Special report: A generalist's guide to the 5-year outlook for healthcare AI

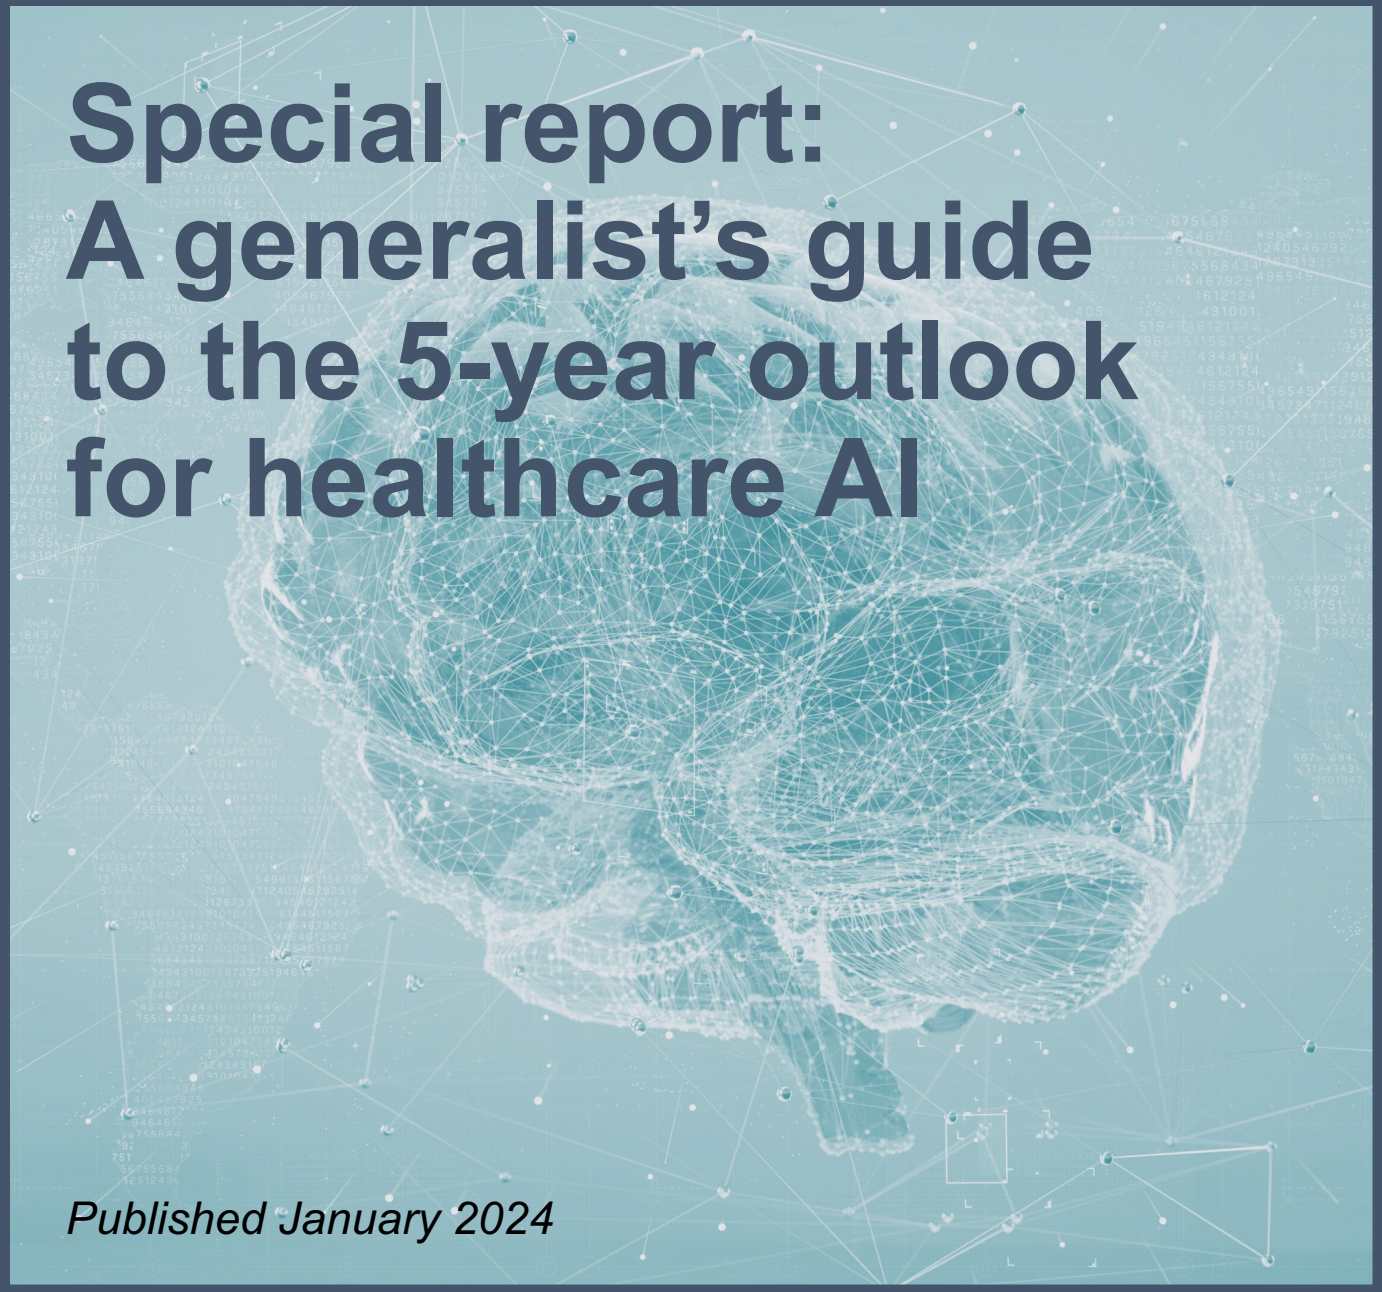

Published January 2024

UNION
HEALTHCARE INSIGHT

Table of contents

1 Introduction 

2 Six insights on the 5-year outlook for AI in healthcare 

3 Resource center: Regulatory guardrails on healthcare AI 

4 AI application feasibility calculator 

5 Conclusion 

Introduction

Achieving meaningful discussion of AI

AI discussion is everywhere, currently infusing the whole of the healthcare news and insight ecosystem. And yet, it remains an exceptionally difficult topic for many healthcare leaders to write, read, or talk about in a meaningful way, for several reasons (fig. 1).

First, pre-existing levels of topic expertise vary *widely*. This dooms most AI content to being set at the wrong level: either overly technical and overwhelming, or elementary and obvious. Second, the topic is particularly prone to stoking both dramatic hopes and fears. Skeptics doubt the practical benefits, worry about the risks, and feel AI is siphoning attention away from more pressing issues. Proponents may be frustrated by these objections because they perceive near-endless opportunities for AI to solve previously intractable industry problems. Finally, AI discussions tend to lack defined timeframes, making almost any impact *theoretically* possible, but also easy to challenge. This unmoors discussions, and feeds polarization—because it puts both dramatic benefits and disastrous consequences on the table.

To bring clarity, this research is aimed at a specific audience and time horizon. As always, we are addressing **the healthcare executive generalist**: individuals who are conversant in the industry at large but who constantly need updated understanding across the many strategically important domains they track. This includes general management executives, board members, and strategy leaders, among others. **We believe that this group is typically best served by an explicit three-to-five year outlook.** This timeframe ensures discussion is forward-looking, but not overly futuristic in nature. It also aligns with a typical strategic planning cycle. **To further ground the discussion, we emphasize a "commonsense" view** and a question/answer structure for the work. More on this shortly.

Fig. 1: Common challenges to having meaningful conversations about AI—and how this research addresses them

These principles (executive generalist audience; three-to-five year timeframe, and a 'commonsense' research frame) have guided both Union's larger project on the future of innovation and this AI-focused report.

Our methodology for this work included both primary and secondary research (fig 2). We interviewed over 20 healthcare executives from across industry sectors. We reviewed general innovation literature, as well as targeted information and viewpoints on AI.

>>>>>>

For more on the general innovation landscape, download "The future of innovation: A 5-year outlook" at unionhealthcareinsight.com/research

Fig. 2: Union's 'future of innovation' research methodology

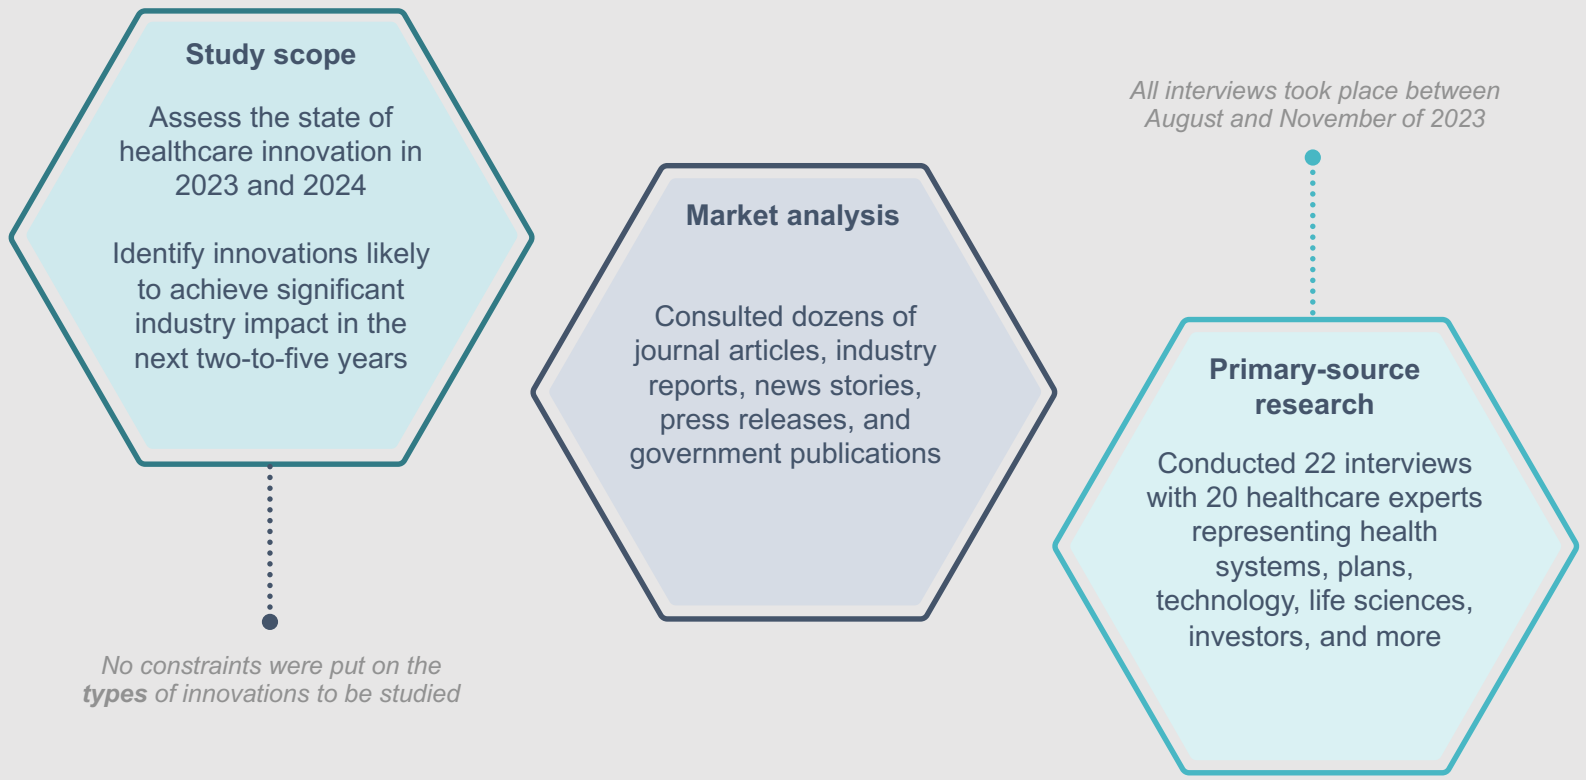

To frame the AI portion of the work, we began with the truism that AI is *both* over-hyped and under-leveraged. It's traveling along the oft-cited innovation hype cycle, in which we tend to overestimate the short-term impact of an innovation, while underestimating the long-term. Our own primary research corroborated this sentiment: AI was the healthcare innovation vector most often picked as likely to make a meaningful impact in the next five years but was also most frequently cited as overhyped (fig. 3).

Clearly, the question is not *whether* AI will be adopted in the next 5 years, but *where*, *how*, and by *how much*. It is equally clear, then, that all executive generalists should have baseline AI knowledge to help guide decision-making about potential AI-based solutions. This report is therefore structured around six commonsense questions that any thoughtful board member or strategy leader might pose—and should have educated answers to (figure 4, page 6).

Fig. 3: Results of Union Healthcare Insight's 2023 innovation survey
N = 20

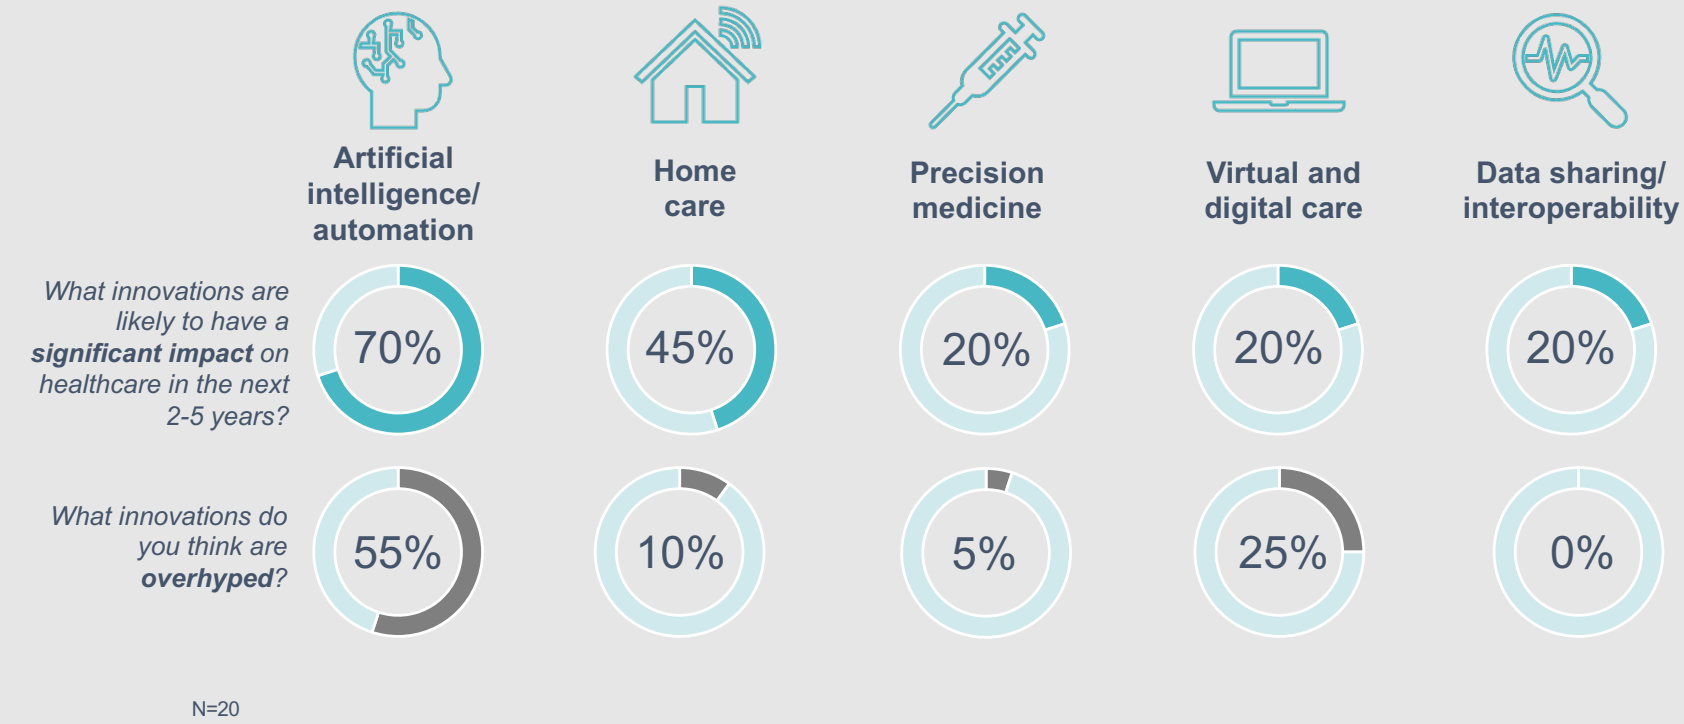

Fig. 4: Six commonsense questions informing the five-year outlook for healthcare AI

01 *"Generative AI is dominating the headlines—is that where we should focus our efforts?"*

02 *"I feel (and hear) lingering clinician/patient concerns; what does the latest data show?"*

03 *"In theory, AI could do many healthcare tasks; which types is it **generally** best suited for?"*

04 *"Generalities aside, where is **specific** uptake in healthcare **currently** highest—and why?"*

05 *"Given its massive 'someday' potential, how should we think about investing in AI today?"*

06 *"It's difficult to track all the news about AI regulations; what's the general picture?"*

Six insights on the 5-year outlook for AI in healthcare

Educated answers to commonsense questions

Six insights on the 5-year outlook for AI in healthcare

1: Surging interest in generative AI will not prompt immediate widespread adoption, but it will increase organizational AI IQ and momentum for more mature applications, such as predictive AI.

2: Clinicians and patients have developed more positive—and more nuanced—views on AI; the probable rate-limiter is no longer 'whether to use it at all' , so much as 'in what way'.

3: In the near term, adoption will follow an inverted bell curve; i.e., highest for low-skilled and super-human tasks, and lowest for highly skilled tasks.

4: Radiology has been the front-runner in healthcare AI adoption, revealing some larger truths about the types of healthcare AI applications with biggest future potential.

5: While it may be tempting to treat AI as an exception to the rule, AI's potential efficiency gains and ROI should be evaluated through the same lens as any other innovation.

6: Regulatory bodies are still in the early stages of assessing and approving AI-powered tools—and will be a wild card in accelerating or slowing AI adoption in the coming years.

"Generative AI is dominating the headlines—is that where we should focus our efforts?"

Insight #1

Surging interest in generative AI will not prompt immediate widespread adoption, but it will increase organizational AI IQ and momentum for more mature applications, such as predictive AI.

While the first generative AI application was developed in the 1960s, it wasn't until the recent launch of ChatGPT in November 2022 that generative AI made its way into the mainstream—and touched off an explosion of interest in AI generally. ChatGPT enabled individuals to interact and experiment with AI in a practical and accessible way. It has also demonstrated generative AI's ability to produce content with some level of human-like nuance, as opposed to the clunkier, more limited-use chatbots that had preceded it.

The spectrum of potential healthcare applications for generative AI is wide, particularly in cases where the inputs can be tightly controlled and trusted.

One particularly clear-cut example that has generated little pushback: it could be tremendously useful in providing language translation services.

Other possible applications are less clear-cut, and also more fraught; for example, early research indicates that patients may find AI-generated responses to be more empathetic that the language of human clinicians. This suggests that there may be interesting applications for tailoring patient communications but has also generated significant concern among physicians.

In fact, despite it's potential, there are a wide range of concerns and open questions about the use of generative AI in healthcare. These range from challenges around ensuring the accuracy of its generated content, to the general wish to preserve the 'human touch' in healthcare, to the desire to protect healthcare professions and employment.

As these ideas, concerns, and questions have played out in both industry and mainstream news, professional interest in AI has skyrocketed. At healthcare organizations, this interest has manifested in several concrete ways:

1. **Mandates from the board**: Boards are proactively approaching their organization's executive teams and IT departments with questions about AI and requests to understand their organization's approach to AI.
2. **Dedicated budget for AI experimentation**: Many organizations have dedicated a (sometimes significant) portion of their innovation budget toward AI specifically.
3. **Increased willingness among frontline workers to try solutions out**: Widening familiarity with generative AI has made previously reticent portions of the healthcare workforce—particularly clinical workers—more open to the idea of at least piloting AI applications (for more on this, see insight #2).

Even if generative AI faces significant hurdles before widespread adoption, these larger organizational shifts are a notable development. ChatGPT has increased awareness of AI as a whole, opening the door to a more mature set of predictive AI technologies that we think will see a meaningful uptick in adoption in the coming years (fig. 5).

In some cases, generative AI can even enhance the power of predictive AI; for example, by more easily enabling the creation of synthetic data sets (which can be an efficient means for enhancing the statistical significance of predictive models) or by helping to structure data for input into predictive models. This work is currently manual.

Fig. 5: Three near-term impacts of the rise of generative AI

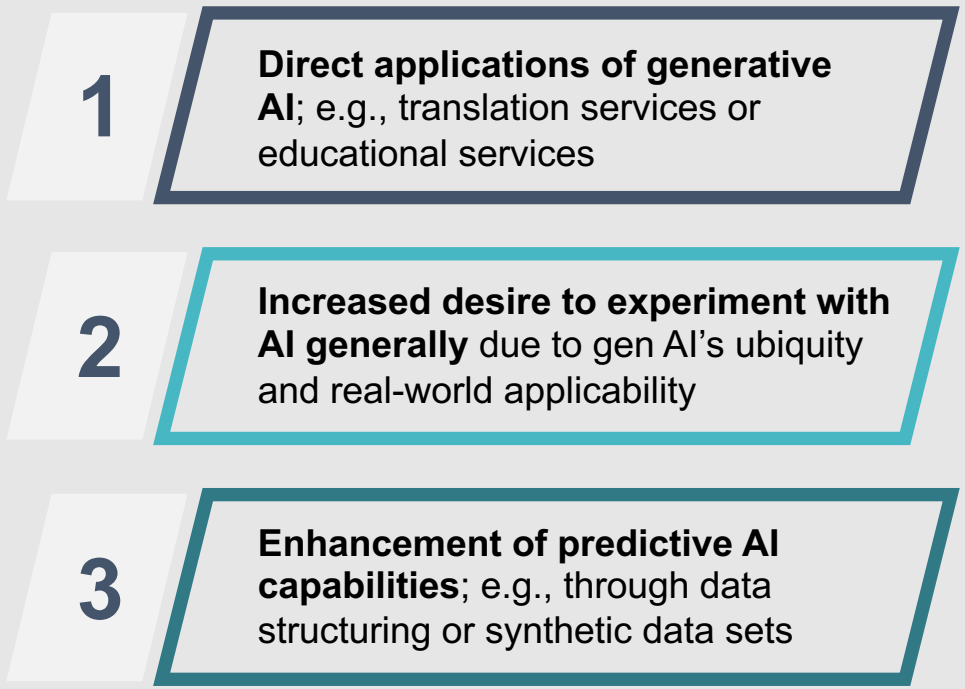

Fig. 6: Term check—generative AI vs. predictive AI

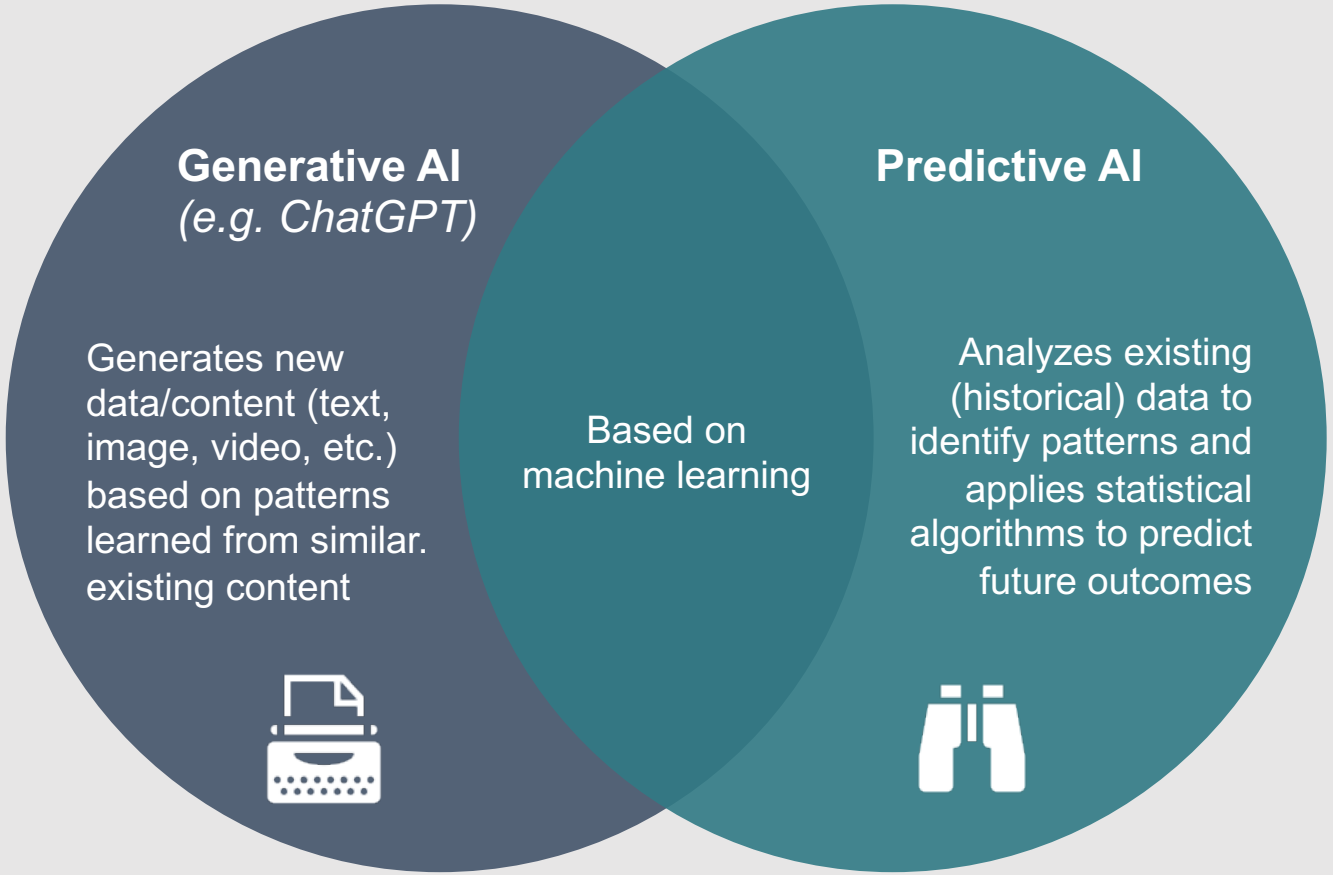

An AI jargon cheat sheet starter pack

Machine Learning (ML)	A type of artificial intelligence that uses algorithms to learn from datasets to make decisions/reach conclusions without human intervention.
Continuous learning	An approach to machine learning in which models can incorporate new information/data without explicit retraining.
Deep Learning	A subset of machine learning that mimics the human brain through the use of multiple layers of inputs/outputs (neural networks).
Large Language Models (LLM)	A subset of machine learning/deep learning that analyzes language inputs (Natural Language Processing/NLP) from massive datasets to generate text.

IBM Watson	The first widely known, AI-powered platform which gained prominence on Jeopardy in 2011; its foray into industries such as healthcare has been slower than expected.
OpenAI / ChatGPT	OpenAI is an organization founded with the goal of advancing AI development through collaboration and public patents; it made the ChatGPT chatbot publicly available in 2022.
Google Bard	Ai chatbot developed by Google and released in response to the launch of ChatGPT in early 2023.
Microsoft Bing	Longstanding Microsoft search engine which was relaunched in early 2023 with an integrated AI chatbot.

"I feel (and hear) lingering clinician/patient concern; what does the latest data show?"

Insight #2

Clinicians and patients have developed more positive—and more nuanced—views on AI; the probable rate-limiter is no longer 'whether to use it at all' , so much as 'in what way'.

A bedrock piece of industry conventional wisdom about AI has long been that the overall, adoption rate-limiter will be clinician or patient acceptance, or lack thereof.

This picture has changed meaningfully in a short time, thanks not only to changing views, but also more and better perception data. The arrival of ChatGPT has spurred more AI opinion research in healthcare; these studies have illuminated important nuances within clinician and patient views. These studies show that many clinicians and patients are still concerned about the use of AI in healthcare. However, overall acceptance about AI generally has grown, especially among clinicians.

In 2019, Medscape reported that nearly half of physicians in the U.S. were uncomfortable with the general idea of AI-powered tools. In contrast, a 2023 Medscape survey of over 1,000 physicians found that only 28% of physicians categorized themselves as "apprehensive". In fact, the plurality of physicians (42%) characterized themselves as "enthusiastic", with the remainder (30%) describing themselves as neutral (fig. 7, page 13).

>>>>>>

As of 2023, only 28% of physicians characterize themselves as generally "apprehensive" toward AI.

A critical nuance in current AI perception data is that surveys now test views on specific applications of AI, not just general sentiment about the technology per se.

Currently, while most physicians remain concerned about the potential use of AI to drive diagnoses independently, the majority are enthusiastic about its potential use as an adjunct to diagnosis.

And, in a stark contrast to the general public (more on this in a moment), many physicians believe that AI will enable them to spend *more* time with their patients, not less.

This is likely a function of the types of tasks that physicians report being most open to using AI for, with greater enthusiasm for administrative tasks and assistance with diagnosis vs. actual patient communication or treatment (figures 8 and 9, page 15).

Fig. 7: Physician attitudes toward AI in 2023
Medscape Physicians and AI Report 2023

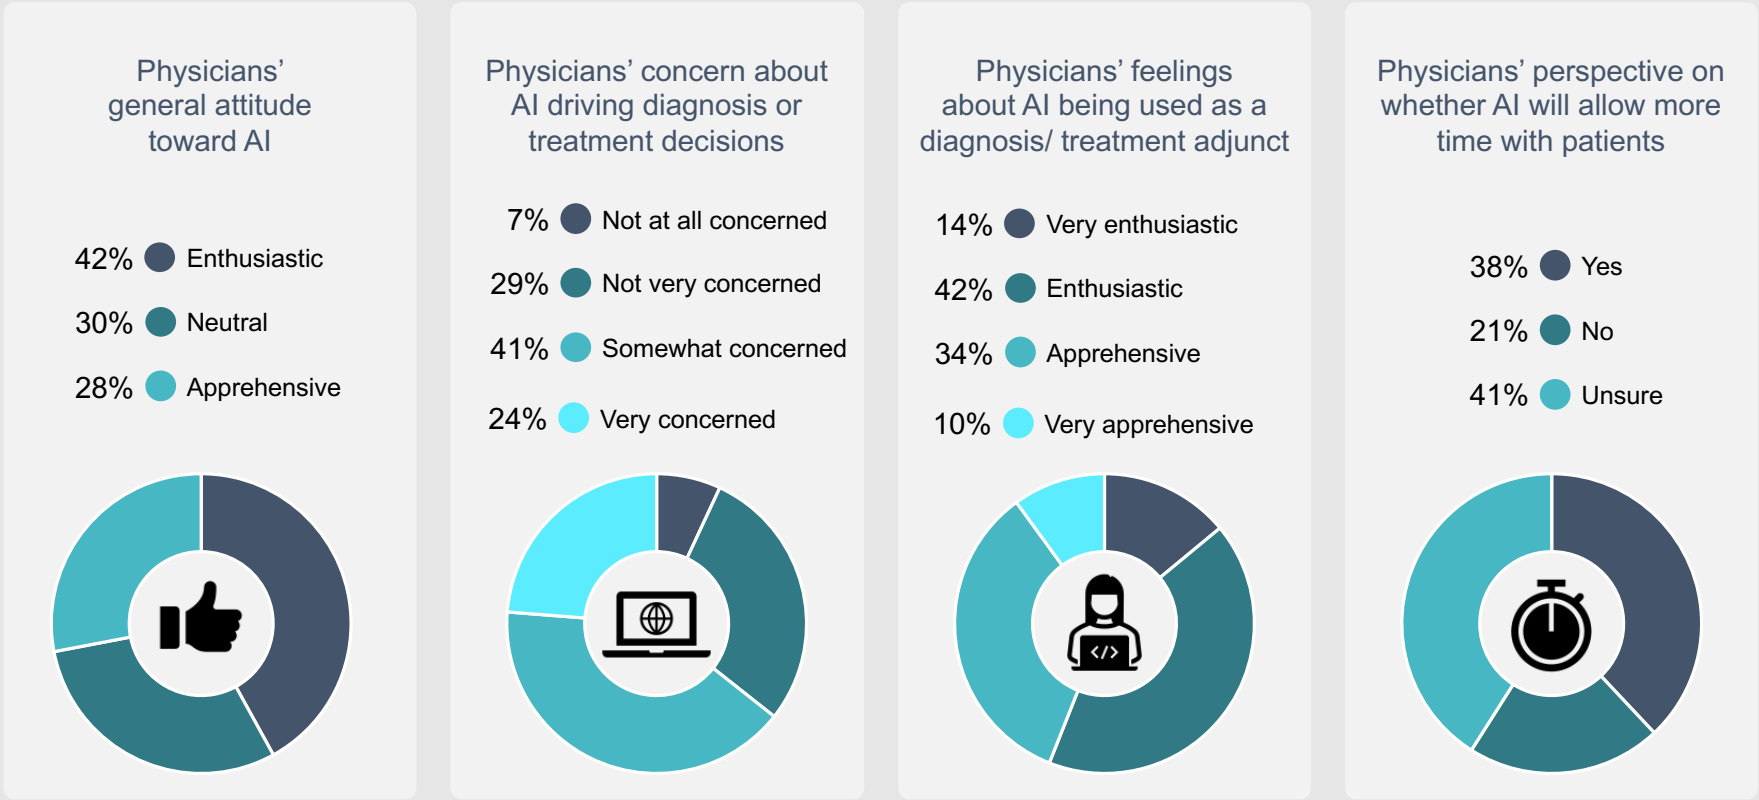

Fig. 8: How medical practices would consider using AI
Medscape Physicians and AI Report 2023

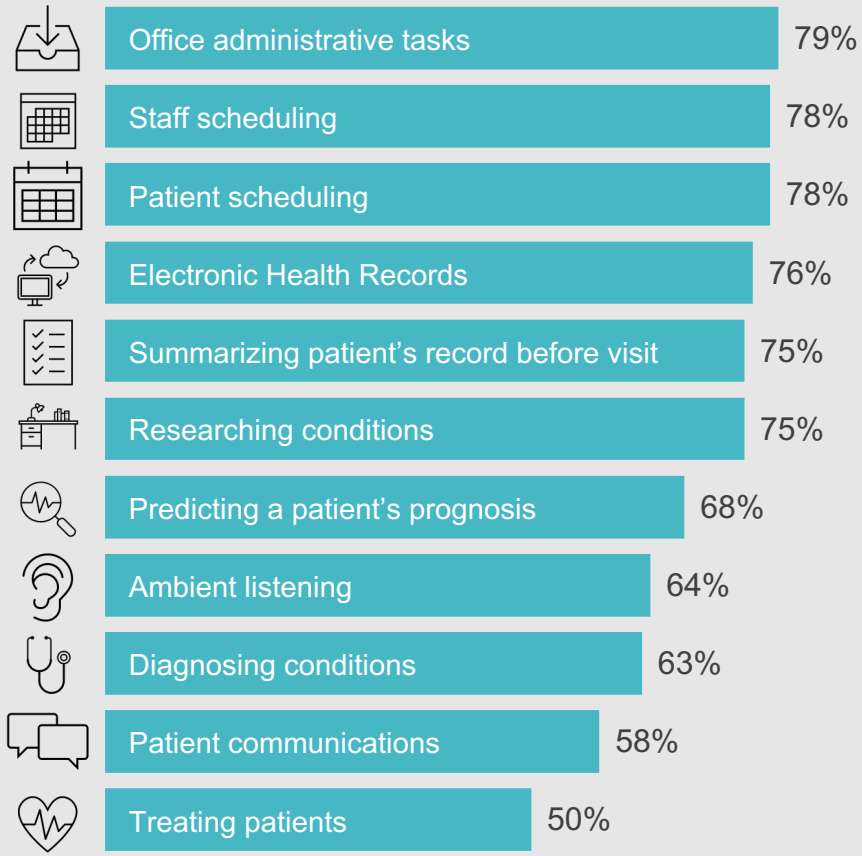

Fig. 9: How physicians feel about AI's use for specific activities
Medscape Physicians and AI Report 2023

Activity	Enthusiastic	Neutral	Negative
Administration and record keeping	67%	29%	5%
Scan interpretations	50%	35%	16%
Diagnosis	42%	37%	21%
Patient communications	37%	38%	25%
Malpractice risk management	36%	43%	21%

It's important to note that fresh data on the patient side is less positive than on the physician side; about 60% of patients report feeling uncomfortable with the idea of AI being used in their care (fig. 10).

However, public opinion data also shows that higher levels of AI familiarity are correlated with greater comfort with its use in healthcare. In a Pew survey of over 11,000 Americans conducted in December of 2022 (i.e. immediately following ChatGPT's launch) respondents who reported knowing "a lot" about AI were evenly split on being comfortable vs. not comfortable with its use in healthcare. On the other hand, among those who knew nothing about AI, a full 70% reported being uncomfortable with its use.

These results suggest that as familiarity with continues to AI grow, acceptance will likely grow too.

More granular data also now show the public recognizing trade-offs between different AI benefits and risks. The largest concerns are information security and potential compromise to personal relationships with providers (a concern that may eventually be allayed given the types of applications physicians are likely to embrace, as noted previously).

On the flip side, respondents are more optimistic about AI's ability to reduce medical errors; this is a notable departure from research published just a few years ago, in which increased medical errors were a top concern. The public also believes in AI's potential to help reduce racial disparities.

Fig. 10: 2023 Pew survey shows mixed public opinion on AI in healthcare

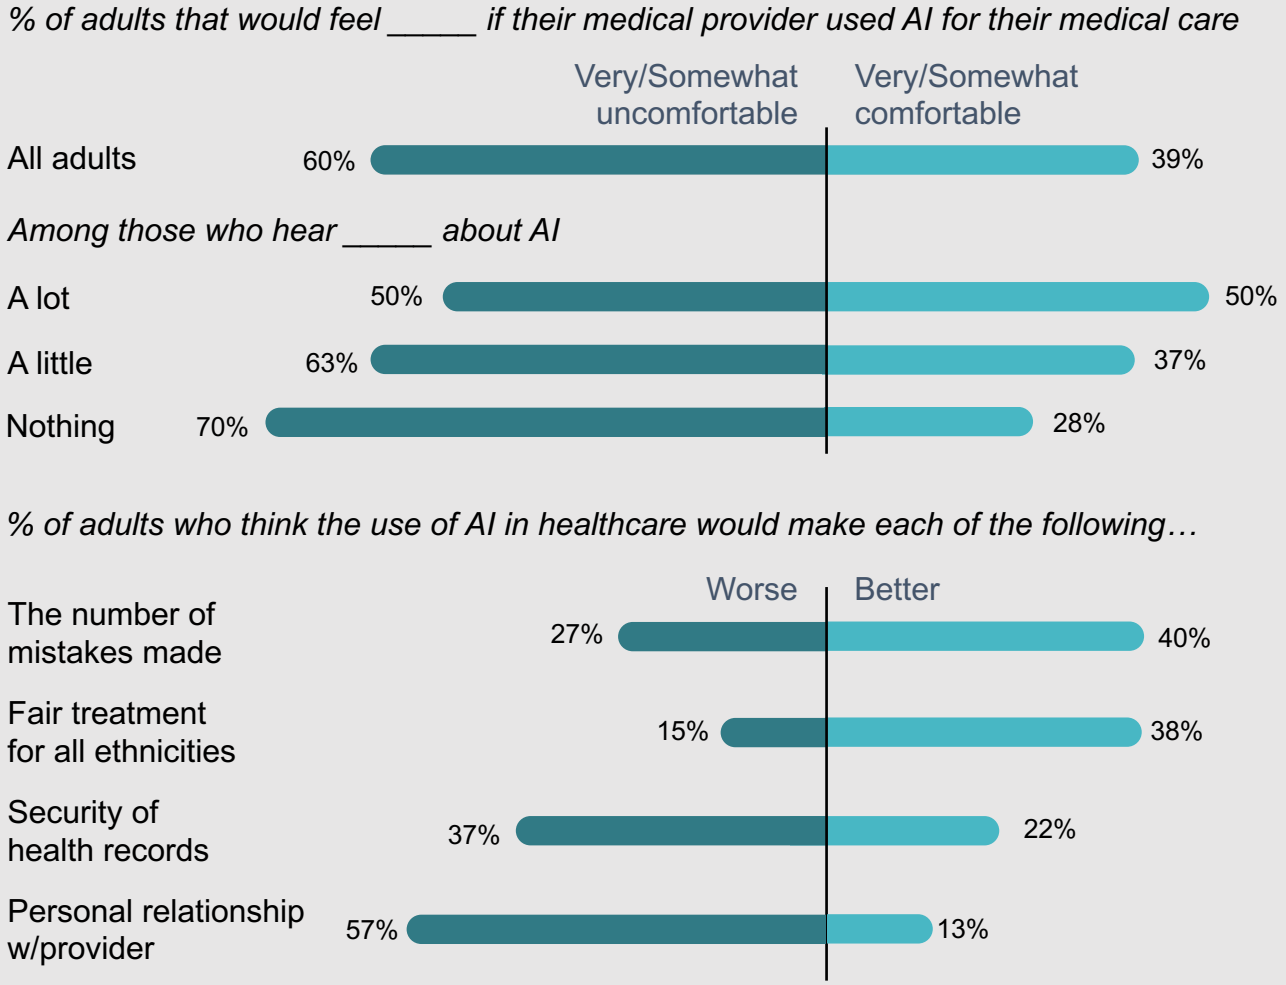

*In theory AI could do many healthcare tasks; which types are **generally** best suited?*

Insight #3

In the near term, adoption will follow an inverted bell curve; i.e., highest for low-skilled and super-human tasks, and lowest for highly skilled tasks.

Fig. 11: Near-term outlook for AI adoption, by task skill level

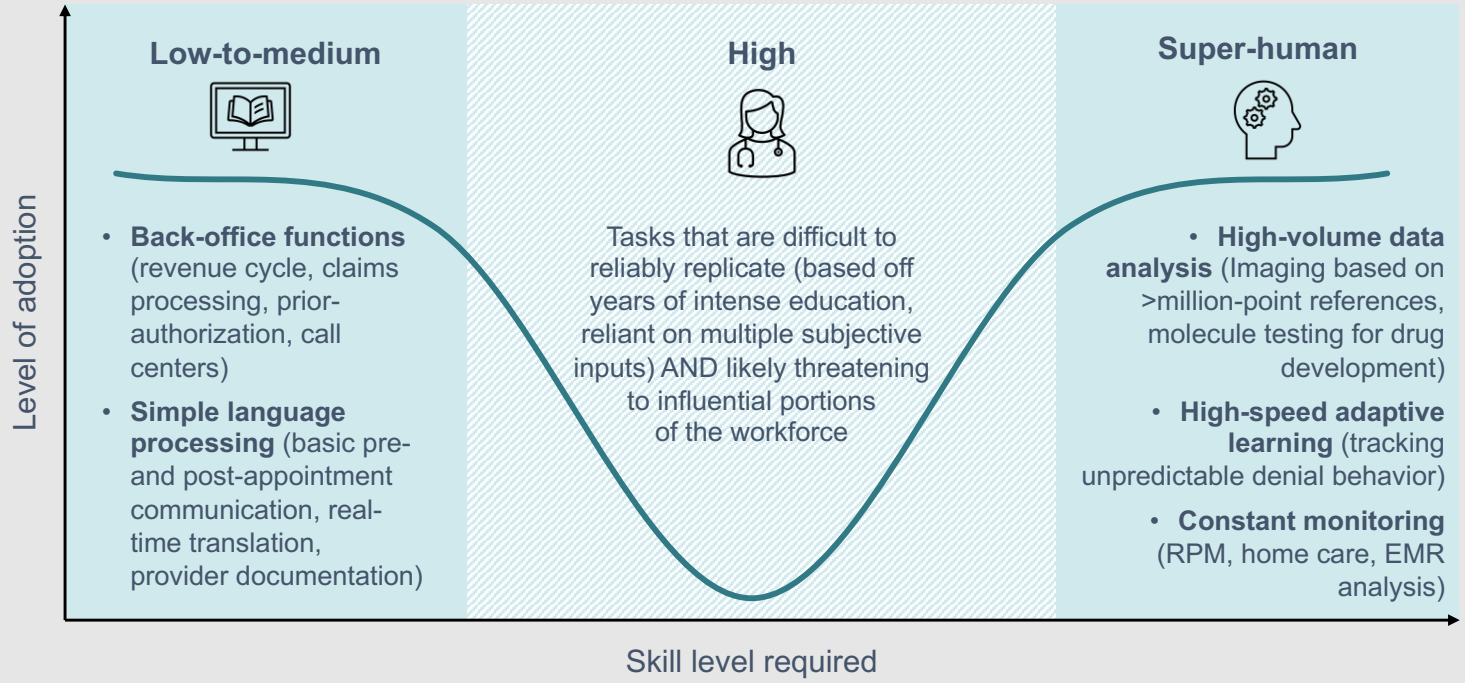

Across both generative and predictive AI applications in healthcare, the most significant three-to-five-year uptake potential lies at two different extremes of the task complexity spectrum. Low-to-moderate complexity tasks on the one hand, and super-human on the other, should see highest uptake (fig. 11, page 17).

That healthcare tasks at the lowest end of the complexity spectrum will see uptake is intuitive enough. These tasks are administrative, or clinical support related (such as patient communication or physician documentation). Adoption here will be speeded by the labor challenges facing the industry today; a silver lining in these challenges is that they hopefully obviate the need for the mass layoffs that many originally feared might be prompted by AI.

Meanwhile, uptake will be slowest in healthcare's highly skilled tasks, such as complex clinical decision-making. This dynamic is rooted in two primary factors.

First, it is more difficult to train AI to make complex decisions as skilled humans do today. Decisions such as those may be based on years of medical training *and* guided by a complex set of data inputs, including relatively "subjective" data such as patient-reported symptoms or lifestyle factors.

Second, highly skilled tasks are more likely to be the domain of highly influential portions of the workforce, such as physicians or unionized nurses. These professions have the power and the motivation to fight any AI adoption that they perceive as threatening their job security and/or the quality and safety of patient care.

At the far end of the complexity spectrum, however, we believe adoption will rise again; super-human applications of AI are ripe for adoption. Why? Because they are non-threatening to clinicians—and offer meaningful potential advances for patient care. Clinicians will likely be excited to pilot these applications.

Note that an AI application may be superhuman along two main dimensions (often simultaneously):

1) **Volume**. AI has the potential to sift through massive amounts of data—far more than a human being could ever be expected to. This could have interesting implications for diagnostics and drug development, in which there is potential for analysis of hundreds or thousands of microscopic data points.

2) **Speed**. AI can also complete many tasks more quickly and efficiently than a human being could ever hope to do. This could have interesting implications for all sectors of healthcare, both clinical and non-clinical. For example, AI could be trained to quickly learn and anticipate shifts in patterns around claims denials.

>>>>>>

AI-based technologies will have highest uptake at the two extremes of the task complexity spectrum: low-to-moderate complexity, and super-human. It will be slowest in the middle, for higher-complexity tasks.

This pattern is likely to be reassuring to both physicians and patients, as discussed in insight no. 2.

*"Generalities aside, where is **specific** uptake in healthcare **currently** highest, and why?"*

Insight #4

Radiology has been the front-runner in healthcare AI adoption, revealing some larger truths about the types of healthcare AI applications with biggest future potential.

In many ways, the biggest healthcare AI adoption story to-date is radiology. The FDA has approved more AI-enabled medical devices for radiology than for any other medical specialty, and by a wide margin: nearly 80% of AI-enabled devices approved between January and July of 2023 were for radiology, followed by cardiology as a distant second.

Many of these applications do, intuitively enough, focus on increasing radiologist efficiency by automating image analysis and reporting. But AI has the potential to enhance all aspects of the radiology care continuum (figure 12, page 20).

On the front end, AI-enabled solutions can help improve the image acquisition process; for example, by walking lower-skilled workers through the process of taking a scan, even using handheld devices, or by enhancing image quality.

AI can also lighten radiologist administrative tasks. For example, it can reduce the manual work to standardize and structure data for reporting and compliance purposes.

Finally, AI-enabled radiology devices can read scans at a level of detail and nuance that is undetectable to the human eye, thereby enhancing diagnostic potential and enabling earlier detection of disease.

>>>>>>

Nearly

80%

of AI devices approved by the FDA (between January and July of 2023 were for radiology, followed by cardiology as a distant second

Fig.12: Multiple opportunities for AI to improve efficiency and quality of medical imaging

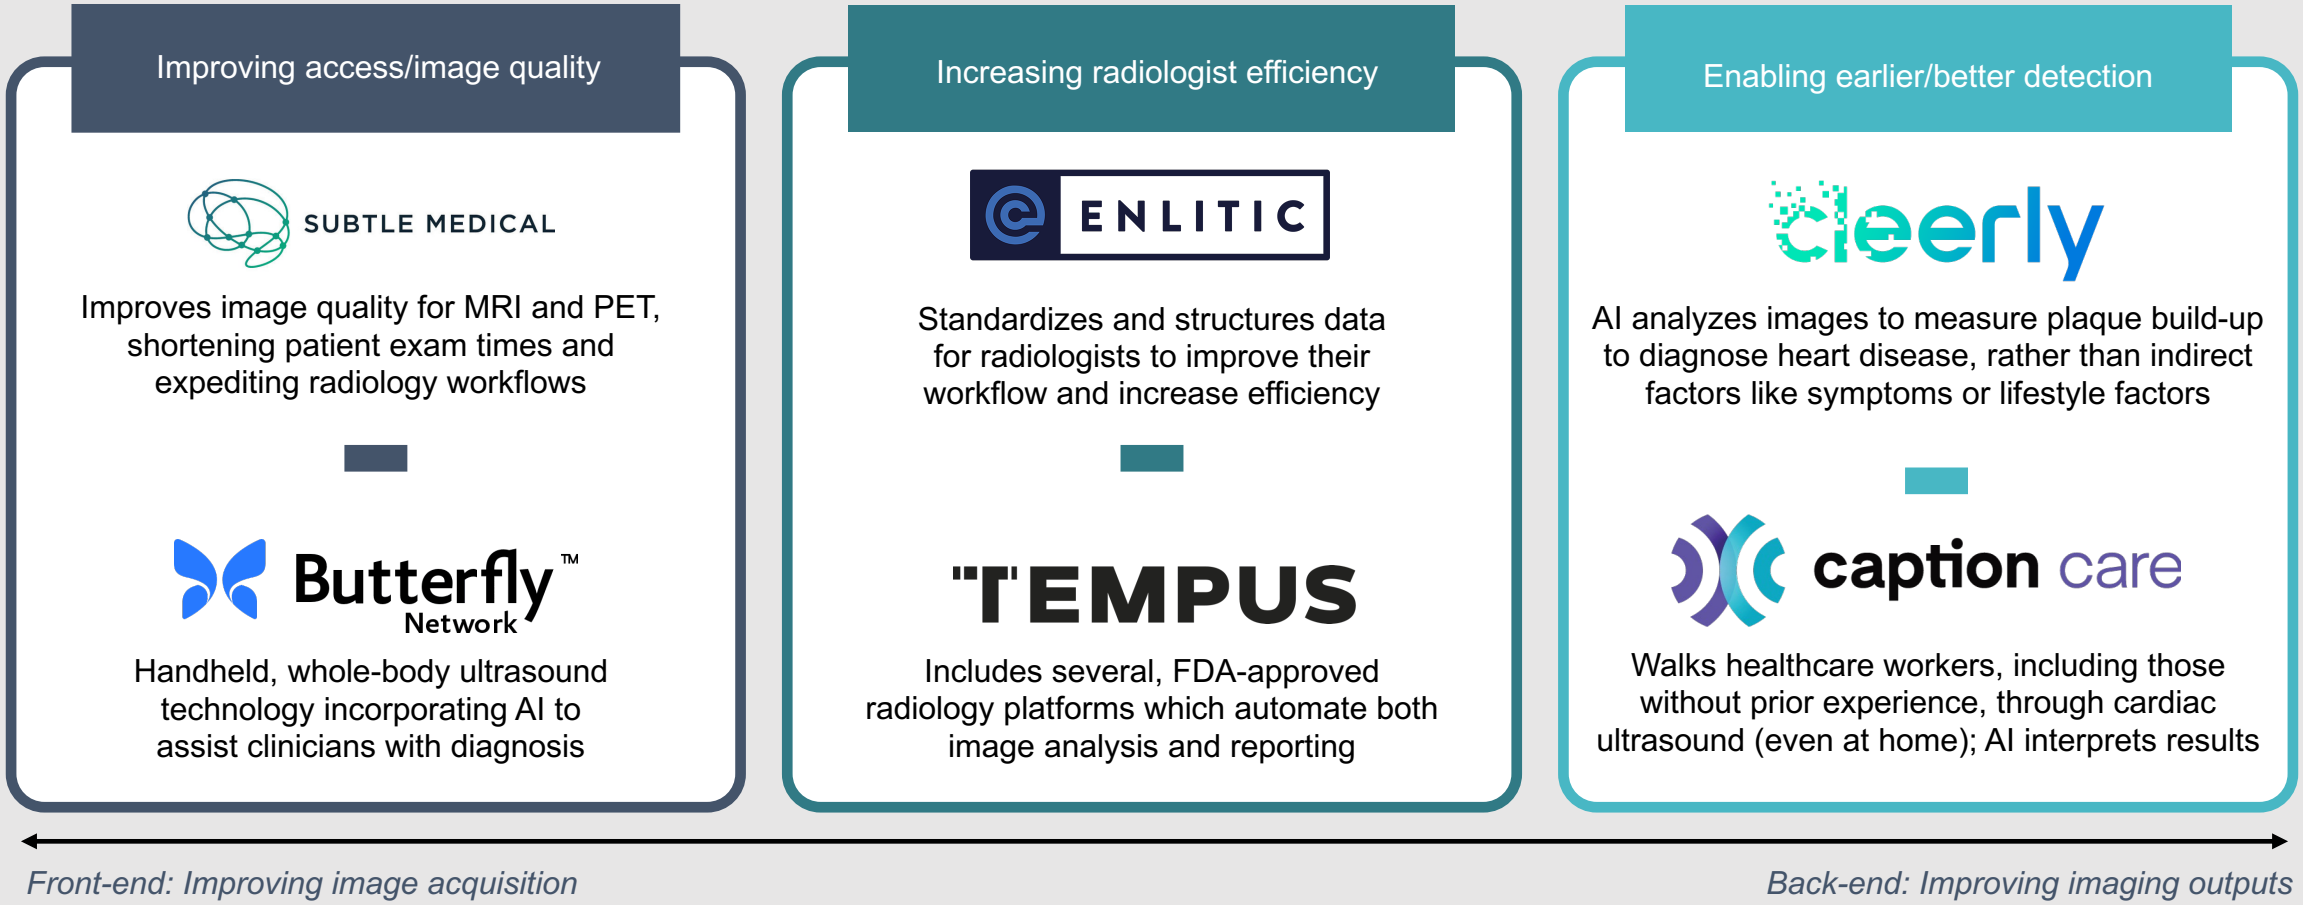

A few key factors have pushed imaging to the forefront. These include the ease with which the data inputs are collected, the relative objectivity of the data, the existence of clear care protocols guiding next steps, and the prevalence of both low-skilled and super-human tasks within the radiology workflow (i.e., targets that are less threatening to the radiology workforce). Recognizing these as 'success factors' should help guide both developers and enterprise buyers in assessing the near-term feasibility of AI applications outside of radiology (figure 13).

It's worth noting that once adoption in an area begins, it is likely to have a self-reinforcing/cyclical effect. This has happened in radiology as researchers and technology companies have built off existing approaches. Precedents also build confidence applications will have a good chance of securing adoption (and, where necessary, regulatory approval).

Fig. 13: Focus on imaging points to larger principles about AI feasibility

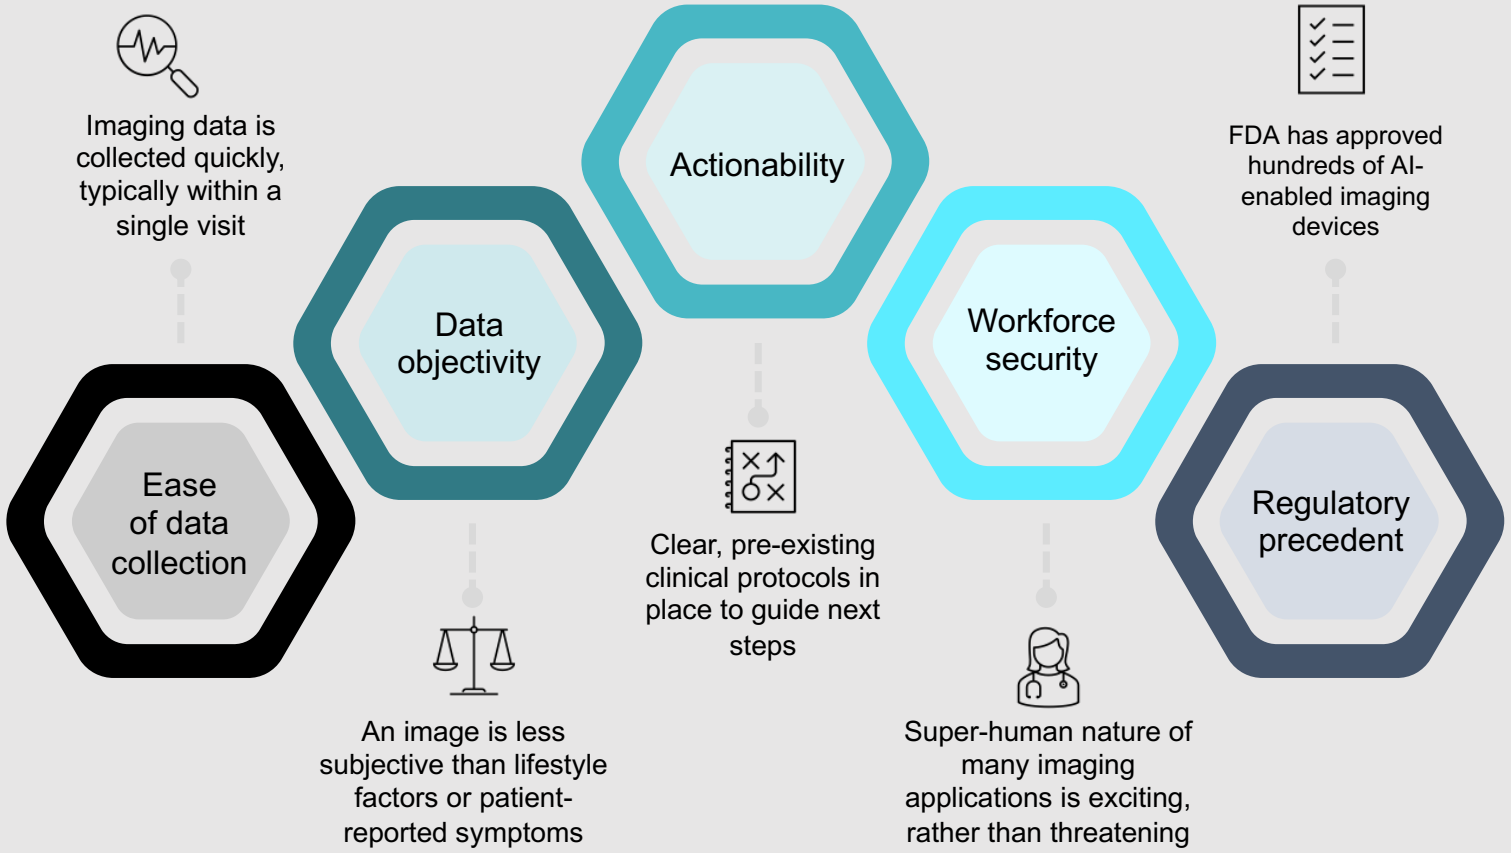

"Given its massive 'someday' potential, how should we think about investing in AI today?"

Insight #5

While it may be tempting to treat AI as an exception to the rule, AI's potential efficiency gains and ROI should be evaluated through the same lens as any other innovation.

To be adopted, most innovations must offer a relatively clear track to ROI for their enterprise buyers. But is AI, with its hype and much-vaunted potential, an exception to this rule?

Some organizations have reallocated large portions of their innovation budgets toward the express purpose of *experimenting* with AI—indicating that a strict path to ROI may not be needed for all AI-related solutions to at least get a foot in the door.

However, while this might be the right answer for a select few, for most healthcare organizations—for example, providers with thin operating margins—AI is best viewed as one of many potential tools for addressing organizational challenges and priorities. As such, it should be held to the same ROI expectations as any other potential solution, especially because there are "higher likelihood of ROI" AI solutions to experiment with, if experimentation is the mandate.

As with any other technology in healthcare, speed and size of ROI is determined not only by the nature of the application itself (e.g., tools that deliver savings within a single clinical episode vs. care management tools which deliver savings years down the line), but also by how quickly and easily the application can be adopted by end-users.

ROI will be achieved more quickly when tools slot easily into existing administrative and clinical workflows (figure 14, page 23)—and in cases where the workforce is explicitly incentivized to adopt tools that enhance productivity and efficiency.

This incentive may either be financial in nature (e.g.; functions that are rewarded based on volume of output) or more abstract (e.g.; the promise of superior survival rates due to the fast-track of life-saving care).

Fig. 14: AI applications with clear next steps will achieve quicker end-user adoption—and therefore faster ROI

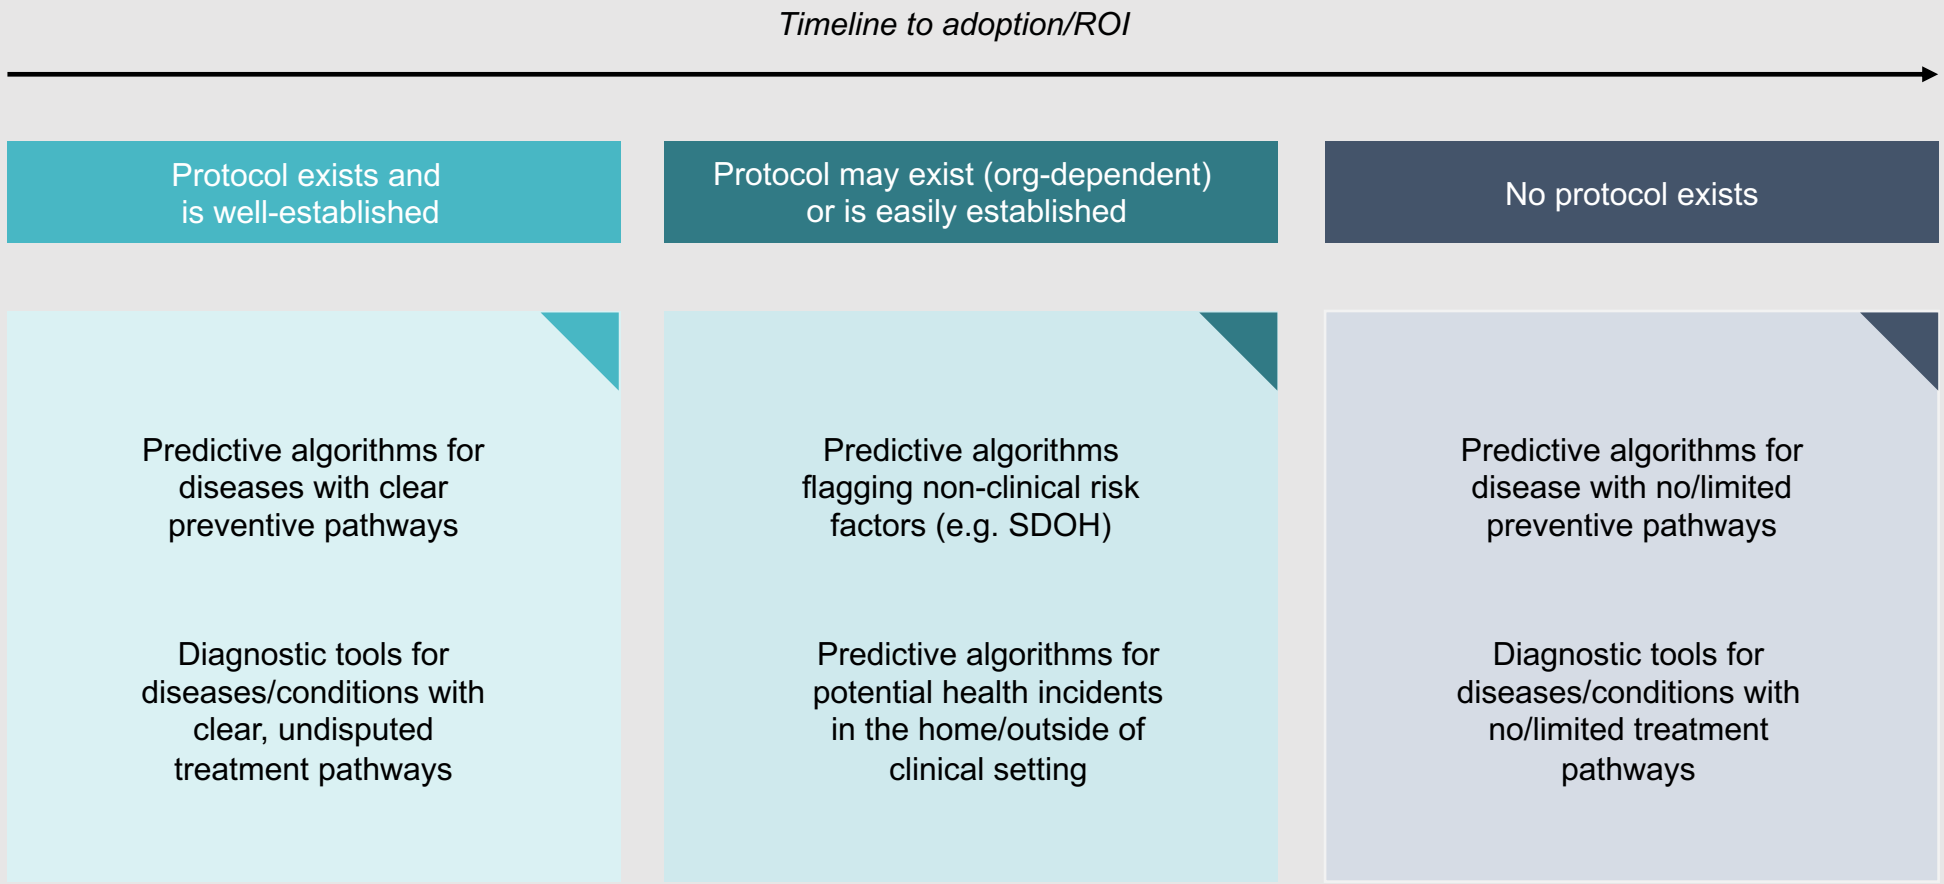

"It's difficult to track all the news about AI regulation; what's the general picture?"

Insight #6

Regulatory bodies are still in the early stages of assessing and approving AI-powered tools—and will be a wild card in accelerating or slowing AI adoption in the coming years.

The U.S. is taking a relatively decentralized approach to AI governance. There are no sweeping laws or regulations guiding the use of AI generally. Instead, a patchwork of efforts is taking shape across various federal agencies. Agencies within HHS (CMS, ONC, FDA, etc.) are all at various stages of issuing guidance and/or developing AI strategies, with the FDA being the furthest along.

Given this, it is more important for healthcare generalists (i.e. non-tech experts) to understand the broad themes that have emerged to-date, rather than to dive deep into individual pieces of draft guidance which are highly likely to become obsolete. Thus far, we have observed three major regulatory "lanes" for AI applications in healthcare (fig. 15, page 25).

1) AI-enabled "Software as a Medical Device" (SaMD)

Prior to the 1970s, medical devices were largely hardware-based. However, the increasing use of software as components of hardware-based medical devices (such as MRIs and pacemakers), prompted the FDA to expand its purview to software as well. The explosion in digital health tools in recent years has further expanded the FDA's oversight to software that is not embedded in a hardware device (i.e. Software as a Medical Device, or SaMD).

In 2019, the FDA released a proposed regulatory framework for AI-enabled SaMD, followed by an action plan in 2021. In the meantime (and as noted earlier), the FDA has approved nearly 700 AI-enabled medical devices on an ad hoc basis.

2) Use of AI for drug development

The FDA has also released some detail (and a request for feedback) on a potential approach to regulating the use of AI in drug development, acknowledging that there are potential use cases for AI at every step in the drug development process, from discovery to postmarket surveillance.

3) Everything else

There are *many* additional use cases for AI in healthcare outside of the preceding two categories, including administrative tasks (e.g. claims processing and management, back office/corporate functions, and non-clinical patient/member communication) as well as the *many* clinically-related tasks that fall short of the FDA's definition of a medical device (e.g. health tracking, many types of clinical decision support, patient-provider communication, etc.).

Fig. 15: Three emerging regulatory "lanes" for healthcare AI

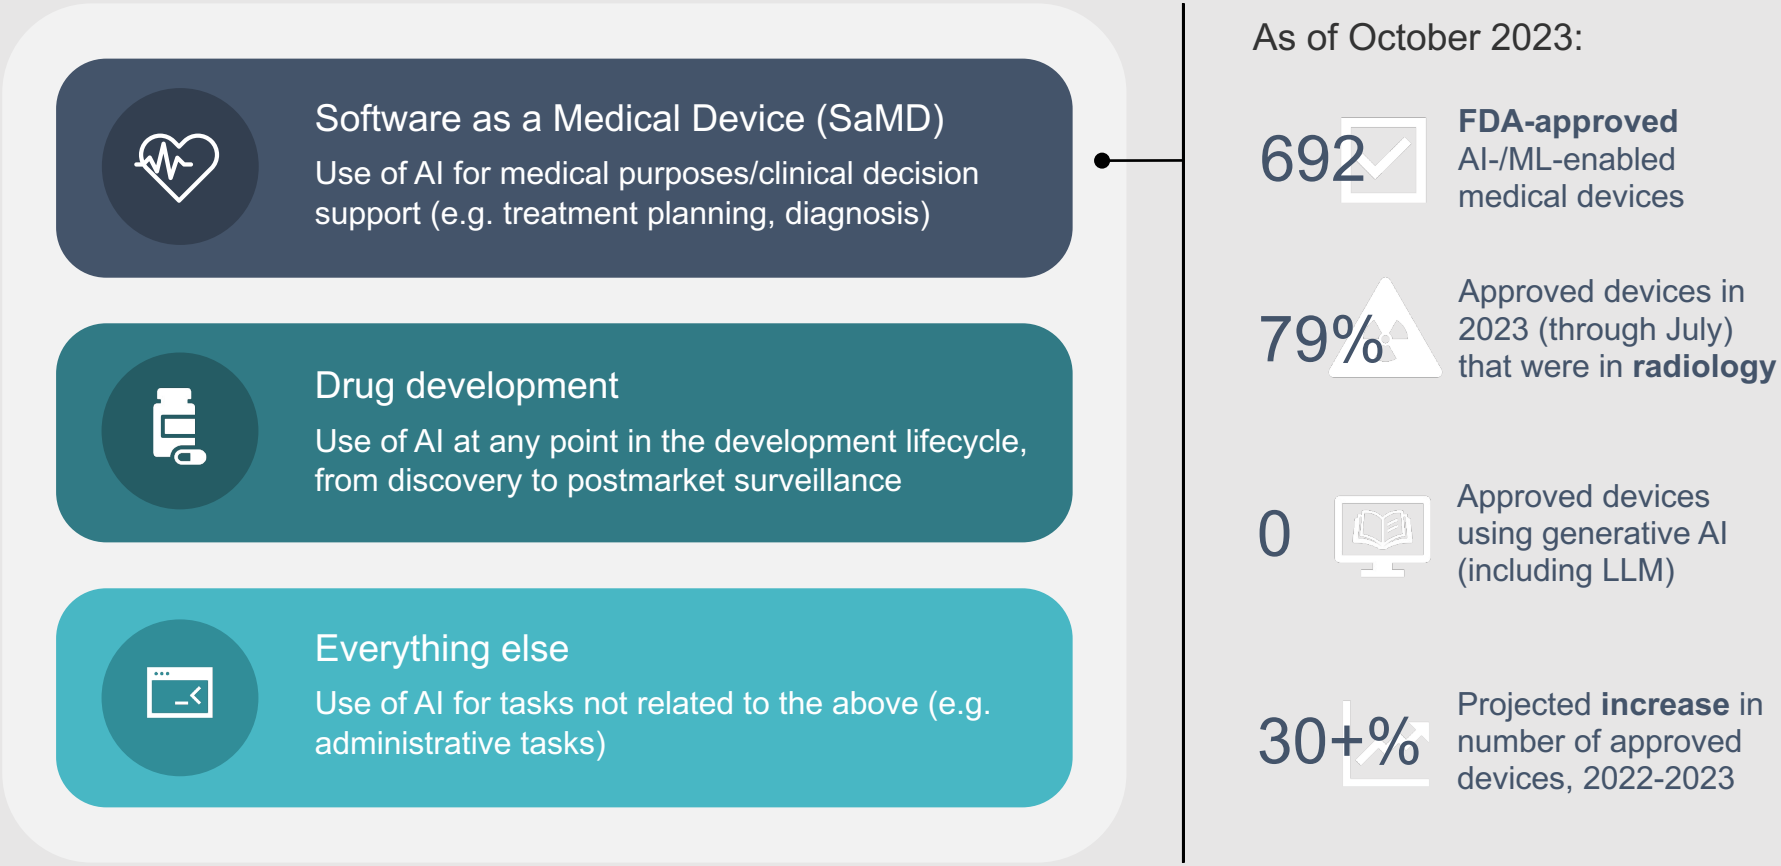

This does not mean that AI applications outside of the SaMD and drug development spaces are not subject to regulation; just that there are no *AI-specific* guidelines for these types of applications—at least as of this writing in January 2024. It is also likely that additional "lanes" will continue to emerge—for example, the ONC just finalized a rule in late 2023 which includes the first specific regulations about the use of AI in health IT systems such as Electronic Health Records.

While distinct policy frameworks are beginning to take shape for at least some specific type of applications, there are a clear set of common themes emerging from across these various regulatory bodies and governmental agencies. The majority of the policy focus tends to center on five areas (fig. 16, page 27).

Safety

First and foremost, policymakers are looking to establish processes and systems to ensure that AI-based healthcare applications don't jeopardize patient safety.

Quality/Efficacy

Beyond ensuring a baseline of basic safety, regulators and lawmakers are looking to establish frameworks for evaluating the quality/efficacy of AI in healthcare, both in an absolute sense, but also in comparison to non-tech-based approaches.

Fairness

A major concern with AI-based platforms is ensuring that they are fair, equitable, and have built-in protections against bias. Because AI-based applications are often trained on human-generated outputs, there are concerns that the technology could learn to mimic or even exacerbate human biases.

Security

As with any sort of healthcare technology, AI-based platforms will face high expectations related to privacy and security, especially in cases where consumer data—and especially patient health information—is involved.

Transparency

Finally, policymakers are actively grappling with questions related to transparency. This includes issues related to user consent (e.g. ensuring that users of AI-based applications are aware of the use of AI) as well as transparency around the design of the technology itself, especially in cases where machine learning and continuous learning are involved (i.e., the application continues to evolve over time).

An additional regulatory layer to consider: Staffing ratios and signing authority

Because AI has the potential to reduce the vast amount of manual, labor-intensive work that is currently done in healthcare, regulators are also working through questions of mandating certain levels of human involvement in healthcare processes. These types of regulations are driven by quality concerns, as well as lobbying efforts on the part of influential groups of clinicians such as physicians and nurses.

AI applications will need to comply with any (evolving) regulations around minimum staffing ratios, or instances in which a human being, potentially with a particular type of licensure, is required by law to review, approve, or sign off on a certain process or request.

These types of regulations are already common not only in the provider space, but among plans who do various types of medical review related to claims processing, and the AI element will continue to be hammered out in this terrain over time.

Fig. 16: Focus areas of select/notable rules, guidance documents, or proposed regulatory frameworks

	Safety	Quality/Efficacy	Fairness	Security	Transparency
White House Blueprint for an AI Bill of Rights	●	●	●	●	●
Coalition for Health AI Blueprint for Trustworthy AI	●	●	●	●	●
ONC HTI-1 Proposed Rule					●
FDA's proposed framework for AI-based SaMD	●	●			●
FDA discussion paper on AI in drug development		●			●
Executive Order on the Safe, Secure, and Trustworthy Development and Use of AI	●	●	●	●	●

Six insights on the 5-year outlook for AI in healthcare

"Generative AI is dominating the headlines—is that where we should focus our efforts?"

1: Surging interest in generative AI will not prompt immediate widespread adoption, but it will increase organizational AI IQ and momentum for more mature applications, such as predictive AI.

"I feel (and hear) lingering clinician/ patient concerns; what does the latest data show?"

2: Clinicians and patients have developed more positive—and more nuanced—views on AI; the probable rate-limiter is no longer 'whether to use it at all' , so much as 'in what way'.

*"In theory, AI could do many healthcare tasks; which types is it **generally** best suited for?"*

3: In the near term, adoption will follow an inverted bell curve; i.e., highest for low-skilled and super-human tasks, and lowest for highly skilled tasks.

*"Generalities aside, where is **specific** uptake in healthcare **currently** highest—and why?"*

4: Radiology has been the front-runner in healthcare AI adoption, revealing some larger truths about the types of healthcare AI applications with biggest future potential.

"Given its massive 'someday' potential, how should we think about investing in AI today?"

5: While it may be tempting to treat AI as an exception to the rule, AI's potential efficiency gains and ROI should be evaluated through the same lens as any other innovation.

"It's difficult to track all the news about AI regulations; what's the general picture?"

6: Regulatory bodies are still in the early stages of assessing and approving AI-powered tools—and will be a wild card in accelerating or slowing AI adoption in the coming years.

Resource center: Regulatory guardrails on healthcare AI

As of January 2024

Governmental Body	Resource	Useful if you want/need…
State Legislatures	Summary of Artificial Intelligence 2023 Legislation	A tracker of recent state legislative efforts/to check on the status of your state
White House	Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence	A broad overview of the federal view on AI + to look ahead to potential next steps for specific healthcare agencies
	Blueprint for an AI Bill of Rights	A succinct, accessible overview of the federal view on AI (i.e. light on policy jargon)
HHS	Artificial Intelligence (AI) at HHS	A broad summary of HHS' approach
	Trustworthy AI (TAI) Playbook	A detailed summary of HHS' approach
	Department of Health and Human Services: Artificial Intelligence Use Cases Inventory	A sense of how healthcare governmental agencies themselves are using AI

Governmental Body	Resource	Useful if you want/need…
FDA	Artificial Intelligence and Machine Learning (AI/ML) Software as a Medical Device Action Plan	A detailed summary of the FDA's current approach to AI SaMD (published January 2021)
	Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices	A list of all FDA-approved AI-enabled "Software as a Medical Device" (i.e. clinical AI)
	Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD)	The historical context on FDA approach to AI-enabled SaMD (RFI that preceded the January 2021 action plan)
	Artificial Intelligence and Machine Learning (AI/ML) for Drug Development	A top-line summary of FDA's approach to the use of AI in drug discovery/development + links to more detailed resources
	AI/ML for Drug Development Discussion Paper	A detailed summary of FDA's current thinking on the use of AI for drug development (in the request for feedback stage)
ONC	Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing	Technical detail on the ONC's potential approach to AI in health IT (proposed rule)
CMS	Artificial Intelligence at CMS	A broad summary of CMS' approach + links to additional resources

AI Application Feasibility Calculator

Less feasible ◄——————————————————————► *More feasible*

	1	2	3	4	5	Score
Output quality	**How superior is the output of the AI to that a human being?**					
	On par *(similar output/speed)*		**Somewhat superior** *(similar output, faster speed)*		**Far superior** *(super-human speed/output)*	
	How quick is the likely ROI?					
	Long-term *(years)*		**Medium-term** *(within a year)*		**Immediate** *(within a care episode)*	
					Total	

Less feasible ◄——————————————————————► *More feasible*

	1	2	3	4	5	Score
Usability	**How difficult is it to collect the necessary inputs for this AI application?**					
	Very difficult *(multi-source, multi-input)*		**Moderately difficult** *(single-source, multi-input)*		**Not at all difficult** *(single input)*	
	How clear is the next step (e.g. care protocol, administrative task) based on the output of this AI application?					
	Not at all clear *(requires new protocol)*		**Somewhat clear** *(protocol needs modification)*		**Very clear** *(slots into existing protocol)*	
	How familiar are people likely to be with this application's technology/interface?					
	Not at all familiar *(limited use to-date)*		**Somewhat familiar** *(wide use in other industries)*		**Very familiar** *(wide use within healthcare)*	
					Total	

Less feasible ← → *More feasible*

	1	2	3	4	5	Score
Workforce impact	**How threatening is this AI application from a job security perspective?**					
	Very threatening *(highly-skilled task)*		**Moderately threatening** *(low-skilled task)*		**Not at all threatening** *(super-human task)*	
	How much organizational or political leverage do the impacted workers have?					
	Significant leverage *(E.g. unionized nurses, MDs)*		**Moderate leverage** *(highly-skilled workers)*		**Limited leverage** *(Low-skilled workers)*	
	To what extent are there labor shortages within the impacted workforce?					
	No shortage *(sufficient supply)*		**Moderate shortage** *(slow attrition/hard to backfill)*		**Severe shortage** *(capacity-limiting shortage)*	
	How desirable is the application likely to be to the average impacted worker?					
	Not at all desirable *(offloads desirable tasks)*		**Somewhat desirable** *(offloads undesirable tasks)*		**Highly desirable** *(fun to use/excitement factor)*	
					Total	

Less feasible ← → *More feasible*

	1	2	3	4	5	Score
Stakes	**How high-stakes is the task at hand?**					
	Very high *(life or death)*		**Somewhat high** *(basic clinical task)*		**Very low** *(basic administrative task)*	
					Total	

Less feasible ◄——————————————————————————————► *More feasible*

	1	2	3	4	5	Score
Regulatory outlook	**How highly-regulated is the task/decision at hand?**					
	Limited regulation *(purely administrative tasks)*		**Somewhat highly-regulated** *(clinical decisions/review)*		**Very highly-regulated** *(requires FDA approval)*	
	To what extent has the application been designed and tested to prevent bias and discrimination?					
	Not at all *(no testing/protections)*		**Somewhat** *(proven not to worsen bias)*		**Significant** *(proven to improve bias)*	
	To what extent has the application been designed and tested to ensure data security and privacy?					
	Not at all *(no testing/protections)*		**Somewhat** *(some testing/protections)*		**Significant** *(HIPAA-compliant/no PHI)*	
	How well-established is the process for ensuring transparency for users of this application?					
	Not at all *(no formal process)*		**Somewhat** *(one-time disclaimer)*		**Significant** *(clear, continuous protocols)*	
					Total	

Less feasible ◄——————————————————————————————► *More feasible*

	1	2	3	4	5	Score
Ease of adoption (for enterprise buyers only)	**To what extent is there existing buy-in from the impacted workforce?**					
	No buy-in *(no clear champion)*		**Moderate buy-in** *(champion, limited influence)*		**Significant buy-in** *(champion w/sig. influence)*	
	How involved was the impacted workforce in building out or testing this tool/application?					
	Not at all involved *(3rd party vendor)*		**Somewhat involved** *(pilot site for vendor)*		**Very involved** *(developed in-house)*	
	To what extent does the application align with broader organizational goals?					
	Not at all aligned *(no link to pre-existing goals)*		**Somewhat aligned** *(tied to department goals)*		**Significant alignment** *(tied to org's strategic plan)*	
					Total	

Conclusion

How to move thoughtfully on AI

Position AI as part of a toolset for existing problems

As noted earlier, it can be tempting, given the hype around AI, to treat it differently than other types of innovation. We believe, however, that AI is best viewed as one of a plethora of potential tools to address existing organizational challenges and opportunities.

As alluded to throughout this report, potential applications for AI are seemingly endless—but thoughtful exploration reveals that certain types of tasks are better-positioned for near-term adoption than others. Leaders should weigh feasibility of adoption against their existing, named organizational priorities. We think some potential areas of overlap may include: clinical trial efficiency, health equity (although, as noted, AI also carried the risk of exacerbating bias), and, of course, workforce shortages.

Target parts of the workforce where AI will be a value-add

While AI (along with more basic forms of automation) can be a powerful tool for addressing workforce challenges, leaders must be careful to target early AI applications to portions of the workforce where AI will be as a value-add, as opposed to a potential threat or an additional technological burden.

This may include departments with high levels of natural attrition (especially if roles are particularly difficult to backfill), as well as any roles that are explicitly compensated or rewarded on the basis of productivity or volume. It may also include portions of the workforce who are excited by the new capabilities that AI can unlock (for example, the radiology applications discussed earlier in this report).

Don't underestimate the value of a clinician champion

The importance of a clinician champion for any evolution to existing clinical protocol is well-established at this point. And its importance has not diminished over time.

This principle remains true for any experimentation with or adoption of AI-based technologies. This means that organizations should actively source opportunities for experimentation directly from physicians and nurses. Or, in instances where innovation or technology leaders have identified a potential solution to a known challenge within the workforce, leaders should work to establish a clear champion before attempting to push broad adoption. Any pilots lacking a clear champion should be approached with significant caution, fi not re-evaluated entirely.

Invest in AI education and governance, not an AI strategy

The most important step an organization can take is to invest in AI education and governance. Doing so will enable thoughtful evaluation of AI-based solutions at all levels of the organization (from executive teams and board to frontline workers), while preventing waste of resources by pursuing AI for the sake of AI.

We hope that this report will serve as a helpful resource for those with baseline knowledge who are actively looking for a more informed perspective. But we also know that there are many who would benefit from 101-level education on AI in order to better understand the technology, navigate the jargon, and learn about the major players. To that end, Union will have a strategy bootcamp module on AI available in Q1 of 2024.

Fig. 17: How to move thoughtfully on AI

- Clinical trial efficiency
- Health equity (also a potential risk!)
- Workforce shortages

- Departments with high attrition
- Roles that are difficult to backfill
- Roles rewarded on productivity

- Source opportunities from physicians/nurses
- Think twice about pilots without a clear champion

- Enable leaders to make smart decisions as needed
- Union strategy bootcamp module coming Q1 2024

Research team
Yulan Egan
Yulan@UnionHealthcareInsight.com

Kirsta Hackmeier
Kirsta@UnionHealthcareInsight.com