challenges to AI in Clinical Practice. In *HUMAN@ AAAI Fall Symposium*.
- [145] Feiyang Yu, Mark Endo, Rayan Krishnan, Jan Pan, Andy Tsai, Eduardo Pontes Reis, Eduardo Kaiser Ururahy Nunes Fonseca, Henrique Min Ho Lee, Zahra Shakeri Hossein Abad, Andrew Y. Ng, Curtis P. Langlotz, Vasantha Kumar Venugopal, and Pranav Rajpurkar. 2023. Evaluating progress in automatic chest X-ray radiology report generation. *Patterns* 4, 9 (2023). <https://doi.org/10.1016/j.patter.2023.100802>
- [146] Kun-Hsing Yu, Andrew L Beam, and Isaac S Kohane. 2018. Artificial intelligence in healthcare. *Nature biomedical engineering* 2, 10 (2018), 719–731.
- [147] Hubert D Zajac, Dana Li, Xiang Dai, Jonathan F Carlsen, Finn Kensing, and Tariq O Andersen. 2023. Clinician-facing AI in the Wild: Taking Stock of the Sociotechnical Challenges and Opportunities for HCI. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023), 1–39.
- [148] Angie Zhang, Olympia Walker, Kaci Nguyen, Jiajun Dai, Anqing Chen, and Min Kyung Lee. 2023. Deliberating with AI: Improving Decision-Making for the Future through Participatory AI Design and Stakeholder Deliberation. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–32.
- [149] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 493–502.

## **10 APPENDIX A: EXAMPLE RADIOLOGY REPORT AND PROTOTYPE FLOWS**

This section contains (a) an example radiology report from the MIMIC-CXR dataset [64], and (b) the click-through prototypes used

for concepts where we had more than a single frame (*Draft Report Generation, Augmented Report Review, Visual Search and Querying*).

EXAMINATION: CHEST (PA AND LAT)  
 INDICATION: \_\_\_ year old woman with ?pleural effusion // ?pleural effusion  
 TECHNIQUE: Chest PA and lateral  
 COMPARISON: \_\_\_  
 FINDINGS:  
 Cardiac size cannot be evaluated. Large left pleural effusion is new. Small right effusion is new. The upper lungs are clear. Right lower lobe opacities are better seen in prior CT. There is no pneumothorax. There are mild degenerative changes in the thoracic spine  
 IMPRESSION:  
 Large left pleural effusion

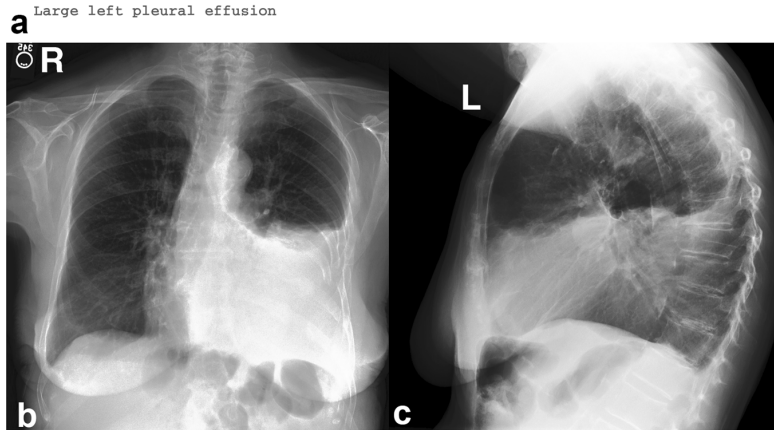


Figure 6: An example image-report pair from the MIMIC-CXR dataset [64] with (a) the radiology report and (b, c) the chest X-rays including the (b) frontal view and (c) the lateral view.

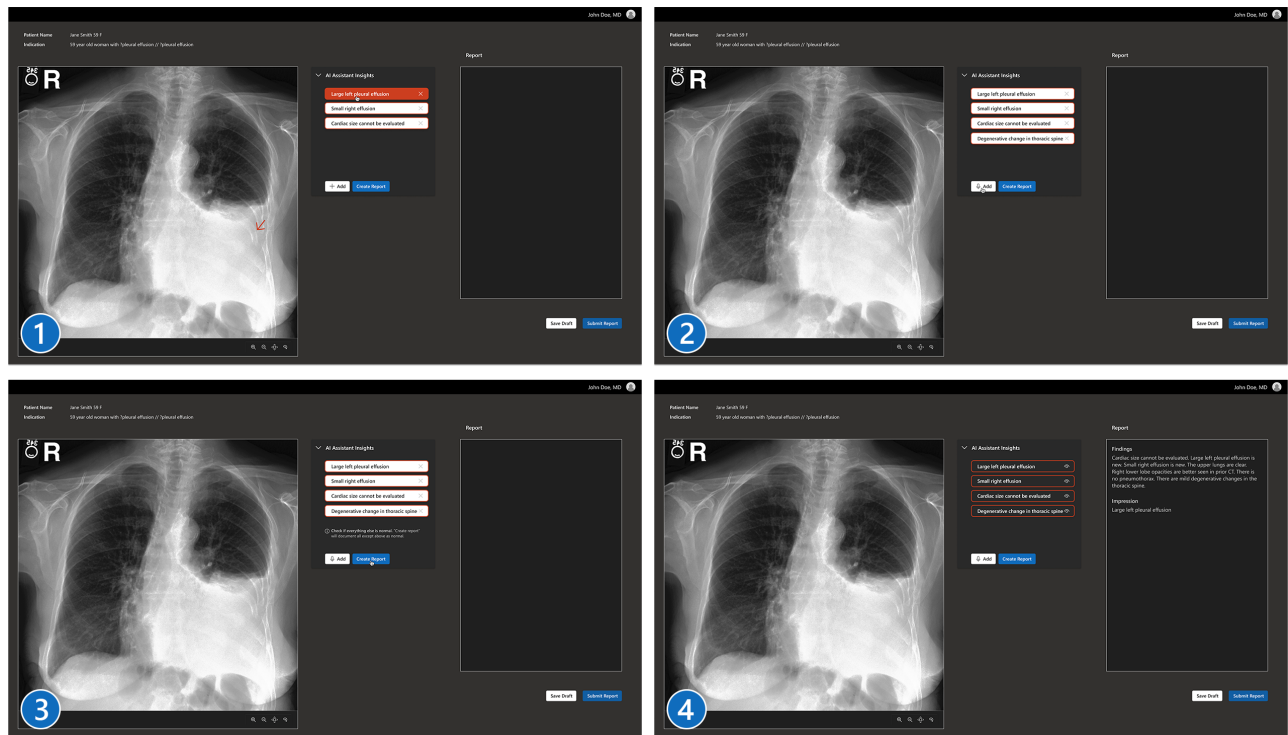


Figure 7: Click-through prototype flow illustrating the *Draft Report Generation* concept.

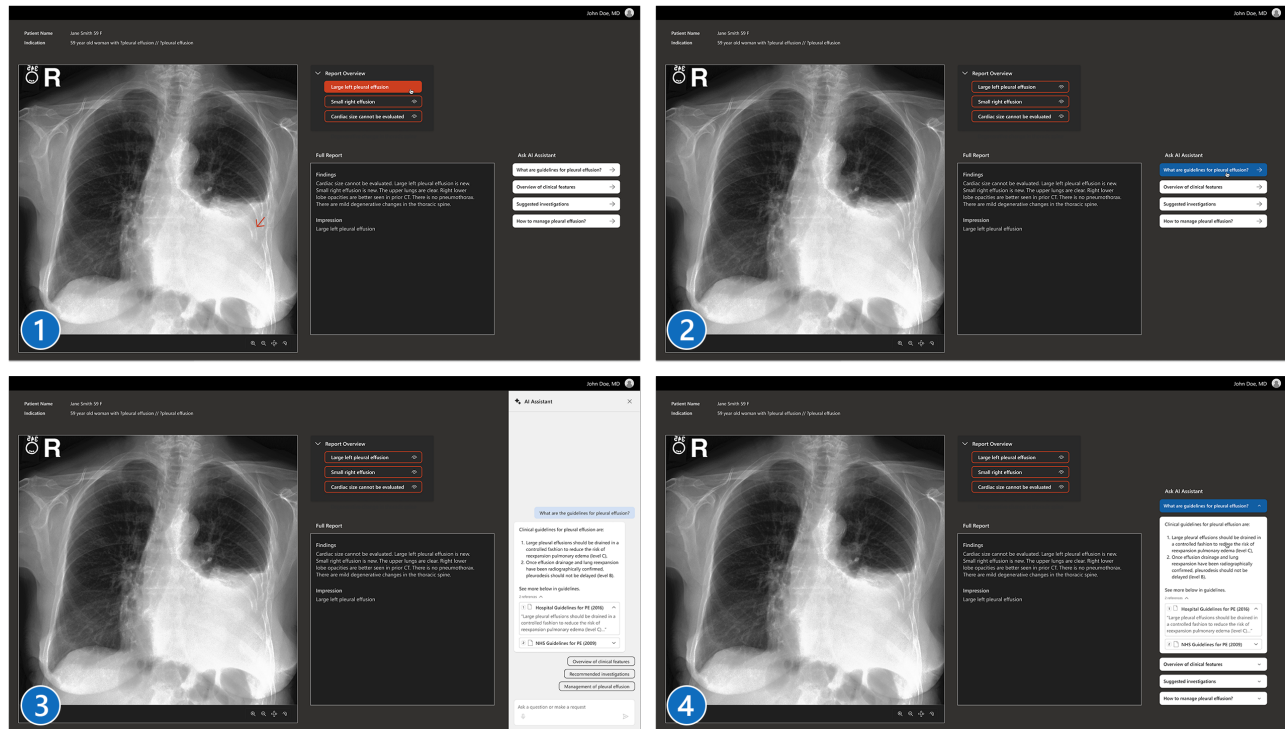


Figure 8: Click-through prototype flow illustrating the *Augmented Report Review* concept.

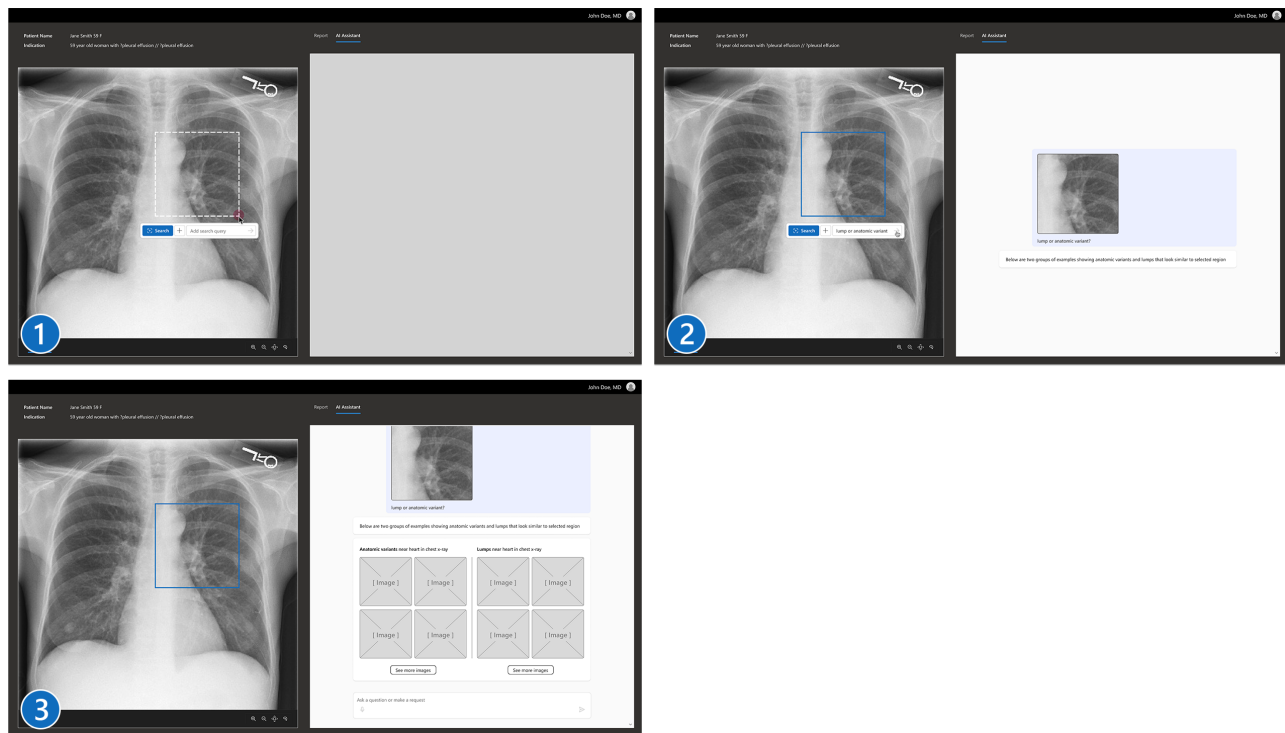


Figure 9: Click-through prototype flow illustrating the *Visual Search and Querying* concept.