UNIVERSITY OF LJUBLJANA
SCHOOL OF ECONOMICS AND BUSINESS

MASTER'S THESIS

THE ROLE OF INTELLIGENT AGENTS IN OPTIMIZING HEALTHCARE BUSINESS PROCESSES: ADOPTION AND INTEGRATION CHALLENGES

Ljubljana, May 2026 WENTING LIU

AUTHORSHIP STATEMENT

The undersigned Wenting Liu, a student at the University of Ljubljana, School of Economics and Business, (hereafter: UL SEB), author of this written final work of studies with the title "The Role of Intelligent Agents in Optimizing Healthcare Business Processes: Adoption and Integration Challenges" prepared in collaboration with mentor doc.dr. Luka Tomat.

DECLARE

- this written final work of studies to be based on the results of my own research;
- the printed form of this written final work of studies to be identical to its electronic form;
- the text of this written final work of studies to be language-edited and technically in adherence with the UL SEB Technical Guidelines for Written Works, which means that I cited and / or quoted works and opinions of other authors in this written final work of studies in accordance with the UL SEB Technical Guidelines for Written Works;
- to be aware of the fact that plagiarism (in written or graphical form) is a criminal offence and can be prosecuted in accordance with the Criminal Code of the Republic of Slovenia;
- to be aware of the consequences a proven plagiarism charge based on the this written final work could have for my status at the UL SEB in accordance with the relevant UL SEB Rules;
- to have obtained all the necessary permits to use the data and works of other authors which are (in written or graphical form) referred to in this written final work of studies and to have clearly marked them;
- to have acted in accordance with ethical principles during the preparation of this written final work of studies and to have, where necessary, obtained permission of the Ethics Committee;
- my consent to use the electronic form of this written final work of studies for the detection of content similarity with other written works, using similarity detection software that is connected with the UL SEB Study Information System;
- to transfer to the University of Ljubljana free of charge, non-exclusively, geographically and time-wise unlimited the right of saving this written final work of studies in the electronic form, the right of its reproduction, as well as the right of making this written final work of studies available to the public on the World Wide Web via the Repository of the University of Ljubljana;
- my consent to publication of my personal data that are included in this written final work of studies and in this declaration, when this written final work of studies is published;
- that I have verified the authenticity of the information derived from the records using artificial intelligence tools.

Ljubljana, May 2026

Author's signature:



ABSTRACT

Intelligent Agents are emerging as transformative tools for optimizing healthcare business processes, from clinical workflows to administrative operations. This thesis investigates how IAs enhance efficiency, reduce costs, and improve outcomes through automation and data-driving decision-making while critically analyzing the adoption of barriers and systemic integration challenges. Drawing on global case studies, the research identifies technical, operational, and ethical bottlenecks that hinder IAs deployment. The findings propose actionable strategies for overcoming these challenges, emphasizing the need for interoperable frameworks, stakeholder collaboration, and adaptive regulatory policies to realize the full potential of IAs-enabled healthcare optimization.

KEY WORDS: Intelligent Agents, healthcare business process, digital healthcare transformation, integration challenges

SUSTAINABLE DEVELOPMENT GOALS



POVZETEK

Inteligentni agenti se pojavljajo kot transformativna orodja za optimizacijo poslovnih procesov v zdravstvu, od kliničnih delovnih procesov do administrativnih operacij. Ta disertacija raziskuje, kako inteligentni agenti povečujejo učinkovitost, zmanjšujejo stroške in izboljšujejo rezultate z avtomatizacijo in odločanjem na podlagi podatkov, hkrati pa kritično analizira ovire in izzive sistemske integracije. Raziskava, ki se opira na globalne študije primerov, opredeljuje tehnična, operativna in etična ozka grla, ki ovirajo uvajanje inteligentnih agentov. Ugotovitve predlagajo izvedljive strategije za premagovanje teh izzivov, s poudarkom na potrebi po interoperabilnih okvirih, sodelovanju deležnikov in prilagodljivih regulativnih politikah, da bi uresničili polni potencial optimizacije zdravstva, ki jo omogočajo inteligentni agenti.

KLJUČNE BESEDE: Inteligentni Agenti, poslovni procesi v zdravstvu, digitalna transformacija zdravstva, izzivi integracije

CILJI TRAJNOSTNEGA RAZVOJA



TABLE OF CONTENTS

1 INTRODUCTION 1

2 LITERATURE REVIEW 4

2.1 Definition of IAs and their characteristics 4

2.2 Relevant Theoretical Frameworks for Analyzing the Adoption and Integration of IAs in Healthcare Business Processes 6

2.2.1 RBV theory 6

2.2.2 DOI theory 8

2.3 IAs and Business Processes 9

2.3.1 Healthcare Business Processes 10

2.3.2 Functional typology of IAs within healthcare processes 11

2.3.3 Market Validation of IAs 13

2.4 Current applications of IAs in healthcare 14

2.5 Business and Economic impact of IAs in healthcare 16

2.5.1 B2B Market segment 20

2.5.2 B2C Market segment 22

2.6 Challenges in IAs adoption in healthcare 22

2.6.1 Technical challenges 23

2.6.2 Operational challenges 24

2.6.3 Legal and Regulatory Concerns 25

2.6.4 Ethical and Societal Concerns 26

3 ADOPTION AND INTEGRATION CHALLENGES OF INTELLIGENT AGENTS IN HEALTHCARE 27

3.1 Description of research approach and research design 27

3.1.1 Research methodology 27

3.1.2 Data collection 28

3.1.3 Data analysis 29

3.1.4 Ethical considerations and limitations 30

3.2 Interviews implementation 30

3.3 Analysis of the results 31

3.3.1 Baichuan Intelligence analysis 31
3.3.2 Alipay Health analysis 32
3.3.3 A tertiary public hospital analysis 32
4 DISCUSSION 33
4.1 Summary and interpretation of key findings 34
4.1.1 The relevant findings of Alipay's deployment of IAs 34
4.1.1.1 Creation: Owning the healthcare "Front Door" 34
4.1.1.2 Enhancement: Empowering Clinicians with a "Digital Twin" 36
4.1.1.3 Replacement: The Emergence of the Virtual Family Doctor 36
4.1.1.4 Challenges and strategic considerations 37
4.1.2 The relevant findings of Baichuan's deployment of IAs 38
4.1.2.1 Creation: A new intelligent interface for pediatric care 38
4.1.2.2 Enhancement: AI augmentation of healthcare workflows 39
4.1.2.3 Replacement: Delegating low-complexity cases to agents 40
4.1.2.4 Challenges and strategic considerations 41
4.1.3 The relevant findings of public hospital adoption of IAs 41
4.1.3.1 Creation: IAs as care coordinators and emotional support tools 42
4.1.3.2 Enhancement: Documentation automation and efficiency 42
4.1.3.3 Replacement: Selective automation of routine administrative tasks 43
4.1.3.4 Challenges and strategic considerations 43
4.2 Key findings and Limitations 44
4.3 Future work 45
5 CONCLUSION 46
LIST OF KEY REFERENCES 47
REFERENCES 48
APPENDICES 57

LIST OF TABLES

Table 1: The design and implementation principles of IAs 6
Table 2: Types of IAs applied with healthcare business processes 11

Table 3: Cost and Benefit model of agent adoption in healthcare organizations 17
Table 4: Business and Economic Impact by Category 18
Table 5: The RACI matrix of deployment of IAs in Healthcare environments 25
Table 6: Interview protocol development focus area 29

LIST OF FIGURES

Figure 1: Core business processes in healthcare 10
Figure 2: The AI Hype Cycle 13
Figure 3: The strategic roles of IAs 16
Figure 4: The interface of AI Health Butler 35

LIST OF APPENDICES

Appendix 1: Interview Transcript – Participant from Alipay (P1) 1
Appendix 2: Interview Transcript – Participant from Baichuan (P2) 10
Appendix 3: Interview Transcript – Participant from hospital (P3) 17

LIST OF ABBREVIATIONS

AI - Artificial Intelligence

API – Application Programming Interface

AR – Augmented Reality

B2B - Business to Business

B2C - Business to Customer

CDSS - Clinical Decision Support Systems

CEO – Chief Executive Officer

CFDA - China Food and Drug Administration

CIO – Chief Information Officer

DOI - Diffusion of Innovations

DRG - Diagnosis-Related Group

EHR - Electronic Health Record

EMR - Electronic Medical Record

HIS - Hospital Information System

HL7 – Health Level Seven International

IAs - Intelligent Agents

IT - Information Technology

IAW-MISS – Intelligent Agent Web-based Medical Information Support System

LLMs – Large Language Models

MASH – Multi-Agent Systems for Healthcare

RACI matrix – Responsible, Accountable, Consulted, Informed matrix

RBV - Resource-Based View

R&D – Research and Development

RNNs - Recurrent Neural Networks

ROI – Return On Investment

RPA - Robotic Process Automation

SaaS – Software as a Service

SMEs – Small-to-Medium-Sized Enterprises

SMPC – Security Multi-Party Computation

WHO - World Health Organization

1 INTRODUCTION

Worldwide, Healthcare systems face systemic inefficiencies that undermine operational efficacy and patient outcomes. According to the World Health Organization (WHO, 2023), the distribution of healthcare spending remains highly unequal among countries. For instance, administrative overheads consume up to 25% of health budgets in some countries, with low-income nations disproportionately affected (11.2% vs. 3.6% in high-income countries). These disparities stem from fragmented governance, redundant processes (e.g., manual claims adjudication), and outdated Information Technology (IT) infrastructures. Such inefficiencies delay care delivery, exacerbate clinician burnout, and widen inequities in resource-poor regions. Concurrently, demographic shifts – such as aging populations and the rising prevalence of chronic diseases – are reshaping healthcare demands. By 2050, 22% of the global population will be over 60, with the number of individuals aged 80+ tripling to 426 million. (WHO, 2025). Older adults incur healthcare costs 3–5 times higher than younger populations, driven by chronic conditions like diabetes and hypertension, which affect 1 in 3 adults globally (Maresova et al., 2019; Buchner & Wasem, 2006). This demographic shift leads to higher per capita healthcare expenditures, driven by the higher costs associated with chronic diseases and end-of-life care. The continuous rise in healthcare costs threatens the fiscal sustainability of healthcare systems, necessitating policy adjustments and reforms. Additionally, the growing need for long-term care services adds to the financial burden, particularly for lower-income individuals (Mason & Miller, 2018). These challenges are compounded by fragmented data, regulatory complexities, and healthcare professional burnout. Traditional reform efforts such as policy overhauls, gradual IT upgrades, and workforce expansion have proven inadequate to address the coexistence of rising demands and systemic inefficiencies. The scale and complexity of these interwoven challenges call for fundamental changes that can learn, adapt, and operate at a pace and precision beyond human capacity alone.

Artificial Intelligence (AI), particularly Intelligent Agents (IAs) – autonomous systems capable of learning, reasoning, and interacting with users – have emerged as pivotal tools to optimize healthcare workflows, reduce costs, and improve service quality. Notable IAs applications now span diagnostic assistance, virtual patient consultations, Robotic Process Automation (RPA), and conversational chatbots. For instance, according to the study by Darbandi et al. (2024), AI-powered diagnostic tools have achieved accuracy rates of up to 98.58% in early breast cancer detection using Recurrent Neural Networks (RNNs). This high accuracy has the potential to directly reduce diagnostic delays and improve patient outcomes. In administrative contexts, RPA platforms like UiPath have cut processing times for scheduling, records management, and billing by automating repetitive tasks, freeing clinician capacities for patient care (Debbadi & Boateng, 2025). However, despite these proven capabilities, the adoption remains limited due to technical, operational, and ethical barriers.

Technical barriers are among the most pressing. Fragmented IT infrastructures and data silos are common pitfalls in healthcare information systems, as most institutions often rely on exclusive protocols and incompatible platforms, which cause significant challenges for cross-platform data exchange. For example, diagnostic agents requiring real-time access to pharmacy inventories or patient histories may fail if subsystems use proprietary protocols (Sousa et al., 2024). A diagnostic tool trained on European populations may misdiagnose Asian patients due to genetic or demographic differences (Jeyaraman et al., 2023). The "black box" nature of these systems also erodes trust: Physicians hesitate to rely on AI recommendations they cannot interpret, such as imaging agents prioritizing non-pathological artifacts (Sadeghi et al., 2023).

Operational challenges also reinforce the limitations. Integrating IAs into clinical workflows often requires restructuring established processes. For example, emergency dispatch agents may conflict with operating room schedulers over resource prioritization, necessitating new monitoring interfaces. User resistance is equally critical: only 38% of physicians fully accept AI-assisted diagnostics, primarily due to liability concerns (Almyranti et al., 2024), while patients often abandon conversational agents (e.g., mental health chatbots) due to perceived "mechanical responses" or a lack of empathy (Milne-Ives et al., 2020). Training gaps worsen adoption; many institutions rarely provide ongoing staff education, leaving end users unprepared to use these tools effectively. Another concern comes from financial sustainability. High upfront costs – from hardware upgrades for model training to Application Programming Interface (API) feeds for large language models and ongoing data annotation – strain budgets, particularly for public hospitals, which have tight budget controls. Insurers demand a proven Return On Investment (ROI), yet healthcare outcomes often manifest over years, complicating cost-benefit analysis.

Ethical and regulatory risks further slow the progress of integration. Medical data sensitivity demands rigorous privacy safeguards; however, encryption methods like and secure Security Multi-Party Computation (SMPC) degrade system performance, rendering them unsuitable for emergencies (Jeyaraman et al., 2023; Elendu et al., 2023). Moreover, ambiguity in liability deters adoption. For example, if an AI diagnostic error stems from flawed training data, unclear protocols may leave developers, hospitals, or physicians vulnerable to legal action. Patient autonomy is another risk area. Studies show 73% of patients follow AI over physician advice when conflicts arise, threatening the integrity of shared decision-making (Stokel-Walker & Van Noorden, 2023).

The purpose of my master's thesis is to explore and understand how IAs and adoption of AI can optimize healthcare business processes. As healthcare systems face challenges like rising costs, resource shortages, and inefficiencies, IAs offer potential solutions by automating tasks, improving decision-making, and enhancing patient care. However, their adoption is hindered by technological, regulatory, and ethical barriers. Therefore, there is a pressing need to assess their effectiveness, challenges, and perceived value from the perspective of end users. This thesis will thus explore the gap between the theoretical promise and practical

implementation of IAs in healthcare by examining their potential to optimize business processes and the challenges hindering their integration. While existing literatures largely focus on the technical capabilities of AI systems, fewer studies explore organizational, regulatory, and human-centric factors that shape successful implementation.

Based on the above explanation, the thesis pursues four goals:

- Analyse the role of IAs in healthcare sector and examine how they are used in healthcare sector in administrative and operational applications.
- Identify the key factors that influence the adoption of IAs.
- Assess the challenges of integrating IAs, including data interoperability, resistance from healthcare professionals, and regulatory compliance.
- Analyse the commercial viability of IAs in terms of business models and business potential.

By applying technical, ethical, and operational considerations, this thesis aims to provide actionable insights for healthcare stakeholders exploring AI-driven automation. Ultimately, the findings are intended to inform decision-makers on how IAs can enhance workflow efficiency, reduce administrative burdens, and improve patient outcomes.

To achieve the research objectives, this thesis will use a mixed-methods approach, combining theoretical analysis with empirical research. The theoretical part will involve a comprehensive literature review to examine the concepts and characteristics of IAs, their current applications in healthcare (e.g., patient management and administrative automation), and the regulatory and ethical challenges they face. This phase will establish a strong contextual background for the study.

To complement the theoretical analysis, the empirical part will include semi-structured interviews with key stakeholders from selected companies in healthcare sector to gain insights on best practices and challenges. The primary qualitative data was obtained through semi-structured, in-depth interviews conducted virtually in a one-on-one format. The interviewees included three senior executives from distinct sectors: a leadership representative from Baichuan AI focusing on clinical AI deployment and trust-building; a digital health lead at Alipay overseeing platform-based IAs; and the informatics director of a tier-1 public hospital involved in institutional AI governance and implementation. Each participant offered insightful opinions shaped by direct experience with IAs integration in industrial, platform, and public healthcare contexts.

During the preparation of this master's thesis, I also used the assistance of generative AI tools namely ChatGPT and DeepSeek. ChatGPT was used to help me brainstorming, support grammar correction and the interpretation of relevant academic literature. DeepSeek was mainly used to assist with the transcription of the recordings. Throughout this process, I ensured that the speakers' content remained authentic and accurately represented. Anyhow,

these GenAI tools are solely served as language and organizational aids to clarify complex concepts and streamline the writing process. All results are carefully reviewed and verified to ensure accuracy and maintain academic standards.

The thesis is organized into five main chapters. Chapter 1 introduces the research problem, establishes the purpose and goals of the study, and outlines the methodology. Chapter 2 presents the theoretical foundations and literature review, defining IAs, exploring relevant theoretical frameworks, further investigates current applications, economic impacts, and challenges in the healthcare sector. Chapter 3 describes the empirical research design, detailing the implementation of semi-structured interviews conducted to gather practical insights from industry stakeholders. Chapter 4 discusses the key findings, interpreting them in light of the theoretical frameworks while acknowledging study limitations and proposing development for future research. Finally, Chapter 5 concludes the thesis by summarising the critical contributions and offering recommendations for optimising the adoption and integration of IAs in healthcare business processes.

2 LITERATURE REVIEW

This chapter establishes the theoretical foundations of the study and reviews relevant literature on the application of IAs in healthcare business processes to support the objectives outlined in Introduction.

2.1 Definition of IAs and their characteristics

The COVID-19 pandemic and accelerating demographic shifts globally are two key factors driving the development of AI-based healthcare solutions. While the pandemic accelerated the overall healthcare industry's evolution and solidified the core technological position of AI, concurrent aging populations and rising chronic disease burdens are increasing service demand while constraining resources (Zamare, 2024). These intertwined pressures create a compelling need for innovation, which AI, as a cornerstone of healthcare transformation, is uniquely positioned to address (Fu & Chen, 2025). This potential is also underscored by international frameworks like the WHO 's *Global Strategy on Digital Health (2020–2025)*, which advocates for human-centered, accessible, affordable, scalable, and sustainable digital health solutions to improve global health outcomes (WHO, 2021).

Among all the innovative technologies, AI stands out as a convincing solution for healthcare organizations to enhance services, advance drug development, optimize health system management, and assist clinical professionals in improving patient care (Hamid et al., 2024). Remarkably, the recent breakthroughs of generative AI represented by ChatGPT have further highlighted its disruptive potential in diagnostics, patient engagement, and administrative efficiency (World Bank, 2023). These developments reflect a broader trend: AI adoption in healthcare is now marked by rapid acceleration, diversification, and

integration into clinical workflows (Haug & Drazen, 2023). The world is in a sprint toward the future of AI-powered health, with nations racing to draft IAs as their star athletes – transforming them into the engines of healthcare's next evolution.

Modern and widely accepted defination of agents

Within this landscape, IAs generally refer to autonomous computational systems which are designed to perceive their environment, process information through reasoning or learning mechanisms, and take actions to achieve specific goals (Russell & Norvig, 2016). Rooted in Minsky's (1986) conceptualization of distributed intelligence – where complex behaviors emerge from simpler interacting components – IAs gained formal traction in the late 1990s as researchers operationalized them into adaptive decision-making systems (Jennings & Wooldridge, 1996). Today, they operate across health domains, from robotic surgery assistants to AI diagnostic tools, exemplifying adaptability, goal orientation, and contextual awareness (Fox et al., 2003). While there is not a universally agreed-upon definition of IAs, they are commonly and broadly characterized by four essential attributes – autonomy, reactivity, proactivity, and sociality (Delgado et al., 2020; Kapoor et al., 2024; Yadav et al., 2020).

- Autonomy: the ability to operate independently, making decisions and taking actions without human oversight. This autonomy enables agents to operate effectively in dynamic environments, handling uncertainty and complexity without human intervention.
- Reactivity: the ability to sense and react to environmental changes in real time. This capacity is critical for agents to make decisions quickly during unexpected events.
- Proactivity: the ability to anticipate and address future needs. Proactive agents engage in strategic planning, set goals, and take action to fulfill them, demonstrate foresight and goal-oriented behaviors.
- Sociality: the ability to interact and communicate with humans, external systems, or other agents. Sociality includes understanding and generating natural language, recognizing communicative signals, establishing cooperative relationships, and coordinating actions in multi-agent systems.

In general, these traits enable IAs to operate across diverse environments, by sensing changes in environment (through sensors or data inputs) and making decisions based on learned knowledge and algorithms, they execute actions that fulfill predefined objectives. In healthcare, this means IAs can be precisely applied to various clinical scenarios, such as imaging diagnostics, disease treatment, medical education, and patient services (Fu & Chen, 2025).

On the other hand, the design and implementation of IAs are grounded in a series of complex principles - perception, reasoning, planning, learning, and execution – depending on functional requirement, environmental constraints, and design objectives (Kumova & Heye, 2017). The decision-making workflow of IAs generally comprises several key stages, which is summarized in Table 1.

Table 1: The design and implementation principles of IAs

Principle	Definition	Process	Examples
Perception	Acquiring environmental data through sensors.	Sensors (physical/virtual) collect raw data → Convert data into interpretable internal representations.	Electronic Health Record (HER) interface capturing patient vitals.
Reasoning	Logical analysis using collected data and stored knowledge base.	Apply rules/logic to data → Resolve contradictions → Make predictions or identify goals.	Predictive analytics forecasting disease outbreaks.
Planning	Formulating actionable steps to achieve objectives.	Evaluate actions/consequences → Generate optimal action sequences.	Treatment plans for chronic diseases.
Learning	Optimizing decision-making through feedback and adaptation.	Analyse outcomes → Update knowledge base (supervised, unsupervised, reinforcement learning).	Personalized treatment models adapting to patient responses.
Execution	Implementing decisions through actuators.	Translate plans into physical/virtual actions → Execute via actuators.	Robotic arms performing sutures.

Source: Own work based on Stryker (2025); Lukic et al. (2025) and Russell & Norvig (2016)

These foundational principles equip IAs to address diverse challenges across domains. However, their effectiveness depends on how these principles are systematically prioritized to align with task complexity and operational hurdles. For instance, while some scenarios demand rapid, rule-driven responses to immediate stimuli, others require strategic foresight, collaborative problem-solving, or iterative self-improvement.

2.2 Relevant Theoretical Frameworks for Analyzing the Adoption and Integration of IAs in Healthcare Business Processes

To comprehensively examine how IAs are being adopted and integrated into healthcare business processes, the assessment is grounded in two prominent theories: the Resource-Based View (RBV) theory and the Diffusion of Innovations (DOI) theory. The strength of combining these two theories lies in their complementary focus: the RBV explains how IAs create value as strategic resources, while DOI illustrates how that value spreads through organizational and social systems. Together, this perspective aligns with the thesis's goal of understanding how IAs are integrated into the healthcare sector, and it helps clarify the process by which technology adoption leads to real-world impact and value realization.

2.2.1 RBV theory

*what a firm **has** and what it **can do** matters more than industry structure alone.*

The RBV theory argues that a company's long-term competitive advantage springs from its unique internal resources and capabilities, not just external market conditions (Barney, 1991). This perspective is especially relevant in the context of healthcare, where organizations are

increasingly turning to IAs as strategic tools for differentiation and operational improvement. RBV highlights two core insights that support this analysis:

- Unique resources: According to RBV, a firm's resources can be divided into tangible and intangible, with intangible resources often serving as the primary drivers of securing competitive advantages (Barney, 1991). In the healthcare context, tangible resources could be physical assets such as medical equipment, facilities, and financial capital, whereas intangible resources encompass elements such as brand reputation, clinical expertise, organizational culture, and accumulated knowledge. Within this framework, the VRIN criteria (Valuable, Rare, Inimitable, Non-substitutable) or the expanded VRIO (which adds Organization) are used to assess a resource's potential for creating strategic value. IAs, with their high-speed data processing, algorithmic precision, and capacity for continuous learning, exhibit clear VRIN characteristics. For example, a hospital's proprietary diagnostic agent for early disease detection could constitute a critical competitive resource if it provides superior predictive accuracy (valuable), is not widely available in the market (rare), cannot be easily reverse-engineered or replicated by competitors (inimitable), and no alternative technology can offer equivalent outcomes (non-substitutable).
- Dynamic capabilities and Integration: Beyond possessing valuable resources, organizations also need to be adaptable to dynamic capabilities. They must continuously learn, adjust, and combine their own strengths with outside expertise to keep up with shifting markets and competition (Teece et al., 1997; Barney, 1991). In healthcare, this applies to the ongoing integration of IAs into complex clinical environments. Medical institutions must develop absorptive capacity such as identify, assimilate, and apply new knowledge to seamlessly integrate IAs into existing workflows. This may include aligning decision-support tools with EHR systems, fostering collaboration across IT and clinical departments, and continuously refining use cases based on feedback and evolving medical standards. A strong example of this can be seen in Alipay's deployment of its AI Healthcare Butler within its healthcare platform. This AI-powered agent connects users with a wide range of offerings, such as doctor recommendations, medical test result interpretation, in-hospital navigation, and personalized medical advice (Ant Group, 2024). The success of this integration requires beyond the technology itself, also dynamic capabilities, such as reconfiguring internal processes, fostering cross-functional collaboration, and continuously adapting the AI services based on user feedback and evolving healthcare standards.

Overall, the RBV emphasizes that sustainable competitive advantage is primarily rooted in the unique internal resources and capabilities of a firm. In the context of healthcare, the deployment of agents for administrative automation, clinical decision support, and patient management creates valuable resources that enhance operational efficiency. These competencies, once incorporated through data-driven learning and adaptation, become difficult for competitors to replicate, thereby reinforcing a sustainable competitive advantage.

2.2.2 DOI theory

DOI explains how new things spread, why some people adopt them quickly and why others resist change.

Whilst the RBV provides a strategic foundation for understanding the internal value creation of IAs within healthcare organizations, it is equally important to consider how these technologies are disseminated and adopted across healthcare systems. The selection of the DOI theory is grounded in its strong explanatory power regarding the technology adoption process. Developed by Everett Rogers in 1962, DOI theory describes the process by which an innovation is communicated through certain channels over time within a social system. It provides a robust framework for understanding how technological advancements permeate through social systems, particularly relevant in examining the adoption in public health contexts (Dearing & Cox, 2018).

DOI emphasizes the role of social networks in spreading innovation, which is particularly relevant in healthcare settings where professionals often influence patterns of technological diffusion. According to the theory, five key attributes of innovations significantly influence their adoption rates: relative advantage, compatibility, complexity, trialability, and observability (Rogers, 2003).

- Relative Advantage: Innovations that offer clear benefits over existing solutions are more likely to be adopted quickly (Rogers, 2003). In the context of IAs, relative advantage is demonstrated by their ability to automate administrative tasks and enhance diagnostic precision.
  - ***is it better than what I use now?***
- Compatibility: It measures how well an innovation aligns with the existing values and practices of potential adopters. Innovations that are consistent with current norms and experiences are more easily accepted because they reduce uncertainty and make adaptation easier (Rogers, 2003). For IAs, compatibility is evident in their ability to integrate with existing EHR systems and established clinical workflows. ***does it fit how I already work?***
- Complexity: It refers to the degree of difficulty in understanding and using the innovation. Innovations that are simpler and easier to use are more likely to be adopted because they require less effort and learning (Rogers, 2003). In healthcare, if users find IAs technically challenging, it is considered as a barrier to adoption. ***is it easy to adopt?***
- Trialability: Trialability is the ability to experiment with an innovation on a limited basis before committing to it. This characteristic reduces uncertainty by allowing potential adopters to evaluate the innovation without fully committing resources (Rogers, 2003). IAs that can be tested in clinical settings before broader implementation as more likely to achieve high acceptance. ***can I try it safely (without risk)?***
- Observability: Observability refers to the visibility of the outcomes generated by an innovation. Innovations whose results are readily observable are more likely to be adopted, as tangible benefits serve as powerful social proof (Rogers, 2003). In healthcare, when the positive impacts such as reduced waiting times or improved diagnostic accuracy are clear, adoption is more encouraged by the organizations.

can I see others benefiting from it?

A practical example of the DOI principles can be seen in a field study of examining the use of AI-powered Clinical Decision Support Systems (CDSS) in rural China. The research found that while clinicians recognized the diagnostic benefits of AI-CDSS (relative advantage), adoption was limited by several factors: poor alignment with local workflows (compatibility), a complex interface with inadequate training (complexity), and lack of transparency in algorithmic decisions (credibility) (Wu et al., 2021). Furthermore, the absence of pilot programs (trialability) and unclear performance outcomes (observability) further reduced clinicians' willingness to adopt the system. This case serves as a good example of considering DOI attributes in the successful diffusion of AI innovations in healthcare settings. Technological advancement alone cannot guarantee widespread adoption without careful attention to human and organizational factors (Dearing & Cox, 2018).

RBV and DOI theory offer a dual perspective for analyzing how IAs are applied in healthcare. From an internal standpoint, RBV helps explain how healthcare organizations can turn intelligent technologies into strategic assets by continuously investing in resources and developing capabilities. On the other hand, DOI focuses on the external environment, offering a structured way to understand how new technologies spread through complex healthcare systems. It considers how medical staff adopt new tools, how organizations adapt to change, and how social networks influence broader acceptance. Thus, this combined internal – external framework not only provides a unique insight into the strategic deployment and the adoption dynamics of IAs, but also helps interpret existing cases in healthcare.

2.3 IAs and Business Processes

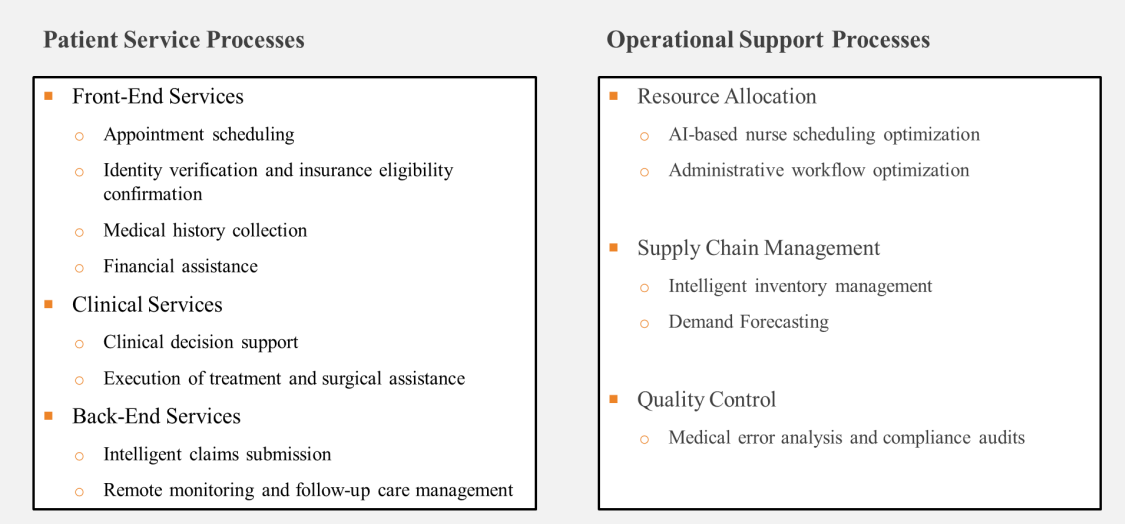
In response to the dynamic global landscape, organizations have embraced fundamental rethinking and reengineering of business processes to enhance value delivery to customers, stakeholders, and internal operations (Davenport, 1993). Healthcare is no exception. Providers must deliver high-quality care while managing costs and adapting to a constantly shifting environment. In this context, healthcare business processes refer to structured sequences of actions designed to achieve specific care delivery objectives. These encompass both administrative tasks such as appointment scheduling and clinical operations such as diagnostics. It also extends far beyond clinical interactions. As emphasized by the WHO (WHO, 2020), daily operations in the health sector also include resource management, communications, regulation, education, financial transactions, and data handling, all of which are essential for effective service delivery. Understanding these processes is essential for evaluating how AI solutions can streamline workflows, assess adoption feasibility, and generate measurable value.

2.3.1 Healthcare Business Processes

Healthcare brings together many different medical fields, such as cardiology, oncology, and pediatrics, each with distinct expertise, methods, and workflows. Defining these processes is challenging due to their inherent complexity, dynamic variability, decentralized structure, and multidisciplinary nature (Antonacci et al., 2021). The sector involves various stakeholders, such as patients, healthcare service providers, administrators, and regulators. While all share the same goal of improving patient care, each group has its own priorities and operational demands. Balancing these competing interests and ensuring clear communication across departments is a constant challenge. Lenz and Reichert (2007) broadly categorize healthcare processes into (1) organizational processes, which handle coordination, resource management, and workflow across healthcare professionals and units; and (2) medical treatment processes, which focus on patient-specific medical procedures and decision-making.

Building on this foundation, this thesis adopts a complementary functional framework, distinguishing patient service processes (direct diagnosis, treatment, and patient interaction) from operational support processes (infrastructure and administrative efficiency of medical institutions). This distinction is grounded in a dual theoretical foundation. First, it draws on established academic research in information systems, particularly the work Lenz and Reichert (2007), which highlights the structural and functional dichotomy within healthcare workflows. Second, it aligns with policy frameworks such as the U.S. Health and Human Service's Medical AI Strategic Plan (2021), which emphasizes optimizing both clinical services and management operations. Patient service processes include front-end services, clinical services, and back-end services, while operational support processes involve resource allocation, supply chain management, and quality control (see Figure 1).

Figure 1: Core business processes in healthcare



Source: Own work based on Lenz & Reichert (2007) and U.S. Department of Health and Human Services (2025)

- ***it captures functions that are UNIVERSAL across healthcare regardless of country, specialty focus or institutional context***
- ***it abstracts away from specialties but preserves core FUNCTIONS/PROCESSES***

These processes are interconnected and necessary to the effective functioning of healthcare organizations. While the specific activities may vary across settings – from primary care clinics to large public hospitals – this categorization captures a core set of business functions that are broadly applicable across the sector. It provides a consistent basis for analyzing IAs applications in subsequent sections.

IAs are usually deployed to:
- ***automate tasks***
- ***support decisions***
- ***coordinate workflows***
- ***optimize resources***

2.3.2 Functional typology of IAs within healthcare processes

To better understand the roles and capabilities of IAs in healthcare environments and align with the core processes outlined above, this thesis proposes a functional typology that classifies agents into three categories: administrative agents, operational efficiency agents, and clinical support agents (see Table 2). Built on real-world implementations and academic perspectives, this approach focuses on how agents contribute to different healthcare processes objectives, with varying degrees of autonomy, adaptability, and goal orientation (Russell & Norvig, 2016). It aligns with the structure of healthcare delivery, which involves clinical expertise, administrative operations, and patient participation. It also reflects the evolving landscape where the adoption of IAs is expected to interact with dynamic environments and manage complex tasks autonomously. By organizing this way, we can systematically evaluate their impact across diverse use cases.

Table 2: Types of IAs applied with healthcare business processes

Category	Focus Area	Examples	Benefits
Administrative Agents	Automation of back-office and patient-centric workflows	Chatbots for appointment scheduling; Automated billing systems	Reduces turnaround time; lowers operational costs; improves patient experience.
Operational Efficiency Agents	Optimization of healthcare institution logistics and resource management	Patient flow management systems; Inventory management	Improves resource utilization; reduces waiting times; cuts waste in logistics.
Clinical Support Agents	Assistance with diagnostic, treatment, and patient treatment planning	AI-powered decision-support systems; Remote patient monitoring tools	Enhances clinician decision-making; provides evidence-based recommendations; improves accuracy in complex cases.

Source: Own work based on Davenport & Kalakota (2019); Google Cloud (2025); Khan et al. (2025) and Russell & Norvig (2016)

- Administrative agents: In administrative areas, IAs handle routine operational and clerical tasks such as billing, appointment scheduling, and insurance claims processing. These agents often rely on RPA and rule-based logic, sometimes enhanced with machine learning for more adaptive execution. They can significantly reduce manual workloads

and enable healthcare staff to focus on more complex, patient-centred activities (Davenport & Kalakota, 2019). Google's 2025 AI business trends report highlights how AI applications include multiple agents with human-in-the-loop and multimodal AI are increasingly used to manage complex administrative workflows and improve the speed and accuracy of backend healthcare functions (Google Cloud, 2025). A representative example is found in Alipay's Health Butler, which supports hospitals by automating routine administrative tasks such as patient inquiries, triage assistance, and report interpretation. According to a senior executive from Alipay's healthcare strategic department (Participant P1, personal communication, April 29, 2025), this agent-powered application helps guide patients through the hospital journey from check-in to medication pickup, while collecting questions and feedback in real time. The information is then sent to doctors for more efficient and personalized care.

- Clinical support agents: Clinical agents operate at the core of patient care. They are designed to support medical decision-making, diagnostics, and treatment planning. They often utilize machine learning, deep learning, and natural language processing to interpret medical data, identify patterns, and recommend evidence-based interventions. Their functions range from diagnosing diseases through image recognition to suggesting personalized treatment plans using precision medicine models (Davenport & Kalakota, 2019). One representative case is Baichuan's AI Doctor, an intelligent medical assistant developed using advanced Large Language Models (LLMs) and machine learning. Baichuan's AI Doctor can interpret patient inputs, analyse clinical symptoms, and provide real-time diagnostic suggestions and personalized health advice. Designed to support both clinicians and patients, it enhances diagnostic accuracy, streamlines patient triage, and reduces the burden on staff (Participant P2, personal communication, May 10, 2025).
- Operational efficiency agents: As healthcare systems face increasing pressure to deliver timely, high-quality care under constrained resources, operational efficiency agents are designed to manage the complexity by optimizing resource use, capacity planning, and process coordination across healthcare facilities (Fox, 2024). In contrast to traditional rule-based automation, these agents often rely on machine learning, multi-modal large models, and real-time data analysis to dynamically adjust to changing conditions thus reduce delays and inefficiencies in service delivery. A real-world example is Agentforce by Salesforce, which help healthcare organizations validate insurance coverage in real time and estimate patient out-of-pocket costs, reducing delays in care and simplifying billing workflows. It streamlines back-office functions and improves both operational performance and patient satisfaction (Khan et al., 2025).

In summary, this typology clarifies distinct IA roles in healthcare delivery. As healthcare continues to shift toward more digital and data-driven, understanding these categories is essential for analysing adoption challenges, implementation strategies, and measurable impacts on efficiency and patient outcomes.

Market Validation - does this (product) solve a real problem for enough people?
it's the process of checking whether customers actually need and willing to pay before you fully invest in it.

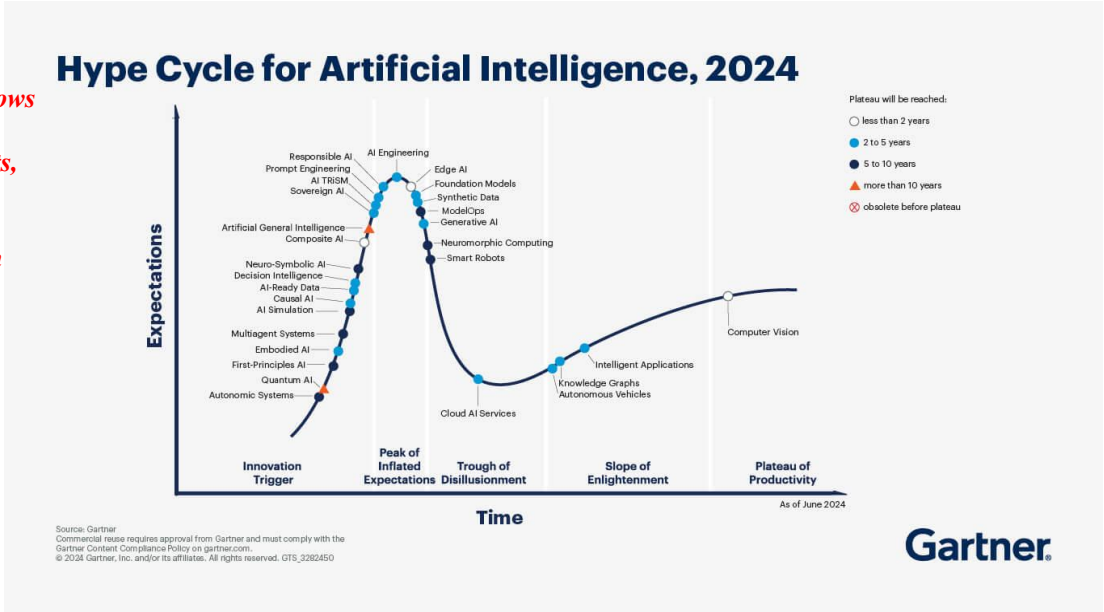
2.3.3 Market Validation of IAs

Recent market research reveals the growing operational value of IAs in healthcare. According to Gartner's 2024 Hype Cycle for Artificial Intelligence (see Figure 2), practical and deployable innovations such as multiagent systems, edge AI, and domain-specific applications are progressing steadily toward mainstream adoption. The hype around generative AI is beginning to settle as businesses move past the excitement over foundation models and shift toward applications that deliver real value. This change in focus is helping to drive momentum toward autonomous AI. While current AI systems still lack full autonomy, research groups are making steady progress in developing IAs that can interact with their environments and work toward defined goals – a process that is expected to unfold gradually over time (Gartner, 2024). This trend highlights the urgent momentum for healthcare organizations to integrate IAs into their business and clinical workflows.

it's a model that shows how expectations about a new Tech change over time.

The AI Hype Cycle shows that AI first gets over-hyped, then disappoints, and finally becomes useful when organizations focus on practical applications

Figure 2: The AI Hype Cycle



Source: Gartner (2024).

However, to fully understand the reasons behind the growing maturity of IAs within healthcare, it is necessary to examine the underlying drivers. Three primary factors support this trend. First, their development is underpinned by decades of AI research and industrial application, which reduces innovation risks and aligns new investments with tested theoretical frameworks. Second, strategic endorsements from major technology firms, such as IBM Watson Health, Microsoft, and Google DeepMind signal market confidence and accelerate development cycles (Miller, 2013). Third, and perhaps most crucially, there is a lack of viable alternatives. While LLMs like GPT-4 have demonstrated impressive generative potential, they are also prone to hallucinations risks. These refer to false or fabricated information generated by the model without a clear factual basis. This is especially

dangerous in healthcare settings, where hallucinations can be subtle and convincing, making it difficult for even experienced physicians to distinguish (Xi et al., 2025). This limitation creates demand for more reliable hybrid systems that combine rule-based logic with generative capabilities (Almyranti et al., 2024; Xi et al., 2025).

In light of these technological advancements and market validations, IAs have moved beyond experimental innovations to become practical enablers of healthcare business process optimization. Building upon this foundation, the next section explores current IAs deployments within the healthcare sector. Through real-world case studies and operational examples, it illustrates how these technologies are reshaping clinical, administrative, and operational workflows across healthcare organizations.

2.4 Current applications of IAs in healthcare

The rapid evolution of AI has increasingly shifted focus from LLMs to IAs, especially within healthcare context. While LLMs such as GPT-4 remain influential for knowledge generation, concerns over hallucination risks, operational inefficiencies, and high maintenance costs have prompted interest in more practical, goal-driven solutions (Gartner, 2024). IAs build on LLM capabilities by introducing autonomous decision-making and modular architectures, have risen rapidly as powerful alternatives (Lukic et al., 2025). These modular architectures enable complex healthcare tasks to be broken down into smaller, manageable components, each handled by a specialized agent. This modularity enhances flexibility and allows organizations to retain control over proprietary data, addressing critical privacy concerns (Yang et al., 2024). As a result, the emphasis is shifting away from "big model wars" – simply building ever-larger models – towards developing intelligent, operational agents that deliver measurable business value and scalable implementation (Purdy, 2024). Against this backdrop, this section explores the current IAs deployment across operational, administrative, and clinical settings.

The strategic positioning of IAs within healthcare business processes can be broadly understood through three domains: Creation, Enhancement, and Replacement (see Figure 3). This conceptual classification is grounded in both the evolutionary logic of technological diffusion and practical implementation patterns observed in the healthcare domain. From a technological evolution perspective, the integration of emerging technologies into healthcare typically follows a progressive trajectory: (1) creation of entirely new business scenarios; (2) enhancement of existing processes; and (3) eventual replacement of manual functions. This classification approach offers several methodological advantages. It is comprehensive, capturing the full life cycle of technological adoption from inception to maturation; distinctive, in that it defines the evolving value propositions of AI at different stages of development; and operationally actionable, offering medical institutions a practical roadmap for phased AI deployment (Huang et al., 2024; Iqbal et al., 2016; Isern & Moreno, 2016; Lin et al., 2012; Moritz et al., 2025). Notably, these purposes are not strictly separated but

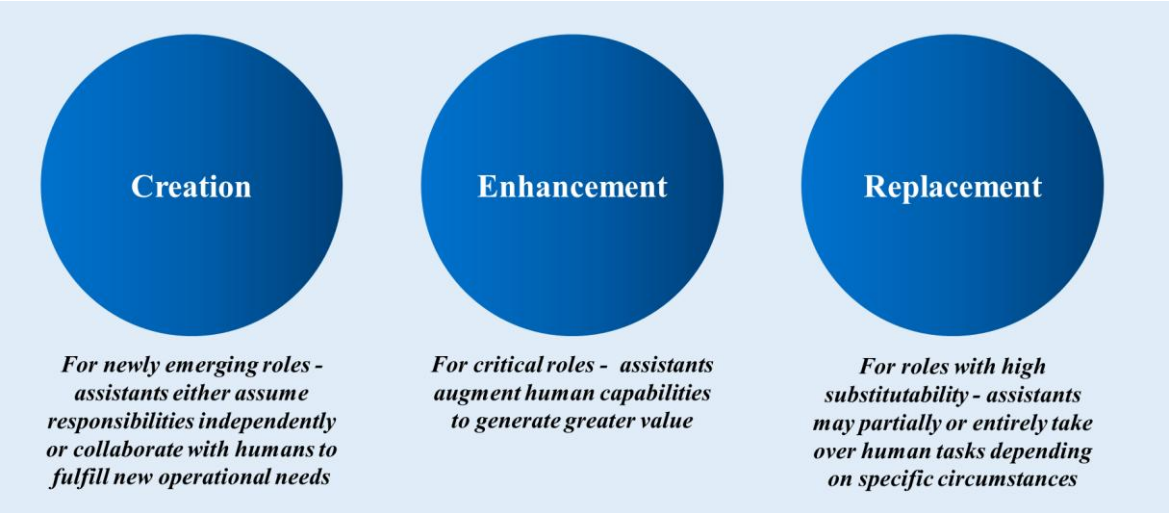
interact and evolve over time. For example, a role initially introduced for enhancement may, over time, displace existing roles as efficiency gains render certain functions obsolete.

- Creation**:** IAs are used in healthcare for building new applications or support systems that were previously infeasible. These often appear in the clinical and operational domains. For instance, Lin et al. (2012) propose an Intelligent Agent Web-based Medical Information Support System (IAW-MISS) designed to assist users in identifying the most suitable medical centers based on specific symptoms or health conditions. It represents a new kind of healthcare decision-support tool, using business intelligence and IAs to process online data and filter relevant options. It illustrates how IAs could enable entirely new service models in patient navigation and care selection, underscoring their creative potential in digital health design. Similarly, Moritz et al. (2025) present the Multi-Agent Systems for Healthcare (MASH), which are decentralized yet coordinated networks of specialized AI agents. These applications illustrate how IAs can create entirely new digital workflows, supporting more personalized and data-driven healthcare. It adds another layer to this idea by showing how coordinated, specialized AI agents can support healthcare services that would otherwise be impossible using only human effort.
- Enhancement: IAs augment existing healthcare workflows by improving coordination and decision support. For instance, Isern and Moreno (2016) highlight how agents have been applied to improve collaboration between healthcare professionals and enhance digital health tools. Through their autonomous and proactive capabilities, agent-based systems facilitate the integration of various healthcare components, such as EHR and decision support systems. Additionally, they enhance collaboration by enabling the secure sharing and management of medical data, ultimately improving the efficiency and effectiveness of healthcare delivery. In the MASH model, AI agents work alongside clinicians by offering real-time recommendations or analysis, helping them make better-informed decisions without taking over their roles (Moritz et al., 2025).
- Replacement: IAs can take over routine or repetitive tasks in healthcare workflows traditionally done by humans. These applications typically consolidate gains achieved during the enhancement phase. For example, Huang et al. (2024) reference a study that demonstrates how AI can be employed to support physicians in their decision-making processes, specifically in prescribing antibiotics to prevent multidrug-resistant infections. By considering a range of variables from patient and hospital data, the AI system can help reduce the overuse of antibiotic treatment. Iqbal et al. (2016) discuss how mobile multi-agent systems can access and manage distributed healthcare services, streamlining processes that once required manual coordination. In logistics-heavy environments, IAs can autonomously manage supply chains, reorder medical inventory, and allocate resources across departments (Jebbor et al., 2022). IAs can also be used to automate clerical roles that once required human intervention like billing and insurance claims processing. They can automate claims processing while ensuring that billing is accurate and compliant with relevant regulations (Thesmar et al., 2019). These applications

showcase how IAs are assuming functions previously performed manually in both clinical and administrative workflows.

Taken together, these examples highlight not only what these agents do, but where they are doing it, and what that means for the future of healthcare transformation. While this section has drawn from academic literature to define these roles, Chapter 3 will examine how these strategic roles manifest in practice through interview-based case studies.

Figure 3: The strategic roles of IAs



Source: Own work based on Huang et al. (2024); Iqbal et al. (2016); Isern & Moreno (2016); Jebbor et al. (2022); Lin et al. (2012); Moritz et al. (2025) and Thesmar et al. (2019)

2.5 Business and Economic impact of IAs in healthcare

Although the adoption of IAs in healthcare is accelerating, monetization strategies remain fragmented and inconsistent. Despite growing interest from both public and private stakeholders, there is no unified consensus on sustainable commercialization models. This has left much of the market in a trial-and-error phase, where pricing, reimbursement, and long-term financial planning vary widely across use cases and regions. As Abràmoff et al. (2022) observe, few structured frameworks exist for valuing autonomous AI in healthcare, and those that do are often ad hoc - lacking standardized methods for capturing costs liability, and equity. This absence of economic clarity poses a significant challenge for stakeholders seeking to scale AI systems in a financially viable manner.

While the primary objectives of agent adoption are to reduce operational costs and enhance patient outcomes, achieving these goals requires more than technological implementation – it demands a clear understanding of financial dynamics and long-term value creation. Wolff et al. (2020) identify a significant gap in the literature regarding thorough economic impact assessments of AI in healthcare. Given the resource-intensive and high-stakes nature of healthcare systems, it is important for stakeholders to justify AI investments through clear

economic reasoning. Thus, there is a need for comprehensive economic analyses to inform decisions regarding AI adoption. Without such clarity, organizations risk adopting tools that may offer short-term novelty but fail to deliver sustainable value. Moreover, addressing the financial challenges associated with AI adoption is just as important for successful integration in care delivery as tackling security concerns (Rony et al., 2024). As such, this section examines the business and economic impact of IAs by analyzing the cost-benefit structure to determine their true business viability in healthcare transformation (see Table 3).

Table 3: Cost and Benefit model of agent adoption in healthcare organizations

Cost Structure		**Benefit Structure**	
Initial Investment	Software development, Infrastructure integration, Model training.	Efficiency Gains	Automation of repetitive work
Operational Costs	Ongoing system maintenance, System update.	Cost Reduction	Optimized workflows result in lower expenses.
Data Management Costs	Costs for securing and managing private data, Cybersecurity.	Revenue Enhancement	Short term revenue gain through improved operational efficiency.
Human Resources	Costs associated with user training and skill development	Improved Patient Outcomes	Enhanced user experience and proactive patient care.

Source: Own work based on Caruso et al. (2023); Rony et al. (2024) and Wolff et al. (2020)

already demostrated the LT economic return

Traditionally, the predominant costs components in tech adoption have been initial investment and ongoing maintenance expenditures. However, growing evidence supports that these upfront investment yield substantial long-term economic returns. McKinsey's generative AI research reveals that pharmaceuticals and medical products stand to gain $60–110 billion annually from AI adoption, primarily through accelerated Research and Development (R&D) cycles (40–50% faster molecule discovery) and optimized clinical trials (Chui et al., 2023). Abràmoff et al. (2022) outline a comprehensive cost structure for AI systems deployment, which includes not only initial R&D and integration expenses but also recurring costs related to liability, performance monitoring, and equity assurance. However, when deployed strategically, agents can generate meaningful resource utilization to deliver sustainable returns, such as introducing a well framework of reimbursement or robust operational logic. Khanna et al. (2022) highlight AI technologies have already demonstrated cost reductions compared to conventional methods, particularly in treatment processes where efficiency gains are more pronounced than in diagnostics. IAs have the proven ability to expedite diagnosis processes and smooth out treatment workflows, which contribute to cost reduction at the system level while sustaining the quality of care.

Even so, reaching the full economic value of agent adoption is far from straightforward. A review of existing AI-in-healthcare studies reveals that most economic evaluations remain

fragmented, often focusing narrowly on isolated cost or cost-saving elements without presenting a complete financial picture (Wolff et al., 2020). Without a comprehensive view of the financial input required, the true ROI remains speculative at best. Moreover, without clear benchmarks for performance, the limitation in credibility of current evaluations leaves stakeholders without a reliable basis for comparing AI solutions to other viable options. Thus, addressing these limitations is essential for moving beyond hype and toward grounded, evidence-based adoption strategies.

These shortcomings are especially evident in real-world market behavior, where commercialization models diverge significantly based on organizational context. Against this backdrop, a recent analysis of the Chinese healthcare AI sector illustrates this clearly, revealing a distinct divide in monetization approaches between large technology platforms and Small-to-Medium-Sized Enterprises (SMEs). The decision to distinguish between large enterprises and SMEs is informed by the stark contrast in their capacities for digitalization transformation. While larger firms often have more financial and technological endorsement required to adopt and scale IAs, SMEs typically operate under tighter resource constraints and limitation in funding. Their opposing characteristics create a unique analytical value, which responds to existing literature that emphasizes the structural and strategic differences in business model formation (Miller et al., 2021; *Parikh & Helmchen, 2022*). Therefore, the following table presents a comparative structure for analyzing the economic and business implications of IA deployment in healthcare contexts (see Table 4).

commercialization model differ based on org

Table 4: Business and Economic Impact by Category

Dimension	Large Platforms	SMEs
Strategic Focus	Ecosystem integration; Cross-industry resource integration; Holistic healthcare solutions	Rapid ROI; Scalable business models
Business Models	Subscription services; Device sales; User stickiness	Subscription services; Pay-per-use models; Software as a Service (SaaS) licencing
Adoption Patterns	Embedded into existing platforms	Targeting underserved markets
Cost Structure	High fixed costs; Long-term R&D costs	Variable cost dominance; Rapid iteration costs
Value Realization Approach	Ecosystem revenue; Network effect	Agile delivery; Pay-for-performance
Risk Profile	Technology lock-in risks; Interoperability with legacy systems	Customer retention challenges; Technological obsolescence

Source: Own work based on Miller et al. (2021) and Parikh & Helmchen (2022)

Large platforms exemplify a strategic approach to IAs integration, focusing on long-term value creation rather than immediate monetization. For these domain players, agents are

viewed as core capabilities that enhance their digital ecosystems instead of direct revenue streams. Leveraging extensive resources and advanced technological infrastructure, they seamlessly integrate agents into broader digital ecosystems. This aligns with Jacobides et al. (2018), who argue that platform orchestrators fundamentally reshape value creation by coordinating modular contributions from various actors. In such systems, monetization shifts from individual components to the cumulative value of the ecosystem, rendering direct revenue extraction from each module strategically secondary. This contrast is also reflected in broader enterprise tech trends. A recent McKinsey report highlights how large firms are shifting toward cloud infrastructures and as-as-service models to realize flexible operating expenditures. There is an increasing trend in leveraging FinOps (financial operations) to automate spending decisions and allocate digital resources more intentionally (Baig et al., 2025). This supports the idea that the position of IAs is less about short-term revenue and more about strategic capability in platform-driven models.

A case in point is Alipay. Based on an interview with an Alipay's employee, the platform's agent-based application, "AI Health Butler" addresses service quality gaps which often observed in the early phase of telehealth. Rather than monetizing the agent directly, Alipay uses it to drive user retention, expand service adoption, and reinforce its ecosystem's value proposition (Participant P1, personal communication, April 29, 2025). Another successful case is Apple, which integrates agents through the Apple Watch and Apple Health ecosystem. Functionalities such as heart monitoring, fall detection, and personalized health insights transform the wearable device into a proactive health management agent. Apple's strategy similarly avoids direct monetization of these AI features, instead focusing on driving hardware sales, promoting subscription services like Fitness, and enhancing user loyalty within the wellness community.

In contrast, SMEs tend to adopt a more agile and resource-efficient strategy. Without the financial buffer or technological infrastructure of larger firms, SMEs often face greater pressure to demonstrate short-term ROI and maintain cash flow stability (Participant P2, personal communication, May 10, 2025). Therefore, these players are driven to prioritize rapid monetization models, such as per-use payments, SaaS licensing, or modular tools targeted at niche healthcare applications, in order to satisfy both market demand and investor expectations (Parikh & Helmchen, 2022). In the digital economy, a sobering reality is that SMEs are typically embedded within ecosystems dominated by large platforms in the role of "complementors". This positioning requires them to focus on vertical niche areas and provide solutions that can quickly generate measurable value. Without such focus, SMEs risk losing their competitive space to standardized services or leading enterprises with broader capabilities (Jacobides et al., 2018). Research (Nambisan, 2017) on digital entrepreneurship advocates that startups should pursue less bounded initiatives and minimal viable product strategies to navigate high uncertainty and limited capital. From this perspective, the push for early commercialization reflects a deeper strategic imperative: their commercialization paths tend to be more reactive and transactional – akin to navigating a

series of short sprints rather than participating a marathon ecosystem strategy, shaped by short-term survivability rather than ecosystem consolidation.

Overall, the technology adoption strategy is not a one-size-fits-all journey. SMEs may benefit more from targeted, cost-efficient solutions that solve immediate needs, while large enterprises can pursue broader strategies that reshape entire ecosystems. For both types of organizations, it is important to align digital initiatives with their specific goals, market environment, and internal capabilities. When these efforts are realistic and well-adapted to the company's situation, digital transformation becomes more effective and sustainable.

To contextualize the commercialization of IAs in healthcare, the next section explores market-specific approaches by distinguishing between Business to Business (B2B) and Business to Customer (B2C) market structures. By examining B2B and B2C segments separately, this analysis highlights the distinct opportunities and constraints that shape IA deployment across different business contexts.

2.5.1 B2B Market segment

According to iResearch's *Healthcare Industry in the Era of Big Data* report (2023), the healthcare big data industry in China is organized into a three-tiered structure: upstream, midstream, and downstream. The upstream segment includes data sources such as hospitals and service providers like cloud computing vendors. Midstream enterprises comprise core technology firms specializing in AI technologies and natural language processing, which are key contributors to value creation within the healthcare ecosystem. The downstream focuses on application delivery, serving both Consumers-end namely individual consumers and B-end namely institutional and governmental client's markets.

Considering that IAs are deployed at the application layer to deliver real-world value, my work focuses on the market analysis in the downstream segment. At this tier, a clear distinction emerges between institutional (B2B) and consumer-facing (B2C) adoption pathways, each shaped by different operational realities and commercialization logic.

!!B2B - short term; B2C - long term!!

Both B2B and B2C markets create opportunities for IAs to create user value. However, B2B markets offer stronger near-term commercialization potential, while B2C markets present broader long-term growth opportunities. According to a report by Boston Consulting Group (2015), initial digital healthcare innovations in China have predominantly targeted B2B markets, driven by hospitals' and healthcare providers' pressing needs to enhance efficiency and operational capabilities using digital solutions such as cloud computing and big-data analytics. Consequently, public hospitals and healthcare providers are anticipated to be frontline adopters of IAs because of their direct clinical and administrative responsibilities (Xia et al., 2015). Taking Chinese market as an example, after undergoing two major technological waves in healthcare digitalization and Internet health, the room for further improvement in public hospitals is very limited, with most advancements reduced to minor

tinkering. As a result, in recent years, digitalization-related budgets have primarily focused on the medical alliances projects and medical community systems, which typically have shorter implementation cycles (Participant P3, personal communication, May 5, 2025).

Given that medical decisions directly impact patient safety and public trust, healthcare leaders adopt a cautious approach toward unproven technologies, particularly within public institutions. Such caution typically necessitates validated benchmarks before committing to large-scale deployment. Safety assurance, critical in healthcare, is often demonstrated through a "safety case," a documented rationale supported by evidence explaining why a system is acceptably safe to operate within its intended environment (Habli & Porter, 2020; Cleland et al., 2012). Normally, they require validated benchmarks before committing to large-scale deployment. Even when agents are introduced, they are often implemented discreetly, favoring internal trials over public endorsement to mitigate reputational and accountability risks (Participant P3, personal communication, May 5, 2025). As highlighted by Song and Kim (2022), digital transformation in hospitals can significantly improve operational decisions such as staffing, patient flow, and supply-chain management. But, such transformation must be built on clearly demonstrated operational value, further reinforcing the need for validated outcomes before widespread adoption. To successfully penetrate this market, two conditions must be met: 1) the agents must demonstrably achieve cost reduction and efficiency enhancement; 2) the agents must seamlessly integrate with the hospital's existing business and technical systems (Maleki & Forouzanfar, 2024).

Beyond hospitals, individual doctors represent a potential sub-segment of the to-B market. According to insights provided by the CIO of Baichuan custom-developed Doctor AI Assistants could emerge as a new category of paid products. While customized AI assistants for doctors could form a distinct product category, in most cases, doctors themselves are unlikely to be direct payers. Instead, the following parties are more likely to bear the cost (Participant P2, personal communication, May 10, 2025):

- Healthcare institutions: Hospitals may integrate IAs services into their larger service packages. This approach allows the hospital's IAs to interface with the doctor's customised agent, enabling seamless access to internal data and business processes.
- Third-party companies (e.g., pharmaceutical and medical device companies): It is common in the industry for third parties to sponsor doctors, as this facilitates direct commercial collaboration. However, given the current emphasis on medical anti-corruption measures, all parties must carefully assess potential compliance risks to avoid turning IAs into channels for improper benefits.
- Platform providers: For a certain period, platforms may offer free agents to doctors as a strategy to quickly promote their products and services, seize market share, and build user loyalty. If these platforms can create synergy with existing services (e.g., online consultations), the possibility of offering permanent free services cannot be ruled out.

2.5.2 B2C Market segment

Historically, in the Telemedicine era, AI products in the healthcare industry have faced challenges in directly monetizing Consumers-end markets, as consumers have been more inclined to pay for physical products than virtual services. However, unlike telemedicine, IAs are not purely software-based from the outset. Patients today are increasingly open to paying for advanced health solutions and services that align with their preventive care and health preferences (Deloitte, 2023). Their inherent ability to perceive and interact with hardware makes integration with physical devices a key strategic direction for expanding into the consumer market. According to Deloitte, the "new engaged patient" prefers personalized, data-driven solutions and is willing to pay for premium offerings that integrate seamlessly into their daily lives, particularly when these solutions emphasize prevention, wellness, and convenience. As such, incorporating IAs into hardware presents a promising path to enhance user experience through improved human-computer interaction.

The major hardware can be categorized as: smart home, smart mobility, and health wearables (Maleki & Forouzanfar, 2024; Nahavandi et al., 2022; Miyashita & Brady 2019).

- Smart Wearables: A wearable device equipped with an agent could provide real-time health monitoring, predictive alerts, and actionable recommendations, transforming it from a passive tool into an active health assistant. Such devices could also integrate seamlessly with other agents (e.g., glasses and headphones) in the same scenario, enabling multi-device collaboration and a more holistic user experience.
- Smart Homes: Integration into home environments enables ambient health monitoring, early intervention alerts, and seamless coordination with healthcare providers. Other possible scenario such as workplaces, agents could monitor employee well-being, reduce stress, and improve productivity through real-time health insights.
- Smart Mobility: Vehicles equipped with health-monitoring agents can detect signs of fatigue, stress, or medical emergencies, extending health management beyond traditional settings.

However, realizing these opportunities will require not only technological progress but also market validation. As Deloitte (2023) notes, medtech companies must verify consumer willingness to pay for combined software-hardware solutions in specific usage scenarios. Only those offerings that pair technological innovation with demonstrable consumer value will achieve sustainable B2C growth in healthcare.

2.6 Challenges in IAs adoption in healthcare

Like a lighthouse guides ships to safety but also exposing hidden rocks, IAs illuminate healthcare's future while it simultaneously reveals a complex array of risks spanning technical, operational, legal, ethical, and societal dimensions. As the reliance on agents grows in critical functions such as diagnosis, treatment, data analytics, and decision support,

it becomes imperative to systematically identify and address these challenges. This section explores the key hurdles that organizations must overcome to fully realize the value of IAs.

2.6.1 Technical challenges

To ensure the safe and effective integration of new technologies in healthcare, one of the foremost hurdles to be addressed is the maze of technical challenges – encompassing data privacy, security risks, and system interoperability.

The adoption of new technologies in healthcare heavily depends on collection, storage, analysis, and management of vast amounts of sensitive data, from health information to genetic blueprints. Like a double-edged sword, such data help clinicians predict diseases, tailor treatments, and improve the accuracy of diagnoses, they also expose healthcare organizations to significant ethical, legal, and operational risks in the event of data breaches or misuse (Santoni & Mecacci, 2021). For instance, with the widespread use of smart wearables and remote healthcare systems, real-time health data is constantly being collected and shared. While this is useful for chronic disease management and emergency response, these data streams often lack sufficient encryption, making them vulnerable to cyberattacks. Cyberattacks targeting healthcare systems have the potential to threaten patient confidentiality, alter prescriptions, or fake medical records, thereby causing irreversible harm (Kaissis et al., 2020). If healthcare institutions or third-party service providers fail to apply robust security protocols, their systems may become gateways for data breaches. To mitigate these threats, robust encryption techniques, stringent access controls, and secure data storage protocols must be deployed throughout the data lifecycle (Esmaeilzadeh, 2024).

In addition to data security risks, system interoperability and integration also present significant technical barriers. The diversity of healthcare information systems and technologies used across institutions leads to fragmentation, making it difficult to achieve seamless data exchange. Each system may use proprietary formats and communication protocols, hindering the ability to integrate and share information effectively (Cardoso et al., 2014). Achieving interoperability necessitates the adoption of international standards, such as Health Level Seven International (HL7), alongside careful planning to integrate agents into existing IT infrastructures without disrupting clinical workflows (Cardoso et al., 2014; Linz & Reichert, 2007). Without effective interoperability solutions, IAs may face restricted access to real-time patient data, limiting their effectiveness and increasing operational risks. Technology can be healthcare's Prometheus, bringing the fire of innovation. But without strict safeguards of security, that very fire could burn the house down.

2.6.2 Operational challenges

Operational challenges associated with IAs integration in healthcare are multifaceted, encompassing workflow disruptions, data quality management, organizational hurdles, and financial barriers.

- Workflow disruption: The integration of IAs into existing healthcare workflows often necessitates significant changes, which can create unintended complexities for healthcare professionals. As Coiera (2019) argues, AI can heighten cognitive loadand and mental fatigue by requiring clinicians to monitor and verify outputs alongside traditional duties. Besides, workflow disruption can also blur lines of clinical responsibility. When decision-making is partially delegated to agents, uncertainty may arise regarding liability for outcomes, especially when AI recommendations conflict with clinical judgement (Morley et al., 2020). Addressing this requires not only technical solutions, but also thoughtful organizational change management strategies that safeguard professional autonomy, clarify responsibilities, and maintain workflow coherence.
- Data quality management: The performance of agent systems depends heavily on the quality and diversity of the training data. However, healthcare datasets often underrepresent certain demographic groups, leading to biased outputs and exacerbating existing healthcare disparities. For instance, a 2018 study at the MIT Media Lab exposed significantly higher error rates for darker-skinned individuals in commercial AI systems (Buolamwini & Gebru, 2018). Similarly, Seyyed-Kalantari et al. (2021) found that chest X-ray diagnostic algorithms underperformed for female, Black, and socioeconomically disadvantaged patients due to biases in training data sourced from well-resourced urban hospitals. Both findings highlight a critical issue in the realm of healthcare AI: the lack of diversity in training data can result in inequitable model performance and jeopardize patient safety.
- Organizational hurdles: User resistance and inadequate training represent major organizational barriers. Healthcare professionals may be skeptical of AI's reliability and accuracy, particularly given the "black box" nature of many agent-based models (Bjerring & Busch, 2021; Char et al., 2018). Especially in high-stake clinical contexts, clinicians may be reluctant to rely on AI outputs they cannot fully interpret or verify (Kahn & Baron, 1995). And if AI systems are not user-friendly or intuitive, their perceived complexity can further erode user confidence. Healthcare professionals require comprehensive training to understand how to effectively use AI tools and interpret their outputs. Addressing these issues requires not only the design of more interpretable agent-based systems but also the provision of comprehensive, role-specific training programs to equip clinicians and administrative staff with the skills necessary to effectively work alongside new AI technologies (Kaissis et al., 2020).
- Financial obstacles: From the financial standpoint, the barriers to AI adoption are often related to feasibility and ROI. The high upfront costs of acquiring hardware and software, coupled with the need for ongoing investment in system maintenance, infrastructure

upgrades and data management can be prohibitive (He et al., 2019). These financial demands can strain healthcare budgets, particularly for startups and SMEs, thereby limiting equitable access to the benefits of IAs. Addressing these financial obstacles requires careful cost-benefit analysis and strategic investment planning to support the sustainable integration into healthcare systems.

2.6.3 Legal and Regulatory Concerns

The deployment of IAs normally involves multiple stakeholders – technology developers, healthcare provider, patients, and policymakers – each operating with distinct often conflicting priorities (Esmaeilzadeh, 2024). Tech developers pursue speed and scalability; healthcare providers emphasize clinical efficacy, safety, and workflow integration; patients expect transparency and privacy; while policymakers focus on public trust and compliance. According to Kuan (2019), ambiguity surrounding liability and the opaque decision-making processes of AI ("black box" concerns) significantly hinder trust and adoption in clinical environments. These concerns are especially pronounced in public institutions, where decisions are heavily scrutinized and tied to reputational and legal risks. As more stakeholders engage in the development and adoption processes, the ecosystem becomes increasingly complex, creating tensions in translating legal, ethical, and operational principles into practical governance mechanisms (Morley et al., 2020). To better illustrate the complex interplay of roles and responsibilities among these stakeholders, a Responsible, Accountable, Consulted, Informed (RACI) matrix (see table 5) is presented below, offering a structured overview of who is responsible, accountable, consulted, and informed at different stages of IAs deployment.

Table 5: The RACI matrix of deployment of IAs in Healthcare environments

Activity	**Healthcare Provider**	**Technology Developer**	**Policymakers**	**Patients**
System design and development	C	R/A	I	C
Data collection and preparation	R	C	I	C
Model validation and testing	R	R	C	C
Deployment and integration	R	R	C	I
Security compliance and data governance	R	C	R	C
Adverse event investigation	R	C	A	C
Patient consent and communication	R	I	C	A
Ongoing monitoring and feedback	R	R	A	C

Source: Own work based on Esmaeilzadeh (2024); Kuan (2019) and Morley et al. (2020)

- Table notes: R = Responsible; A = Accountable; C = Consulted; I = Informed. Multiple stakeholders may share responsibilities at different stages of IAs deployment.

The RACI matrix reveals the complexity of stakeholder interactions during the deployment of IAs in healthcare. Notably, overlap responsibilities between technology developers and healthcare providers during system design, deployment, and ongoing monitoring may create ambiguity in decision-making authority; Patient involvement in consent processes is essential, yet often insufficient, risking reduced trust in new technologies; Limited early engagement of policymakers to oversee the adoption processes may lead to misalignment between technological capabilities and regulatory requirements. These observations highlight the importance of clearly defined roles and continuous coordination to support the successful and ethical integration of IAs into healthcare systems.

2.6.4 Ethical and Societal Concerns

Another primary concern associated with agent integration is the potential erosion of patients' rights to informed consent and autonomous decision-making. The collection, use, and sharing of medical data presents critical privacy risks, particularly when sensitive medical records are commercialized without transparency (Bjerring & Busch, 2021). A compelling example is the 2019 'Project Nightingale' scandal, where Ascension hospitals shared identifiable health data including diagnoses, lab results, and hospitalization histories with Google under vague business associate agreements. Fifty million patient records were transferred without consent triggering public outcry and regulatory scrutiny, exposing gaps in data ownership protections (Schneble et al., 2020).

Beyond privacy concerns, the involvement of AI in healthcare may also depersonalize service delivery, undermining the trust traditionally built through direct doctor-patient communication. Patients receiving AI-generated diagnoses or treatment recommendations without sufficient explanation may experience confusion or skepticism (Longoni et al., 2019). Particularly in remote care context, limited insight into how systems operate or manage sensitive data can further erode trust. If the output of a system deviates significantly from patient expectations, it may raise concerns about fairness and accuracy. Furthermore, the uneven distribution of AI tools could exacerbate existing healthcare inequities. A 2021 study demonstrated that AI models trained on data from high socioeconomic areas underperformed when applied to lower-income populations, largely due to differences in screening frequency and clinical record completeness (Seyyed-Kalantari et al., 2021). As well-resourced institutions may rapidly adopt cutting-edge systems, underserved regions lag behind, deepening existing divides in care quality and outcomes (Celi et al., 2022).

In summary, the adoption of IAs in healthcare faces a multi-dimensional array of challenges. Technical and operational barriers complicate implementation and performance, while legal, ethical, and societal concerns pose risks to trust, accountability, and equitable access.

Addressing these interconnected hurdles is critical to releasing the full potential of IAs in creating a safer, fairer, and more effective healthcare ecosystem.

3 ADOPTION AND INTEGRATION CHALLENGES OF INTELLIGENT AGENTS IN HEALTHCARE

This chapter describes the design of an empirical study that uses semi-structured interviews to gain practical insights from healthcare industry stakeholders. This chapter details the methodology, data collection process, and analytical framework used in the previous chapters.

3.1 Description of research approach and research design

The main purpose of this research is to investigate how IAs can optimize healthcare processes when it comes to improving efficiency, decision-making, and patient care. It also explores what it takes for these technologies to be successfully adopted and integrated amid technical, operational, and regulatory challenges. The aim is to understand both the strategic value and real-world barriers of IAs adoption across diverse healthcare contexts.

To contextualize this thesis, Chapter 3 introduces three case-study entities: (1) AI technology supplier; (2) platform ecosystem provider; and (3) clinical institution. Due to confidentiality agreements, interviewees are anonymized, such as “Baichuan CIO” or “Hospital Informatics leader”. Anonymization does not compromise data validity or inhibit actionable recommendations.

3.1.1 Research methodology

For years, researchers have acknowledged the value of qualitative methods in enhancing their study. These approaches are especially useful for uncovering contextual insights, refining theoretical frameworks, and developing culturally appropriate measurements (Creswell & Poth, 2016). Since the subject of this master thesis is to explore the challenges surrounding the adoption and integration of IAs in healthcare business processes, a qualitative approach (semi structured interview) is particularly suitable. There are three main reasons for this selection:

First, to explore complex and subjective contexts of challenges, a qualitative research design is essential. The barriers to IA adoption are multifaceted, ranging data interoperability to professional resistance, all of which are difficult to quantify. As a qualitative tool, interviews stand out for their ability to dig beneath the surface. By capturing stakeholders' real-world experiences, hidden concerns and decision-making rationale, we can uncover the "why" and "how" behind the challenges (Creswell & Poth, 2016).

Second, the healthcare ecosystem involves a diverse cast of stakeholders, from frontline nurses and physicians to IT specialists, administrators and policymakers, each with different incentives, vocabularies and concerns. A qualitative approach allows us to capture multiple perspectives in participants' own words, compare points of view across roles, and bridge the gap between the key stakeholders (Creswell & Poth, 2016). This richness is critical for developing theories that are sensitive to the dynamics of healthcare environment.

Lastly, there is potential for laying the foundation for future theories. As the integration of IAs in healthcare is still an emerging field, existing theories are limited. Qualitative method is ideally suited to grounded theory building, enabling us to generate or refine conceptual models that can guide both future research and practice (Richards, 2007). In general, a qualitative design provides the depth and flexibility required to fully explore the adoption and integration challenges of IAs in healthcare business processes.

3.1.2 Data collection

To support the empirical phase of this work, qualitative data collection is carried out through semi-structured interviews. While a survey-based approach was initially considered, the complexity of the topic and limited availability of specialized experts lead to focusing on semi-structured interviews to yield deeper and more contextually meaningful insights. This method is particularly effective as it balances guided structure with adaptive questioning. By allowing participants to articulate naturally while focusing on core themes, it's a tradeoff beyond standardized instruments.

Primary data is gathered through remote interviews conducted via online conferencing tools. To ensure a well-rounded understanding of IA adoption across the healthcare value chain, a total of three interviews is completed with professionals across different stakeholder groups representing distinct organizational perspectives:

- Chief Information Officer, Baichuan AI – representing the technology supplier perspective, this participant provides insights into model development, diagnostic performance optimization, and co-development partnerships with hospitals.
- Executive Leader of AI Business Development, Alipay Health – representing the platform ecosystem perspective, with a focus on patient-facing applications, service integration, and commercial scalability.
- Informatics Director, Tier-1 Public Hospital (China) – offering the clinical institution viewpoint, including frontline adoption constraints, clinician workflows, and regulatory compliance issues.

This purposive sampling approach supports data triangulation across the innovation-to-implementation spectrum. The cases are chosen not to represent a statistically generalizable sample but to provide maximum variation and conceptual contrast across implementation contexts. This design increases the theoretical richness of the study and allows for the

identification of both shared and divergent adoption challenges. As suggested in Guest, Bunce, & Johnson (2006) interviews continued until no substantial new themes emerged from additional data. Interviews are conducted one-on-one between March and May 2025, each lasting approximately between 40 to 50 minutes. All sections are audio-recorded with permission and transcribed verbatim. An interview protocol is developed in advance to ensure alignment with the thesis objectives while allowing for the emergence of new themes. The protocol covers five main areas (table 6):

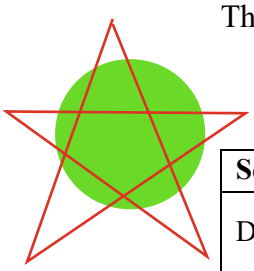


Table 6: Interview protocol development focus area

Section	Focus	Sample Questions
Domain Context	Organizational strategy and AI evolution in healthcare	What differentiates AI-driven healthcare from internet-based health services?
Use Cases & Value	Specific IA applications in administrative and operational processes	What key IA tools and agents you've deployed and the tangible benefits for clinicians/patients?
Adoption Drivers & Barriers	Stakeholder incentives, literacy gaps, resistance factors	What motivates hospitals to adopt these IAs, and what common objections have you encountered?
Integration & Governance	Data interoperability, system integration workflows, regulatory compliance	How do you navigate data standards and privacy regulations when integrating IA with existing IT systems?
Commercial Models	Revenue streams, partnership structures, ecosystem monetization	What business models your IA offerings, and how do you structure partnerships to capture value?

Source: Own work

In addition to the primary interview data, secondary sources in the form of organizational press releases and industry-specific publications are reviewed to provide contextual depth and validate emergent themes. The combination of primary and secondary sources provides a solid foundation for analyzing both the strategic intensions behind IA adoption and the real challenges observed during the implementation.

3.1.3 Data analysis

The study intends to apply a thematic content analysis, following below processes:

- Transcribing the interviews
- Translating into English
- Initial coding of the transcripts to label significant excerpts
- Grouping codes into broader categories reflecting the research objectives (e.g., "integration challenges," "adoption barriers," "business models")

- Identifying key themes and interpreting their implications for the adoption and integration of IAs in healthcare

The transcription of the interviews can be seen in Appendices 1 to 3.

3.1.4 Ethical considerations and limitations

All participants are informed of the research purpose and their rights prior to the interviews. Informed consents are obtained in writing or verbal. Participants are assured of confidentiality and the voluntary nature of their participation. Data are anonymized during analysis and reporting to protect participant privacy.

While the semi-structured interviews yield rich qualitative insights, they also limit the ability to generalize findings to the broader population. The relatively small sample size and subjective interpretations may introduce certain bias. However, the study deliberately prioritizes depth over breadth, aiming to provide actionable insights for healthcare stakeholders rather than statistically representative conclusions.

3.2 Interviews implementation

During the interview preparation phase, I took several steps to ensure high-quality data collection. I emailed each interviewee a detailed package that included an official invitation and a list of pre-formulated topics developed based on academic literature and industry reports. This preparation helped ensure that the interviews remained focused and that the data collected would be both valid and reliable.

- Baichuan AI: The first part of my conversation with Baichuan's leadership covered foundational questions about why the company has doubled down on healthcare AI for clinical scenarios. I focused on their "AI doctor vs. copilot" strategy; the technical and regulatory measures they use to build trust; and the organizational shift from being model-centric to medical-expertise focused.
- Alipay Health: In the Alipay interview, we opened the conversation by comparing the internet-era of connectivity to AI's focus on operational efficiency. Key questions explored their vision for the "AI Health Butler" ecosystem; and their data-sourcing, quality-control practices, and monetization via platform integration.
- Tier-1 Public hospital: My discussion with the hospital's informatics director began by exploring the drivers behind the recent rapid AI rollout: policy mandates, technical maturity, and operational demand. I zeroed in on themes around governance (leadership structures and priority lists), data-security protocols (local processing, anonymization, tiered access), and the feedback loops that enable AI to learn from clinical experts.

Despite varied contexts, five key questions recurred across every interview:

- Why AI, and why now? What's driving your organization's decision to invest in IAs today?
- Which cases deliver the most value? Which specific applications (e.g., record summarization, triage assistants, diagnostic copilots) are you prioritizing, and why?
- What are the main adoption barriers? From data security and interoperability to clinician trust and infrastructure gaps, what hurdles are hardest to overcome?
- How do you build trust and ensure data quality? What protocols are in place to make AI output reliable?
- What's the commercialization model? How do you capture value and what makes that approach sustainable?

By centering on these five questions, I directly assessed how IAs are perceived, implemented, and monetized across the healthcare settings.

3.3 Analysis of the results

In this part, I briefly sort out and analyse the interview results. The full transcripts of the interviews are attached in Appendices 1 to 3 of the thesis. Through the analysis of three organizations, the aim is to summarize more systematically the key obstructive factors and successful experiences that played a key role in the implementation of IAs in healthcare.

3.3.1 Baichuan Intelligence analysis

In the interview with Baichuan, the participant pointed out that in the face of the shortage and uneven distribution of healthcare resources in China, Baichuan has taken healthcare as a strategic priority. Its core mission is to improve the supply capacity of healthcare services through large-scale model technology, especially the research and development of AI doctors. As the CIO explained, the company's agent-powered pediatrician, jointly developed with Beijing Children's Hospital, is a pilot project to address major pain points in pediatric diagnosis and treatment. He admitted that although the AI Pediatrician outperformed most online consultation platforms for diagnostic accuracy during testing, they still lag behind top professionals.

During the discussion of the barriers to the adoption of IAs in healthcare, he emphasized that while governments and medical institutions are actively promoting AI in clinical settings, relevant innovations must comply with strict regulatory requirements. He commented that Baichuan is currently piloting agent-powered family doctors in key areas to deeply integrate them into community healthcare environments. From the industry point of view, he also predicted that the future healthcare would shift from hospital-centric to home-centric, from doctor-centric to patient-centric systems, where AI doctors will collaborate with existing healthcare systems to reshape the diagnosis and treatment process. To this end, Baichuan is exploring a sustainable business model to promote the popularization and adoption of

personalized AI doctors. He concluded that the success of healthcare AI depends not only on the practicality of the technology, but also on winning the triple trust of users, doctors and regulators, and finally realizing the vision of "smart life" by solving practical healthcare challenges.

3.3.2 Alipay Health analysis

On the other hand, the insights from Alipay representative highlighted the unique advantages of the platform-based strategy. The participant explained that AI in the healthcare sector can not play merely as an extended role for doctors or clinical institutions but also operate independently in specific scenarios. This dual attribute is reshaping the care service model and improving the overall efficiency. As he explained then, different from the Telehealth era, which focuses on the characteristics of connectivity, the core value of IAs shifts to efficiency optimization. The participant explicitly emphasized that agents have unique advantages in improving public health literacy and strengthening the service capacity of healthcare institutions.

Followingly, when talking about the barriers to IAs adoption in clinical environments, he stated that although healthcare agents have a bright future, we still face regulatory roadblocks and a trust gap between doctors and patients. He commented that a more pragmatic approach is to start from the back office and gradually penetrate the frontline diagnosis and treatments. When we proceeded the discussion to the key factors for physicians to embrace IAs, he mentioned an interesting phenomenon: physicians generally welcome AI tools for initial patient screening and process automation but have reservations about systems that may intervene in clinical decision-making. This selective acceptance reflects the cautious stance of the professional community on the application of technology. Taking Alipay's AI Health Butler as an example, he explained in detail its innovative value: equipping each individual with a personal health assistant. He also highlighted that the differentiation strategy of Alipay lies in its ecosystem integration capabilities, comprehensive service design, and deep synergy with the public healthcare system.

3.3.3 A tertiary public hospital analysis

The informatics director from a tertiary public hospital explained that public institutions are now investing in AI due to three key driving forces: policy orientation, technological maturity, and operational needs. China's current national policies on health innovation have created a favourable environment for technology development, and the unique value of AI in expanding the coverage of healthcare resources and addressing service gaps at the grassroots makes it a key enabler for hospital capacity building. In terms of the technology evolution path, he divided the development of AI healthcare into two stages: the current stage will primarily focus on scenario testing and tool development, and the next three years will focus on breakthroughs in augmenting clinical decision support and ecosystem

integration. This phased development insight is not only in line with the technology maturity curve but also aligned with the pace of digital transformation in medical institutions.

Followingly, when talking about the role public hospitals should play in adopting AI technologies such as IAs, he particularly emphasized the construction of a systematic innovation framework within the hospital context, including a strong governance structure and five key support mechanisms. He used practical cases to illustrate that the cooperation model between hospitals and technology companies, through the "three-layer protection system," can accelerate the implementation of AI while ensuring data security and compliance. He also shared that the continuous evolution of AI agents in clinical settings requires a "two-engine" approach, coupling technological innovation with real clinical practice.

Towards the end of the interview, he mentioned the challenges of AI adoption involve clinical validation, infrastructure limitations, data quality issues, and talent shortages. He emphasized that successful adoption requires a comprehensive approach spanning from upstream support, downstream buy-in, and a safety net of policies and regulations.

In Chapter 4, I combine these empirical findings with the theoretical framework to generate new sparks. As one interviewee pointed out incisively: "We are not deploying a certain system but writing the new genes of future healthcare."

4 DISCUSSION

Based on the analyses presented in the previous chapters, this chapter focuses on interpreting the practical findings obtained from the deployment and adoption of agents in diverse healthcare scenarios. It particularly centers on the strategic roles of IAs categorized as creation, enhancement, and replacement for understanding their functional impact. Each section uses the strategic role classification framework to interpret interview data, clarifying how IAs are currently creating new capabilities, enhancing existing processes, and replacing traditional functions within healthcare context. To complement this analysis and provide a more nuanced view of adoption dynamics, I also apply DOI theory across each case. By evaluating agents around the five core attributes of relative advantage, compatibility, complexity, trialability, and observability, this chapter sheds light on how innovation is diffused within the distinct operational, regulatory, and organizational contexts of digital platforms, AI startups, and public hospitals. These analyses aim to set the stage for actionable recommendations and clarify the areas for future research direction.

4.1 Summary and interpretation of key findings

4.1.1 The relevant findings of Alipay's deployment of IAs

As a leading digital payment company in China, Alipay has leveraged its extensive ecosystem and digital infrastructure to enter the healthcare market with its newly launched agent-based application, AI Health Butler. Built into Alipay's wallet and insurance interfaces, the Health Butler offers over thirty intelligent services – symptom triage, Augmented Reality (AR)-guided hospital navigation, automated report interpretation, and a dedicated health–insurance Q&A tool, directly to more than 600 million insured users (Ant Group, 2024). This innovative approach – integrate the features directly into its existing ecosystem, has created a highly scalable testing ground for innovative healthcare deliveries.

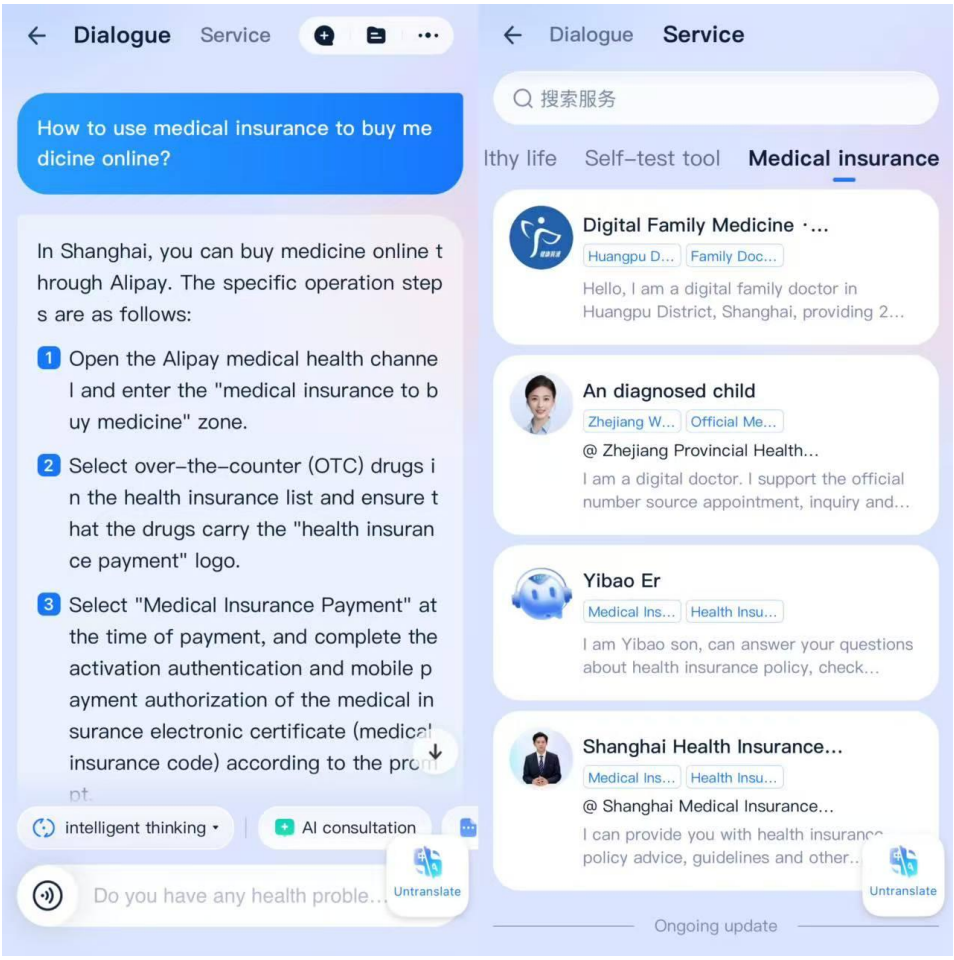
I selected Alipay as an instructive case study of IAs deployment for three key reasons: (1) its extensive user base and established business processes; (2) its strategic partnerships with healthcare institutions and public hospitals; and (3) it was an early mover in rolling out agent-based services in real-world environments. Alipay's intelligent-agent initiatives illustrate all three strategic roles – Creation, Enhancement, and Replacement, offering a window into how large digital platforms can reshape healthcare delivery.

4.1.1.1 Creation: Owning the healthcare "Front Door"

Alipay's documentary materials emphasize the Health Butler's role as a new "gateway" to care (Ant Group, 2024). According to the 2024 press release and the interview with Alipay's executive, it was designed to "shift user behavior from search engines to a conversational AI interface".

The launch of AI Health Butler exemplifies the Creation strategy, introducing a role absent in traditional healthcare models. Alipay's move represents a strategic leap toward building an integrated healthcare ecosystem powered by IAs. Unlike the proliferation of generic medical consultation chatbots flooding the market, which has resulted in largely homogeneous services, Alipay differentiates itself by connecting healthcare offerings with its pre-existing digital payment networks and insurance platform. With access to over six hundred million medical insurance users, Alipay differentiates itself by adding features like a dedicated "health insurance Q&A" tool (see Figure 4).

Figure 4: The interface of AI Health Butler



Source: Alipay application (2025).

This newly created role supports three strategic objectives:

- Short-term user acquisition: The shift in user behavior from traditional search engines to AI-driven interfaces presents an opportunity to seize control of the healthcare “entry point”, making IAs highly competitive.
- Mid-term platform development: If successful, the launch of AI Health Butler could evolve into an AI-era equivalent of an "Application store", fostering a multi-agent ecosystem where developers and services collaborate to deliver diversified healthcare solutions.
- Long-term exploration of "AI doctors": Current stringent regulations and clinical practice limitations mean that AI won’t be substituting physicians anytime soon. But the concept of an AI doctor – combining extensive medical databases with continuous learning algorithms – presents a viable vision to address issues such as healthcare shortages and regional disparities in quality care.

4.1.1.2 Enhancement: Empowering Clinicians with a "Digital Twin"

AI Health Butler is a clear example of how IAs can enhance rather than replace existing healthcare roles. Undoubtedly, doctors serve as the "brain" of the healthcare system, the most irreplaceable players. They don't just analyse test results and symptoms to make accurate diagnoses; they also weigh risks and benefits to tailor the best treatment plan for each patient.

One major enhancement is what Alipay calls the "digital twin" concept – an AI that mirrors a doctor's workflow to assist with documentation, routine communication, and decision support. As the executive put it during the interview, the agent's job is to "digest the unstructured mess" so doctors don't have to. For instance, after consultation, the Health Butler can remind patients about medications or upcoming tests, helping maintain continuity of care without adding to a physician's workload.

Beyond individual doctors, the agent also helps close gaps in the healthcare system by connecting online triage tools with offline hospital visits. Instead of making patients start over every time they switch platforms, the AI keeps track of their records and relays them between services.

4.1.1.3 Replacement: The Emergence of the Virtual Family Doctor

Adopting a functional replacement approach, this strategy involves utilizing agents to streamline administrative and operational tasks within healthcare settings. Alipay's Anzhen'er application – a virtual family doctor application co-developed with the the Zhejiang Provincial Health Commission – showcases the replacement role in taking over non-clinical workflows. While the primary focus of AI Health Butler is expanding ecosystem, Anzhen'er is designed with a deliberate focus on replacing human roles in standard workflows. It automates services typically managed by general practitioners, nursing staff, and hospital administrators, including intelligent triage, appointment scheduling, queue management, online payments, and access to medical reports (Chen, 2024). Its strategic value manifests in complementary benefits for both healthcare providers and service recipients:

- Clinical operational impact: The agent reduces the administrative burden on medical staff by automating repetitive, non-clinical tasks. This redesign allows healthcare professionals to redirect their expertise toward complex, high-value patient care, thereby optimizing both human capital deployment and facility throughput.
- Patient experience: The agent reorients service delivery around patient needs by providing accessible, timely, and personalized services. Patients benefit from reduced waiting times, simplified processes, and continuous support throughout the journey.

Alipay's real-world progress and firsthand accounts clearly show three key steps: creating new digital helpers, improving how doctors work, and replacing repetitive tasks. This proves that established tech platforms can make healthcare smarter by adding agent-based assistants where none existed before; upgrading medical professionals' routines; and taking over tedious paperwork so humans focus on patients.

To further interpret the effectiveness of Alipay's agent deployment, I also draw upon Roger's DOI theory. By situating the five attributes of innovation adoption: relative advantage, compatibility, complexity, trialability, and observability atop this creation-enhancement-replacement framework, it becomes possible to clarify why the adoption has gained traction across the ecosystem.

Its relative advantages are mainly in terms of boosting efficiency (e.g., automating administrative workflows) and democratizing healthcare services with the help of agents. Compatibility is ensured by deeply integrating into China's public healthcare infrastructure and respecting the role of doctors, avoiding disruptive clinical overhauls. To reduce perceived complexity, Alipay starts with low-risk applications (e.g., administrative tasks, simpler specialties) and user-friendly interfaces. Trialability is embedded via non-clinical features (e.g., Health Q&A) and back-office pilots, allowing users and institutions to test AI incrementally. Finally, observability is achieved through tangible, end-to-end services (e.g., AR hospital navigation, report analysis) and quantifiable outcomes (e.g., 10M+ user visits), making AI's value concrete and accelerating trust.

4.1.1.4 Challenges and strategic considerations

Despite its strengths, Alipay's deployment of the AI Health Butler faces several hurdles, as highlighted in the interviews:

- AI Literacy Gaps: Many users "don't know what to ask or where to go," which hampers adoption even with an intuitive interface. Addressing this requires ongoing user education and guided onboarding efforts.
- Data Sourcing & Privacy Constraints: Alipay "doesn't store or own a lot of medical data" and must rely on anonymized, partner-provided datasets. This limits its ability to continuously retrain models on real-world data without complex data-sharing agreements.
- Ecosystem Coordination: With "over 3,600 public hospitals" integrated, aligning disparate IT standards and workflows remains time-consuming. The initial agent-catalog page was even "a bit unclear," underscoring the need for clearer partner onboarding and standardized APIs.
- Monetization Uncertainty: Alipay's indirect, ecosystem-bundled model eschews direct subscriptions because "doctors and patients won't pay" for standalone AI tools. While

this lowers user friction, it also creates opaque ROI metrics for partners, who still demand clear cost-benefit evidence before investing.

While Alipay's early innovations are groundbreaking, future success depends on moving beyond initial market capture to addressing systemic gaps – a challenge echoed by global initiatives such as the WHO's "Digital Health for All" program (WHO, 2021). As ethical considerations and regulatory oversight intensify within the worldwide healthcare landscape, sustainable implementations must demonstrate both technological sophistication and robust governance frameworks. This dual requirement positions IAs to evolve beyond their current experimental status, potentially becoming institutionalized components that support more equitable care delivery systems.

4.1.2 The relevant findings of Baichuan's deployment of IAs

Founded in April 2023, Baichuan AI is a leading Chinese tech company specializing in AI solutions for healthcare institutions. It's core offerings include diagnostic AI for medical imaging analysis, intelligent Electronic Medical Record (EMR) processing, and predictive analytics for operations. Unlike consumer-facing health tech platforms, the primary target customers of Baichuan are hospitals and healthcare providers, with its AI models certified under China's stringent medical device regulations.

One of Baichuan's most notable initiatives is the partnership with Beijing Children's Hospital to launch the first IAs for pediatric – the "AI Pediatrician". This groundbreaking AI pediatrician is China's first clinical assistant capable of dual-mode support – serving both as a research collaborator and frontline clinical partner (China Daily, 2025). The agent is trained on over 40,000 clinical guidelines and 38 million research articles (Wang et al., 2025). In a diagnostic accuracy trial, the AI Pediatrician agent achieved 82% accuracy – 4 percentage points higher than the human physicians' average. While not yet surpassing elite professionals in diagnostic quality, the AI Pediatrician has proven more effective than many existing digital health services, offering a strong proof-of-concept for intelligent clinical agents.

The empirical findings from the Baichuan interview align closely with the strategic roles – Creation, Enhancement, and Replacement.

4.1.2.1 Creation: A new intelligent interface for pediatric care

The introduction of Baichuan AI Pediatrician represents a brand-new healthcare delivery service model in China. This innovation goes beyond mere digitalization of traditional workflows, fundamentally redefining healthcare service provision through improved service reach, decision-making efficiency and resource allocation.

- Expanding service radius: Baichuan executive highlighted that the agent was specifically created to address China's critical pediatric care challenges, such as overcrowded urban hospitals and a severe shortage of pediatric specialists in rural regions. In his words, the agent establishes a new frontline in pediatric care. It extends standardized diagnostic capabilities to underserved areas, which enables families to receive early evaluations without immediate travel to tertiary hospitals.
- Enhancing decision-making efficiency: Traditional pediatric consultations are often lengthy due to detailed history-taking, further strain already resource-limited and overwhelmed hospitals. Baichuan's agent incorporates structured interview templates to streamline this process, significantly reducing consultation duration and minimizing subjective variations.
- Optimized resource allocation: In contrast to typical symptom-checkers or chatbot triage tools, the agent was developed jointly with top-tier clinicians and tested against real-world consultation platforms. Its advanced triage algorithm dynamically distributes cases by severity, redirecting mild conditions to primary care and reserving tertiary hospital resources for urgent or complex cases.

4.1.2.2 Enhancement: AI augmentation of healthcare workflows

Though Baichuan's long-term ambitions include developing autonomous AI clinicians, its current strategy still prioritizes augmenting physician capabilities. The agent currently operates in three key supportive roles:

- Conducting standardized patient intake interviews.
- Generating preliminary clinical evaluations.
- Triaging patient cases to the appropriate department prior to appointments.

Empirical validation of this positioning comes from its co-developer Beijing Children's Hospital, where the AI Pediatrician serves as a first-pass evaluator, with outputs integrated directly into the hospital's clinician dashboards (China Daily, 2025). This integration streamlines workflows from two aspects: (1) minimizing the mismatch between patient needs and specialist resources; and (2) improving diagnostic efficiency and accuracy through better pre-visit data collection.

Critically, Baichuan acknowledges clinicians' concerns about the opaque nature of AI algorithms, which often referred to as the "black box" problem. As I learnt from the interview, Baichuan added a "Doctor Asks Why" button to enhance the transparency of AI-assisted diagnostics. When clinicians question AI-generated diagnoses or treatment recommendations, this function can instantly reveal the algorithm's clinical reasoning path. In doing so, it helps to demystify the AI process and foster greater trust between doctors and patients. Moreover, the agent can generate medical documentation drafts and treatment suggestion templates, supporting clinicians in both decision-making and administrative tasks.

Baichuan's AI enhances care delivery without replacing human judgment, reinforcing the vision of AI as a professional copilot rather than a solo practitioner.

4.1.2.3 Replacement: Delegating low-complexity cases to agents

Although full replacement of physicians is not yet feasible at present, Baichuan is actively targeting high-volume, low-complexity cases for partial automation. This directly points to the core pain point of China's healthcare system. As executives note, the ultimate bottleneck in Chinese healthcare isn't access to information, but a shortage of qualified manpower. The problem becomes painfully clear during the flu season, when parents often face agonizingly long waits at overcrowded children's hospitals. As one hospital administrator noted, 'In many rural areas, we're still struggling with basic pediatric care capabilities.'

A scalable AI diagnostic system designed to fill this critical manpower gap is Baichuan's solution to the problem. As highlighted in the interview, the company envisions future AI agents handling the majority of routine pediatric consultations, freeing up human doctors for complex or emergent cases. A household version of an AI pediatric assistant could serve as a "24/7 on-call consultant". When a child feels unwell, preliminary screening could help determine whether home observation or immediate medical attention is needed. To achieve this, Baichuan is prioritizing three things: pursuing China Food and Drug Administration (CFDA) certification, integrating with public insurance systems, and cooperating with local public health commissions for pilot projects. The CIO noted an intriguing observation from their trial data: the AI Pediatrician demonstrated diagnostic accuracy exceeding that of mid-level doctors, yet at significantly lower costs compared to traditional outpatient services. This is especially appealing for budget-limited hospitals and strained provincial health bureaus. Consequently, Baichuan's strategy emphasizes intelligent delegation, assigning routine diagnostic tasks to AI, allowing clinicians to concentrate on exceptional and complex cases.

To assess how Baichuan's approach facilitates adoption within clinical and public health environments, it is also valuable to map the agent's deployment against the five core DOI attributes.

Its relative advantage is reflected in the use of AI-driven early screening and triage, which effectively address the shortage of qualified pediatricians. Strategic collaborations with top healthcare facilities such as Beijing Children's Hospital, and alignment with national health priorities, have demonstrated compatibility by ensuring that the technology can support, rather than disrupt, existing clinical workflows. To reduce complexity, Baichuan has isolated tasks into manageable modules, starting with low-risk scenarios such as symptom screening and management coordination. Trialability is achieved through phased deployment in public health settings, allowing stakeholders to assess its effectiveness before scaling up. Among the measurable outcomes, observability is evident, including the number of pediatric cases managed, user feedback from clinical partners, and significant improvements in early

detection and patient flow – all of which provide the foundation to build trust and encourage further adoption.

4.1.2.4 Challenges and strategic considerations

Even though Baichuan has made its progress in the healthcare vertical market, its development path is still full of challenges:

- Regulatory constraints: Unlike software platforms that can scale rapidly, medical AI products require CFDA certification and continuous post-market validation. Baichuan's iterative model development must align with a slower, risk-averse regulatory pace.
- Data interoperability: While Baichuan works closely with select hospitals, broader adoption demands harmonization with fragmented EMR systems and regionally siloed medical data.
- Professional acceptance: While doctors are open to pilot collaborations, some remain skeptical about AI's clinical decision-making transparency. Trust-building remains an ongoing priority.
- Business model complexity: Baichuan relies on government partnerships and insurance integration for monetization. This multi-stakeholder funding model introduces administrative friction and limits pricing flexibility.

Still, Baichuan holds a unique edge, as the core believe of the company Chief Executive Officer (CEO), Baichuan is built into the very plumbing of China's healthcare system instead of just another AI company. While consumer apps chase patient subscriptions, Baichuan plays a different game: its currency isn't clicks, but real-world efficiency gains for hospitals and health officials. Success means proving it can move the needle where it matters – delivering measurable operational value to healthcare institutions and regulators. Baichuan's agent deployment pathway shows how to roll out AI in healthcare without tripping over red tape. They turn pilot projects into tools hospitals actually rely on.

4.1.3 The relevant findings of public hospital adoption of IAs

Public hospitals, traditionally regarded as risk-averse institutions within China's healthcare system, are now racing to adopt IAs as assistants. A powerful mix of government policies, overworked staff, and increasingly capable AI is transforming these cautious institutions into key testing grounds for China's healthcare revolution (Olcott, 2023). Based on an in-depth interview with the Director of Informatics at a leading tertiary public hospital, we explore how IAs are being integrated into hospital operations, with a specific focus on their evolving roles across Creation, Enhancement, and Replacement.

This case study was selected for three primary reasons: (1) the hospital's early experimentation with AI in both administrative and clinical settings; (2) its leadership in

piloting AI deployment under strict public-sector constraints; and (3) the candid reflections it offers on the socio-technical complexity of implementing intelligent systems in real-world clinical environments.

4.1.3.1 Creation: IAs as care coordinators and emotional support tools

While AI has yet to replace clinicians in diagnostic or treatment scenarios, public hospitals are actively introducing agents into roles previously unimagined in traditional hospital workflows. The most notable example is the conceptualization of the *"care planner"* role – an IA designed to assist in triaging patients based on urgency, geography, and cost-effectiveness, then routing them to appropriate care levels. In the words of the informatics director: "Most people don't need a top specialist; they need a well-informed, local provider… an AI assistant could help align patients with the right care at the right time."

This role creation extends beyond logistics. Hospitals are exploring the use of AI agents in addressing unmet needs in mental health and emotional support, especially for patients who require consistent engagement but not necessarily psychiatric intervention. With traditional mental health services stretched thin, AI agents are being framed as *"compassionate listeners"* – low-barrier support tools that can stabilize emotional distress before it escalates.

4.1.3.2 Enhancement: Documentation automation and efficiency

The most common and widely accepted application of IAs in public hospitals involves the enhancement of existing clinician workflows. IAs are increasingly being used to:

- Auto-draft medical records and discharge summaries
- Support radiological diagnosis with AI-preliminary reads
- Recommend insurance billing codes through Diagnosis-Related Group (DRG) integration
- Assist doctors with triage suggestions in outpatient or app-based consultation systems

The hospital's informatics team shared an example where a feedback loop was established between an AI imaging model and radiologists, enabling the system to iteratively retrain itself based on expert corrections. This approach, described as *"a digital resident-in-training,"* highlights a broader vision: AI agents not as stand-alone authorities but as learning assistants embedded into clinical workflows. According to the interviewee: "The model keeps learning not just from data, but from real clinical judgment… that's what real AI evolution looks like."

4.1.3.3 Replacement: Selective automation of routine administrative tasks

Though clinical decision-making remains human-led, IAs are already replacing human effort in several administrative domains. The hospital interview outlines early successes in using AI to achieve below targets:

- Auto-read lab results and notify anomalies,
- Manage appointment scheduling and queue notifications,
- Route patients based on their symptoms using digital front-end systems.

Importantly, the hospital's leadership is deliberately prioritizing low-risk, high-frequency tasks as the entry point for replacement efforts. As the director explains: "Start with auto-reading lab reports. Prove the value early before moving into complex areas."

This careful-yet-purposeful approach reflects a clear-eyed view of AI's current limitations and potential. Agents are deployed to reduce bureaucratic overhead and improve throughput, while leaving complex medical decisions exactly where they belong – in human doctors' hand.

Unlike private-sector innovators, public hospitals operate within rigid institutional structures, where the adoption of IAs is shaped as much by organizational inertia and compliance requirements as by technological merit. To better understand how such innovations gain acceptance in this context, it is also insightful to apply DOI theory framework for the hospital case.

Its relative advantage is to improve operational efficiency and free clinical staff from repetitive, non-value-added tasks such as document processing, triage triage, and information retrieval. Compatibility is reflected in the integration of agents into existing hospital workflows and the support (rather than interference) of physician decision-making, which is in line with institutional norms and regulatory expectations. In response to perceived complexity, hospitals have adopted a modular and risk-sensitive implementation strategy, starting with low-risk functions such as record aggregation and queue management. Trialability is achieved through internal pilots in specific departments and administrative units before it is widely adopted. Observability is reinforced through concrete outcomes such as reduced processing time, increased patient throughput, and staff feedback, providing visible proof of value and building internal momentum for continued agent-based transformation.

4.1.3.4 Challenges and strategic considerations

While excitement around AI potentials in healthcare is mounting, the hospital's experience also highlights stubborn roadblocks to agent adoption:

- Infrastructure mismatch: Many departments lack terminals capable of running agent systems.
- Cloud restrictions: Strick data sovereignty rules prevent hospitals from leveraging public cloud processing.
- Shortage of hybrid talent: Most clinicians don't have the time to train AI models or annotate data.
- Fragmented ecosystems: Each hospital builds its own system with different vendors, leading to duplication and inefficiencies.
- Unproven economics: Most AI tools are still viewed as cost centers with unclear ROI.

As the director bluntly notes: "Hospitals can't send data out, but also can't process it in-house. So we're stuck in the middle."

In terms of commercialization, public hospitals are not looking for flashy demos or one-off purchases. They seek co-built ecosystems and long-term partnerships. Their procurement playbook follows strict rules when assessing vendor solutions: (1) No data leaves the premises; (2) Every AI interaction must be logged and auditable; and (3) Total integration costs cannot exceed 15% of Hospital Information System (HIS) budgets. These conservative metrics underscore a crucial point: cutting-edge capabilities mean nothing without trust, safety, and accountability.

The findings from this public hospital demonstrate how IAs can be woven into public-sector care models not through disruption, but through incremental alignment with operational needs and policy requirements. By treating policy constraints as design requirements, the public hospital sector is crafting a uniquely cautious but scalable path forward for AI in healthcare.

4.2 Key findings and Limitations

Analysis of interviews with Alipay, Baichuan Intelligence, and a tertiary public hospital supports the research objectives. These cases represent three distinct sectors – platform ecosystems, private AI companies, and public healthcare institutions, which jointly outline the panorama of IA deployment.

Each entity presents a different operating model and strategic approach: Alipay leverages its ecosystem to achieve service innovation and optimize patient experience; Baichuan embodies the ability of emerging technology companies to quickly iterate and optimize diagnosis; the public hospital reveals compliance challenges in traditional healthcare environments. This variation highlights how organizational context shapes IAs adoption and implementation.

Nevertheless, there are limitations to qualitative research. First, sample representativeness. Although the in-depth analysis of the three typical cases can provide rich qualitative insights, the limited sample size limits the statistical generalizability of the research conclusions. In particular, when extrapolating the adoption rules of a wider range of clinical institutions, it is necessary to take into account that there may be significant differences in the patterns of different levels of hospitals (e.g., secondary hospitals/community medical centres) and institutions of different ownership systems (e.g., private specialized hospitals).

Second, contextual scope. All cases come from the specific context of the Chinese healthcare system. Three unique characteristics of China are noteworthy: (1) the resource allocation model under the hierarchical diagnosis and treatment system, (2) the promotion of digital health policies, and (3) the social and cultural background of the doctor-patient relationship. Findings may not directly translate to other health systems, such as the commercial insurance-led model in the United States or the universal health care system in Europe.

Third, methodological bias. Expert interviews can carry risks of (1) the possible narrative bias of the respondents as representatives of the organization, (2) the interpretation differences in technical terminology, and (3) the conceptual classification bias in the coding process. Although cross-validation with secondary data in analysis is employed, the hermeneutical nature inherent in qualitative research makes complete value neutrality difficult.

These limitations do not diminish the value of this thesis but suggest that follow-up research can be expanded in three directions: (1) expanding the sample coverage and number for quantitative verification, (2) conducting cross-country comparative studies to identify the influencing mechanisms of institutional factors, and (3) using longitudinal tracking methods to capture the dynamic process of technology adoption. This continuous accumulation of knowledge will help to establish a more comprehensive theoretical framework for medical AI applications.

4.3 Future work

Through the analysis of current IAs' adoption in the healthcare sector, we have identified several directions worthy of in-depth exploration. The findings from this research open up opportunities for future research on AI adoption in healthcare, particularly from the following perspectives.

First, there is clear value in expanding the empirical scope. In the future, more different types of healthcare institutions can be included, especially cases from other national health systems. For example, compare the situation in developed and developing countries to see how factors such as the level of digitalization, policies and regulations, and patient needs affect the effectiveness of IAs adoption.

Second, long-term follow-up studies are essential. As healthcare AI applications mature, future studies should continuously assess the sustainable impact on clinical workflows, patient outcomes, and institutional performance. Especially in complex environments such as public hospitals, long-term observation can reveal the interaction between usage habits, technological improvements, and management adjustments.

Third, a variety of research methods should be used in a comprehensive manner. On the basis of the existing case analysis, quantitative methods such as experimental tests can be added to specifically measure user satisfaction, efficiency improvement and other indicators. This allows for a deeper understanding of the specific situation and for more general conclusions.

Additionally, there is a need to carry out economic evaluation research on healthcare IAs adoption. The current deployment of IAs often lacks systematic cost-benefit analysis, resulting in a structural imbalance between technology investment and value return. Future research should construct a full life cycle assessment framework and payment mechanism innovation to compensate for AI services.

It would surely be interesting to also pay attention to standardized management of agent applications. As IAs participate in more and more healthcare decision-making, it is necessary to establish corresponding mechanisms for responsibility identification, risk management and control.

Overall, these follow-up research directions can not only verify the findings of this thesis but also promote the intelligent healthcare from technology implementation to standardized adoption.

5 CONCLUSION

In the healthcare industry, AI has gradually transformed from an initial auxiliary tool to a transformative force in the industry. Driven by the waves of demand, the rapid development of technology and the accumulation of medical data, it plays an increasingly important role in key areas such as early detection of diseases, personalized treatment planning, and optimal allocation of medical resources. IAs are particularly prominent in this regard, providing great help to both clinicians and patients, and as such have been the focus of industry and academia. Agentic applications or systems targeting accessible, efficient, and high-quality care delivery will become the key directions for future development. The agent theory will promote the progress of human-machine collaborative decision-making and intelligent interaction.

In the context of the surge in demand for AI solutions, this thesis conducts an empirical-based multi-dimensional analysis through three in-depth cases: a digital platform (Alipay), a specialized AI company (Baichuan Intelligence), and a tertiary public hospital. Using the Creation-Enhancement-Replacement framework, the study reveals the differentiated strategic role of IAs: Alipay has built a new digital medical entrance (creation), Baichuan

Intelligent has achieved service upgrading (enhancement) through AI clinical support in underresourced fields, and public hospital cases have reduced the operational load by replacing routine administrative tasks (replacement). These role divisions provide an operational perspective for the functional positioning of IAs in a multi-organizational model.

The integration of empirical and theoretical results in three core findings: first, the successful adoption of IAs depends not only on technical performance, but also on ecosystem maturity, regulatory adaptability and user trust construction; Second, platform companies and AI startups present completely different deployment strategies (end-to-end integration vs vertical specialization), and the difference is rooted in organizational genes and stakeholder interaction patterns. Third, although public institutions are facing structural frictions, they can still optimize the adoption trajectory through phased promotion, innovation in governance mechanisms and policy empowerment.

While acknowledging its limitations in sample size, country focus, and methodology, the conclusions provide a key foundation for subsequent research, which can be extended to cross-country comparative studies, the evaluation of the long-term impact of IAs on healthcare outcomes, or the exploration of the evolutionary mechanism of human-robot collaboration in clinical practice.

LIST OF KEY REFERENCES

1. Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. Future Healthcare Journal, 6(2), 94–98. https://doi.org/10.1123/jpah.6.6.805
2. Isern, D., & Moreno, A. (2016). A systematic literature review of agents applied in healthcare. *Journal of Medical Systems, 40*, 1-14. https://doi.org/10.1007/s10916-015-0376-2
3. Jeyaraman, M., Balaji, S., Jeyaraman, N., & Yadav, S. (2023). Unraveling the Ethical Enigma: Artificial Intelligence in Healthcare. *Curēus (Palo Alto, CA)*, *15*(8). https://doi.org/10.7759/cureus.43262
4. Milne-Ives, M., de Cock, C., Lim, E., Shehadeh, M. H., de Pennington, N., Mole, G., Normando, E., & Meinert, E. (2020). The Effectiveness of Artificial Intelligence Conversational Agents in Health Care: Systematic Review. *Journal of Medical Internet Research, 22*(10). https://doi.org/10.2196/20346
5. Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Pearson.
6. Sousa, R., Abelha, V., Peixoto, H., & Machado, J. (2024). Unlocking Healthcare Data Potential: A Comprehensive Integration Approach with GraphQL, openEHR, Redis, and Pervasive Business Intelligence. *Technologies*, *12*(12), 265. https://doi.org/10.3390/technologies12120265
7. World Health Organization. (2021). *Global strategy on digital health 2020–2025*. World Health Organization. https://iris.who.int/handle/10665/344249

REFERENCES

1. Abràmoff, M. D., Roehrenbeck, C., Trujillo, S., Goldstein, J., Graves, A. S., Repka, M. X., & Silva III, E. "Zeke." (2022). A reimbursement framework for artificial intelligence in healthcare. *NPJ Digital Medicine*, *5*(1), 72–72. https://doi.org/10.1038/s41746-022-00621-w
2. Almyranti, M., Sutherland, E., Ash, N., & Eiszele, S. (2024). Artificial Intelligence and the health workforce: Perspectives from medical associations on AI in health. *OECD artificial intelligence papers*. https://doi.org/10.1787/9a31d8af-en
3. Ant Group. (2024, September 4). *Ant Group launches AI-powered mobile app Zhixiaobao at 2024 INCLUSION·Conference on the Bund*. https://www.antgroup.com/en/news-media/press-releases/1725465600000
4. Antonacci, G., Lennox, L., Barlow, J., Evans, L., & Reed, J. (2021). Process mapping in healthcare: a systematic review. *BMC Health Services Research*, *21*(1), 342–342. https://doi.org/10.1186/s12913-021-06254-1
5. Baig, A., Kaplan, J., Lewis, J., & Prieto, P. (2025). *The new economics of enterprise technology in an AI world*. McKinsey Digital. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-new-economics-of-enterprise-technology-in-an-ai-world
6. Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, *17*(1), 99–120. https://doi.org/10.1177/014920639101700108
7. Bjerring, J. C., & Busch, J. (2021). Artificial Intelligence and Patient-Centered Decision-Making. *Philosophy & Technology*, *34*(2), 349–371. https://doi.org/10.1007/s13347-019-00391-6
8. Buchner, F., & Wasem, J. (2006). "Steeping" of health expenditure profiles. *The Geneva Papers on Risk and Insurance-Issues and Practice*, *31*(4), 581-599. https://doi.org/10.1057/palgrave.gpp.2510100
9. Buolamwini, J., & Gebru, T. (2018, January). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research, 81*, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html#:~:text=@InProceedings%7Bpmlr%2Dv81%2D,=%20%7BFriedler%2C%20Sorelle%20A.
10. Cardoso, L., Marins, F., Portela, F., Santos, M., Abelha, A., & Machado, J. (2014). The next generation of interoperability agents in healthcare. *International Journal of Environmental Research and Public Health*, *11*(5), 5349–5371. https://doi.org/10.3390/ijerph110505349
11. Caruso, P. F., Greco, M., Ebm, C., Angelotti, G., & Cecconi, M. (2023). Implementing Artificial Intelligence: Assessing the Cost and Benefits of Algorithmic Decision-Making in Critical Care. *Critical care clinics*, *39*(4), 783–793. https://doi.org/10.1016/j.ccc.2023.03.007
12. Celi, L. A., Cellini, J., Charpignon, M.-L., Dee, E. C., Dernoncourt, F., Eber, R., Mitchell, W. G., Moukheiber, L., Schirmer, J., Situ, J., Paguio, J., Park, J., Wawira, J. G., & Yao,

S. (2022). Sources of bias in artificial intelligence that perpetuate healthcare disparities—A global review. *PLOS Digital Health*, *1*(3). https://doi.org/10.1371/journal.pdig.0000022

13. Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care — addressing ethical challenges. *The New England Journal of Medicine*, *378*(11), 981–983. https://doi.org/10.1056/NEJMp1714229
14. Chen, Y. (2024, September 18). *Zhejiang leads way with AI health assistant*. China Daily. https://global.chinadaily.com.cn/a/202409/18/WS66ea30fda3103711928a85ed.html
15. China Daily. (2025, February 15). *AI pediatrician starts work at Beijing Children's Hospital*. https://www.chinadaily.com.cn/a/202502/15/WS67b051faa310c240449d5734.html
16. Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey Digital. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier
17. Cleland, G. M., Sujan, M. A., Habli, I., & Medhurst, J. (2012). *Evidence: using safety cases in industry and Healthcare.* The Health Foundation.
18. Coiera, E. (2019). The price of artificial intelligence. *Yearbook of Medical Informatics*, *28*(1), 14–15. https://doi.org/10.1055/s-0039-1677892
19. Creswell, J. W., & Poth, C. N. (2016). *Qualitative inquiry and research design: Choosing among five approaches*. Sage publications.
20. Darbandi, M. R., Darbandi, M., Darbandi, S., Bado, I., Hadizadeh, M., & Khorram Khorshid, H. R. (2024). Artificial intelligence breakthroughs in pioneering early diagnosis and precision treatment of breast cancer: A multimethod study. *European Journal of Cancer (1990)*, *209*. https://doi.org/10.1016/j.ejca.2024.114227
21. Davenport, T. H. (1993). *Process innovation: reengineering work through information technology*. Harvard business press.
22. Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. Future Healthcare Journal, 6(2), 94–98. https://doi.org/10.1123/jpah.6.6.805
23. Dearing, J. W., & Cox, J. G. (2018). Diffusion of innovations theory, principles, and practice. *Health Affairs*, *37*(2), 183–190. https://doi.org/10.1377/hlthaff.2017.1104
24. Debbadi, R. K., & Boateng, O. (2025). Enhancing cognitive automation capabilities with reinforcement learning techniques in robotic process automation using UiPath and Automation Anywhere. *Int J Sci Res Arch*, *14*(2), 733-752. https://doi.org/10.30574/ijsra.2025.14.2.0450
25. Delgado-Hernández, X. S., Morales-Rodriguez, M. L., Rangel-Valdez, N., Cruz-Reyes, L., & Castro-Rivera, J. (2020). Development of Conversational Deliberative Agents Driven by Personality via Fuzzy Outranking Relations. *International Journal of Fuzzy Systems*, *22*(8), 2720–2734. https://doi.org/10.1007/s40815-020-00817-w

26. Deloitte. (2023). *Beyond the device: Medtech's expansion into Everything-as-a-Service*. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-beyond-the-device.pdf
27. Elendu, C., Amaechi, D. C., Elendu, T. C., Jingwa, K. A., Okoye, O. K., John Okah, M., Ladele, J. A., Farah, A. H., & Alimi, H. A. (2023). Ethical implications of AI and robotics in healthcare: A review. *Medicine (Baltimore)*, *102*(50). https://doi.org/10.1097/MD.0000000000036671
28. Esmaeilzadeh, P. (2024). Challenges and strategies for wide-scale artificial intelligence (AI) deployment in healthcare practices: A perspective for healthcare organizations. *Artificial Intelligence in Medicine*, *151*. https://doi.org/10.1016/j.artmed.2024.102861
29. Fox, A. (2024, September 11). *Salesforce to launch pre-built AI tools for healthcare*. Healthcare IT News. https://www.healthcareitnews.com/news/salesforce-launch-pre-built-ai-tools-healthcare
30. Fox, J., Beveridge, M., & Glasspool, D. (2003). Understanding intelligent agents: analysis and synthesis. *Ai Communications*, *16*(3), 139–152. https://doi.org/10.3233/EAI-2003-284
31. Fu, C., & Chen, Q. (2025). The future of pharmaceuticals: Artificial intelligence in drug discovery and development. *Journal of Pharmaceutical Analysis*. https://doi.org/10.1016/j.jpha.2025.101248
32. Gartner (2024, July 18). *Gartner identifies key AI trends in 2024 Hype Cycle for artificial intelligence*. Gartner. https://www.gartner.com/en/newsroom/press-releases/2024-08-21-gartner-2024-hype-cycle-for-emerging-technologies-highlights-developer-productivity-total-experience-ai-and-security
33. Google Cloud. (2025). *Google cloud AI business trends 2025*. https://services.google.com/fh/files/misc/google_cloud_ai_trends.pdf
34. Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods, 18*(1), 59–82. https://doi.org/10.1177/1525822X05279903
35. Habli, I., Lawton, T., & Porter, Z. (2020). Artificial intelligence in health care: accountability and safety. *Bulletin of the World Health Organization*, *98*(4), 251–256. https://doi.org/10.2471/BLT.19.237487
36. Haug, C. J., & Drazen, J. M. (2023). Artificial Intelligence and Machine Learning in Clinical Medicine. Reply. *The New England Journal of Medicine*, *388*(25), 2398–2399. https://doi.org/10.1056/NEJMc2305287
37. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, *25*(1). https://doi.org/10.1038/s41591-018-0307-0
38. Huang, K. A., Choudhary, H. K., & Kuo, P. C. (2024). Artificial Intelligent Agent Architecture and Clinical Decision-Making in the Healthcare Sector. *Curēus (Palo Alto, CA)*, *16*(7). https://doi.org/10.7759/cureus.64115

39. Iqbal, S., Altaf, W., Aslam, M., Mahmood, W., & Khan, M. U. G. (2016). Application of intelligent agents in health-care. *Artificial Intelligence Review, 46*, 83-112. https://doi.org/10.1007/s10462-016-9457-y
40. iResearch. (2023). *中国医疗 SaaS 行业研究报告*. iResearch Consulting Group. https://www.scribd.com/document/687864101/H3-AP202112201535703225-1
41. Isern, D., & Moreno, A. (2016). A systematic literature review of agents applied in healthcare. *Journal of Medical Systems, 40*, 1-14. https://doi.org/10.1007/s10916-015-0376-2
42. Jacobides, M. G., Cennamo, C., & Gawer, A. (2018). Towards a theory of ecosystems. *Strategic Management Journal*, *39*(8), 2255–2276. https://doi.org/10.1002/smj.2904
43. Jebbor, S., Raddouane, C., & El Afia, A. (2022). A preliminary study for selecting the appropriate AI-based forecasting model for hospital assets demand under disasters. *Journal of Humanitarian Logistics and Supply Chain Management*, *12*(1), 1–29. https://doi.org/10.1108/JHLSCM-12-2020-0123
44. Jennings, N., & Wooldridge, M. (1996). Software agents. *IEE Review*, *42*(1), 17–20. https://doi.org/10.1049/ir:19960101
45. Jeyaraman, M., Balaji, S., Jeyaraman, N., & Yadav, S. (2023). Unraveling the Ethical Enigma: Artificial Intelligence in Healthcare. *Curēus (Palo Alto, CA)*, *15*(8). https://doi.org/10.7759/cureus.43262
46. Kahn, B. E., & Baron, J. (1995). An Exploratory Study of Choice Rules Favored for High-Stakes Decisions. *Journal of Consumer Psychology*, *4*(4), 305–328. https://doi.org/10.1207/s15327663jcp0404_01
47. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*. https://doi.org/10.1038/s42256-020-0186-1
48. Kapoor, S., Stroebl, B., Siegel, Z. S., Nadgir, N., & Narayanan, A. (2024). *AI agents that matter*. https://doi.org/10.48550/arxiv.2407.01502
49. Khan, N., Reddivari, P., & Ghiya, G. (2025, January 21). *How to use Agentforce in life sciences*. ZS Associates. https://www.zs.com/insights/how-to-use-agentforce-life-sciences
50. Khanna, N. N., Maindarkar, M. A., Viswanathan, V., Fernandes, J. F. E., Paul, S., Bhagawati, M., Ahluwalia, P., Ruzsa, Z., Sharma, A., Kolluri, R., Singh, I. M., Laird, J. R., Fatemi, M., Alizad, A., Saba, L., Agarwal, V., Sharma, A., Teji, J. S., Al-Maini, M., ... Suri, J. S. (2022). Economics of artificial intelligence in healthcare: diagnosis vs. Treatment. *Healthcare*, *10*(12), 2493. https://doi.org/10.3390/healthcare10122493
51. Kuan, R. (2019). Adopting ai in health care will be slow and difficult. *Harvard Business Review Digital Articles*, 2–5. https://hbr.org/2019/10/adopting-ai-in-health-care-will-be-slow-and-difficult
52. Kumova, B. I., & Heye, S. B. (2017). A survey of robotic agent architectures. *2017 International Artificial Intelligence and Data Processing Symposium (IDAP)*, 1–8. https://doi.org/10.1109/IDAP.2017.8090280

53. Lenz, R., & Reichert, M. (2007). IT support for healthcare processes – premises, challenges, perspectives. *Data & Knowledge Engineering*, *61*(1), 39–58. https://doi.org/10.1016/j.datak.2006.04.007
54. Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, *46*(4), 629–650. https://doi.org/10.1093/jcr/ucz013
55. Lukic, V., de Bellefonds, N., & Kropp, M. (2025, May 2). *AI agents can be the new all-stars on your team*. Boston Consulting Group. https://www.bcg.com/publications/2025/how-ai-can-be-the-new-all-star-on-your-team
56. Maleki Varnosfaderani, S., & Forouzanfar, M. (2024). The role of ai in hospitals and clinics: transforming healthcare in the 21st century. *Bioengineering (Basel)*, *11*(4). https://doi.org/10.3390/bioengineering11040337
57. Maresova, P., Javanmardi, E., Barakovic, S., Barakovic Husic, J., Tomsone, S., Krejcar, O., & Kuca, K. (2019). Consequences of chronic diseases and other limitations associated with old age - a scoping review. *BMC Public Health*, *19*(1), 1431–1431. https://doi.org/10.1186/s12889-019-7762-5
58. Mason, C. N., & Miller, T. (2018). International projections of age specific healthcare consumption: 2015–2060. *The Journal of the Economics of Ageing*, *12*, 202–217. https://doi.org/10.1016/j.jeoa.2017.04.003
59. Miller, A. (2013). The future of health care could be elementary with Watson. *CMAJ*, *185*(9), E367–E368. https://doi.org/10.1503/cmaj.109-4442
60. Miller, K., McAdam, M., Spieth, P., & Brady, M. (2021). Business models big and small: Review of conceptualisations and constructs and future directions for SME business model research. *Journal of Business Research*, *131*, 619–626. https://doi.org/10.1016/j.jbusres.2020.12.036
61. Milne-Ives, M., de Cock, C., Lim, E., Shehadeh, M. H., de Pennington, N., Mole, G., Normando, E., & Meinert, E. (2020). The effectiveness of artificial intelligence conversational agents in health care: systematic review. *Journal of Medical Internet Research, 22*(10). https://doi.org/10.2196/20346
62. Minsky, M. L. (1986). *society of mind* (p. 339). Simon and Schuster.
63. Miyashita, M., & Brady, M. (2019). The health care benefits of combining wearables and ai. *Harvard Business Review Digital Articles*, 2–5.
64. Moritz, M., Topol, E., & Rajpurkar, P. (2025). Coordinated AI agents for advancing healthcare. *Nature Biomedical Engineering*, *9*(4), 432–438. https://doi.org/10.1038/s41551-025-01363-2
65. Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From what to how: an initial review of publicly available ai ethics tools, methods and research to translate principles into practices. *Science and Engineering Ethics*, *26*(4), 2141–2168. https://doi.org/10.1007/s11948-019-00165-5
66. Morley, J., Machado, C. C. V., Burr, C., Cowls, J., Joshi, I., Taddeo, M., & Floridi, L. (2020). The ethics of AI in health care: A mapping review. *Social Science & Medicine (1982)*, *260*, 113172–113172. https://doi.org/10.1016/j.socscimed.2020.113172

67. Nahavandi, D., Alizadehsani, R., Khosravi, A., & Acharya, U. R. (2022). Application of artificial intelligence in wearable devices: Opportunities and challenges. *Computer Methods and Programs in Biomedicine*, *213*, 106541–106541. https://doi.org/10.1016/j.cmpb.2021.106541
68. Nambisan, S. (2017). Digital Entrepreneurship: Toward a Digital Technology Perspective of Entrepreneurship. *Entrepreneurship Theory and Practice*, *41*(6), 1029–1055. https://doi.org/10.1111/etap.12254
69. Olcott, E. (2023, October 15). How AI is reshaping China's healthcare system. *Financial Times*. https://www.ft.com/content/c1fe6fbf-8a87-4328-9e75-816009a07a59
70. Parikh, R. B., & Helmchen, L. A. (2022). Paying for artificial intelligence in medicine. *NPJ Digital Medicine*, *5*(1), 63–65. https://doi.org/10.1038/s41746-022-00609-6
71. Purdy, M. (2024, December 12). *What is agentic AI, and how will it change work?* Harvard Business Review. https://hbr.org/2024/12/what-is-agentic-ai-and-how-will-it-change-work
72. Richards, L. (2007). The Discovery of Grounded Theory: Strategies for Qualitative Research. *Qualitative Research Journal*, *6*(2), 183–187. https://doi.org/10.3316/QRJ0602183
73. Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
74. Rony, M. K. K., Akter, K., Debnath, M., Rahman, M. M., Johra, F. tuj, Akter, F., Chandra Das, D., Mondal, S., Mousumi Das, Uddin, M. J., & Rina Parvin, M. (2024). Strengths, weaknesses, opportunities and threats (SWOT) analysis of artificial intelligence adoption in nursing care. *Journal of Medicine, Surgery, and Public Health*, *3*. https://doi.org/10.1016/j.glmedi.2024.100113
75. Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Pearson.
76. Sadeghi, Z., Alizadehsani, R., Cifci, M. A., Kausar, S., Rehman, R., Mahanta, P., ... Pardalos, P. M. (2023). A brief review of explainable artificial intelligence in healthcare. *arXiv preprint arXiv:2304.01543*. https://doi.org/10.48550/arxiv.2304.01543
77. Santoni de Sio, F., & Mecacci, G. (2021). Four Responsibility Gaps with Artificial Intelligence: Why they Matter and How to Address them. *Philosophy & Technology*, *34*(4), 1057–1084. https://doi.org/10.1007/s13347-021-00450-x
78. Schneble, C. O., Elger, B. S., & Shaw, D. M. (2020). Google's Project Nightingale highlights the necessity of data science ethics review. *EMBO Molecular Medicine*, *12*(3), 0–2. https://doi.org/10.15252/emmm.202012053
79. Seyyed-Kalantari, L., Zhang, H., McDermott, M. B. A., Chen, I. Y., & Ghassemi, M. (2021). Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine*, *27*(12), 2176–2182. https://doi.org/10.1038/s41591-021-01595-0
80. Song, H., & Kim, S. H. (2022). How digital transformation can improve hospitals' operational decisions. *Harvard Business School Cases*, 1–2227.
81. Sousa, R., Abelha, V., Peixoto, H., & Machado, J. (2024). Unlocking Healthcare Data Potential: A Comprehensive Integration Approach with graphql, openehr, Redis, and

Pervasive Business Intelligence. *Technologies*, *12*(12), 265. https://doi.org/10.3390/technologies12120265

82. Stokel-Walker, C., & Van Noorden, R. (2023). What chatgpt and generative AI mean for science. *Nature (London)*, *614*(7947), 214–216. https://doi.org/10.1038/d41586-023-00340-6
83. Stryker, C. (2025). *Types of AI agents*. https://www.ibm.com/think/topics/ai-agent-types
84. Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic management journal*, *18*(7), 509-533. https://doi.org/10.1002/(SICI)1097-0266(199708)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z
85. Thesmar, D., Sraer, D., Pinheiro, L., Dadson, N., Veliche, R., & Greenberg, P. (2019). Combining the Power of Artificial Intelligence with the Richness of Healthcare Claims Data: Opportunities and Challenges. *PharmacoEconomics*, *37*(6), 745–752. https://doi.org/10.1007/s40273-019-00777-6
86. U.S. Department of Health and Human Services. (2025). *Strategic plan for the use of artificial intelligence in health, human services, and public health*. https://digitalgovernmenthub.org/examples/u-s-department-of-health-and-human-services-strategic-plan-for-the-use-of-artificial-intelligence-in-health-human-services-and-public-health/
87. Wang, X., Wang, Z., Chen, X., & Liu, C. (2025, March 22). Beijing Children's Hospital launches grassroots version of "AI pediatrician". *China Youth Online*. https://news.cyol.com/gb/articles/2025-03/22/content_6zmdeGij90.html
88. Wolff, J., Pauling, J., Keck, A., & Baumbach, J. (2020). The economic impact of artificial intelligence in health care: Systematic review. *Journal of Medical Internet Research*, *22*(2). https://doi.org/10.2196/16866
89. World Bank. (2023). *Digital-in-health: unlocking the value for everyone*. https://doi.org/10.1596/40212
90. World Health Organization & International Telecommunication Union. (2020). Digital health platform handbook: building a digital information infrastructure (infostructure) for health. World Health Organization. https://www.who.int/publications/i/item/9789240013728
91. World Health Organization. (2021). *Global strategy on digital health 2020–2025*. World Health Organization. https://iris.who.int/handle/10665/344249
92. World Health Organization. (2023). *Global spending on health: Coping with the pandemic*. World Health Organization. https://www.who.int/publications/i/item/9789240086746
93. World Health Organization. (2025). *Ageing and health*. https://www.who.int/news-room/fact-sheets/detail/ageing-and-health
94. Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., …Gui, T. (2025). The rise and potential of large language model-based

agents: A survey. *Science China Information Sciences, 68*(2), 121101. https://doi.org/10.1007/s11432-024-4222-0

95. Xia, M., Poon, L., Choy, M., & Wong, J. (2015). *China's digital health-care revolution.* Retrieved June 3, 2025, from https://www.bcg.com/publications/2015/biopharmaceuticals-medical-devices-technology-chinas-digital-health-care-revolution
96. Yadav, S. P., Mahato, D. P., & Linh, N. T. D. (Eds.). (2020). *Distributed artificial intelligence: A modern approach*. CRC Press.
97. Yang, Y., Peng, Q., Wang, J., Wen, Y., & Zhang, W. (2024). LLM-based Multi-Agent Systems: Techniques and Business Perspectives. https://doi.org/10.48550/arxiv.2411.14033
98. Zamare, R. A. (2024). Elderly Healthcare IoT through Data Analytics and Artificial Intelligence. *2024 IEEE 4th International Conference on ICT in Business Industry & Government (ICTBIG)*, 1–4. https://doi.org/10.1109/ICTBIG64922.2024.10911206

APPENDICES

Appendix 1: Interview Transcript – Participant from Alipay (P1)

Q1: Compared with the changes brought by the Internet to the healthcare industry, what transformative changes do you think AI might bring?

P1: I come from a healthcare services background myself, and this is something I often think about. In my view, the heart of healthcare is the connection between people. Whether it's AI, large language models, or IAs, it's a form of technology. And as technological tools, they have the capacity to take on certain roles that free up human bandwidth. Think of it less as a replacement and more as a powerful support system. Because of this, we believe it has the potential to boost healthcare productivity and reframe how services are delivered, especially within hospitals and clinics. In essence, AI can function as an extension of a doctor or an institution, and in some cases, even operate independently to provide services to patients. That's our fundamental perspective on the role AI could play.

And it's not just within hospitals workflows. AI is already showing up in public-facing areas, such as health Q&A, knowledge sharing, even how people look up medical news. That's reshaping how people perceive healthcare and how they interact with it. So, I see AI as a key factor in both improving public health literacy and enhancing the capabilities of healthcare institutions. Since hospitals still remain the backbone of healthcare service in China, AI can streamline service delivery and management within hospitals, ultimately making the entire operation becomes more efficient and responsive.

Q2: What are the main differences between the era of internet healthcare and the emerging AI-powered healthcare?

P1: We've been doing Internet healthcare for over a decade, and it did a great job creating platforms to connect patients with service providers. But its focus was mainly on "connectivity". With AI, we've shifted to "efficiency". Now, it's about helping healthcare systems work smarter. We believe it can enhance the productivity of the entire industry, from medical imaging interpretation to IA-based services.

Compared to the early days of Internet healthcare or Telehealth, where there was a lot of skepticism – AI is getting more acceptance by the society. Ever since ChatGPT landed, people are a lot more open-minded. The conversation has shifted from "does it work?" to "how do we use it?" From a business development perspective, we no longer need to convince hospital decision-makers about the merits of our products and services. It isn't just digitalizing the workflows; it's changing how we think about care delivery through practical and large-scale applications.

Right now, the reality is that online and offline healthcare services are still pretty disconnected. The question hovering around us is: can AI bridge the chain? Can AI create a feedback loop that strengthens both? I believe this could be a key differentiator in the future. For example, an AI-based online consultation could seamlessly support in-clinic visits by

carrying over information and context. From a broader perspective, we're moving toward hybrid care models where certain routine roles are performed by AI, allowing doctors to focus on their core clinical responsibilities. It's a collaborative model across both online and offline environments, and we envision humans and IAs working together as an ecosystem. Then, the question we should be asking is can we reach this level of integration within the next ten years?

Q3: Since the emergence of ChatGPT and DeepSeek, there's been a growing interest in how IAs are being applied across different industries. Specifically in healthcare, what are some concrete use cases you've seen?

P1: From both Ali's perspective and my own personal observations, we've definitely seen some promising developments, particularly overseas, where the first wave of agent applications in healthcare has taken shape. One of the most stunning strengths is its ability to process large volumes of documents and improve information retrieval. Many international institutions are developing research assistant agents, helping doctors with scientific research and study. In my view, this represents one of the first real, high-impact applications of agents in the medical field.

Building on that, the next major area of development is medical record generation and patient case management. Applications such as designed to support physicians in writing and managing patient records. These tools are evolving into clinical workflow agents, or what we often refer to as "copilots" for the doctors. My take is, this is a major enhancement to both institutional infrastructure and doctor workstations. On the patient side, the applications are relatively more straightforward. We're talking about high-quality conversational agents. In some overseas examples, these agents are designed to better understand patient needs, recommend the right healthcare providers, and integrate the end-to-end service experience offered by hospitals.

So overall, we're seeing two clear directions: enhancing the workflow and capabilities of healthcare professionals and improving the patient experience through more personalized support.

Q4: We've talked a lot about the potential of AI in healthcare. In terms of the real-world adoption, where do you see the most immediate opportunities for impact right now?

P1: Honestly, I think both China and the U.S. are arriving at a similar inflection point. The fastest and least risky entry for AI in healthcare right now is in the administrative domain. That's where you see less clinical resistance and immediate ROI. Things that eat up time like documentation automation, scheduling optimization, or claim processing, are solving real bottlenecks. AI's ability to digest mountains of unstructured medical texts also has made it a great assistant for doctors, such as helping them find clinical trials, review literature, or dig up supporting evidence fast. We're reducing doctors' paperwork burden without getting caught in clinical validation hell. Now, AI handles a lot of that groundwork. For patient-

facing applications, we've seen growth in intelligent Q&A services. Things that let them check symptoms, get medical advice, or even find the right doctor or hospital. But in general, we believe the breakout moment in healthcare starts from the back office and moves forward, not the other way around.

Q5: You don't think AI applications in clinical environment will be the growth driver?

P1: For sure it's promising, but not in the short term. Clinical agents are exciting on paper, but it's a regulatory minefield. Let me tell you something interesting. Earlier last year, the U.S. Department of HHS released a nearly 200-page blueprint titled the "AI Healthcare Strategic Plan," It details how the U.S. government plans to aggressively advance the application of AI across the healthcare industry, painting a picture of AI as a transformative force across the full life cycle. The last time a concept carried this level of strategic importance was probably a decade ago, during the rise of Telehealth. So it's not hard to see why HHS is placing such big bets on AI. Beyond genuine belief in AI's potential to deliver value, it's also clear the current market hype and investor frenzy played a role in making AI healthcare one of HHS's first major bets.

In the plan, "Healthcare Services" is arguably the most crucial business segment. And incidentally, also the most debated area when it comes to AI in healthcare, especially in countries like the U.S. and China. HHS conducted a survey regarding the attitude towards adopting AI in clinical settings. There's still a huge trust gap, both from doctors and from patients. What's far more pragmatic right now is a back-office-to-frontline approach, especially if you're a startup or even a legacy tech vendor looking to scale in healthcare AI. You're more likely to survive in the market by becoming the Oracle or SAP of healthcare first. So, I will say, stop moonshot thinking and start solving the real pain points. And once you earn a footprint in operations, the clinical doors start to crack open. We've seen this pattern before in other industries, and AI in healthcare won't be different just because it is under the spotlight.

Q6: From your experience, what conditions lead doctors to adopt and trust AI agents? Have you seen any turning points?

P1: When we discuss doctor agents with hospitals, the response is generally supportive. Hospitals, especially with well-known doctors, feel responsible for enabling better service. One major challenge for them is appointment scarcity: top doctors' appointments are hard to get, but hospitals lack tools to match patients with the right level of care. Agents can help screen patients more effectively, ensuring the right cases reach the right doctors, benefiting all parties involved.

From the doctors' perspective, one of the most pressing needs is better post-consultation management—follow-ups, medication reminders, and patient education. These tasks are necessary but repetitive, taking up valuable time. AI has strong potential here to automate routine workflows and free doctors to focus on more complex clinical work. Another

motivator for doctors is personal influence. Many specialists with unique methods or successful treatments want their expertise to be recognized and accessible. AI agents can help scale their knowledge while preserving their voice and medical philosophy.

However, a major hurdle we met is resistance to adoption. Despite people generally see value in AI, many of them are reluctant to use systems that might appear to challenge their authority or clinical judgement. To be fair, it's rooted in work culture. No one enjoys tools that constantly point out your errors, even if they're right.

That said, the areas where agents are most welcomed are those that offload tedious, low-value tasks such as data entry, documentation, routine follow-ups. When agents reduce the mental load of repetitive work, doctors can focus more on meaningful care, which improves both patient outcomes and job satisfaction. In large hospitals, some of this burden is handled through support staff, but these resources are costly and usually reserved for senior doctors. In smaller or lower-tier hospitals, or among younger doctors, that support is rare. These physicians often handle everything themselves, which leads to burnout and inefficiencies. For them, AI could be a game-changer.

Q7: Since we've discussed broader topics, let's shift to something more hands-on. Alipay launched the AI Health Butler last year, which has generated quite a bit of buzz. Could you walk me through what the application does and what makes it stand out from other AI health assistants in the market?

P1: We built the Health Butler around a simple and strong belief: to equip each individual with a personal health assistant. That's to say, each Chinese citizen should have a "butler" to look out for their health. Now, this butler doesn't necessarily have to be powered solely by AI – AI might serve as a key component, but the core engine, the interaction layer, or even just the "lubricant" that enhances the experience. The product goes far beyond tech, in fact, we see it as part of a broader ecosystem.

With that vision in mind, we target all demographics across the health journey, whether it's prevention, diagnosis, treatment, or recovery. Particularly for those citizens living in underserved areas who have poor access to credible information, and the lack of continuity in healthcare. The core problem we aim to address is about people's growing attention to healthcare, and the experience and trustworthiness of healthcare processes. We look at this through two primary lenses: medical touchpoints – moments when users are actively seeking care; and non-medical scenarios – such as daily wellness and preventive health management. These are distinct yet complementary, and together they shape our approach.

We define the Butler's functions in three user-centered roles:

Health Q&A assistant: it offers daily health advice and manages personal wellness, when there's no immediate need for a doctor.

Digital care companion: this role comes into place during medical encounters, whether online or in-person hospital visits. It provides end-to-end guidance and support throughout the care journey.

AI digital doctor: This is not meant to replace real physicians, but rather to assist with them for decision-making and support.

We're also building tools for doctors. On the healthcare service provider side, we help them handle patient inquiries, manage workload, and even expand their influence through AI-powered agents. For example, while the assistant agent is accompanying a patient during the visit, it's also collecting the patient's questions or identifying points of confusion. This data can then be relayed to the doctor, helping them understand the patient more efficiently and delivering more personalized care. That's to say, the role of our AI Health Butler is to serve users, support doctors, and integrate intelligently across touchpoints.

What really sets us apart is how we work within an ecosystem. We know by heart that healthcare is complex and deeply specialized. Ali alone cannot solve everything. We want to bring domain expertise into a collaborative ecosystem. That's why we are proactively working with hospitals, doctors, research institutions and companies across the industry to build the platform together. Agent in this context is not just a chatbot, it serves digital extensions of healthcare service providers. And it's a portal to an intelligent, integrated healthcare experience. That's our long-term vision.

Q8: I found some features of the AI Health Butler quite interesting. For instance, the smart triage function wasn't a human assistant, but rather an AR-based navigation tool; there was also the report interpretation feature, where I received a very detailed health analysis. I've learnt that many companies are developing similar applications. What makes Alipay different? What's its unique value proposition in the market?

P1: Sharp question. Let me start with our core positioning, which I believe is central to how we approach AI healthcare. From the beginning, Alipay has been grounded in the idea of inclusion – starting with inclusive finance and now expanding into inclusive healthcare. That means our mission is to make high-quality services accessible to as many people as possible, especially those who may never have engaged with AI before.

We want users' first experience with AI to be dependable and beneficial. Many users might be encountering AI for the first time on Alipay, and we see it as our responsibility not just to provide the service, but also to educate and build trust in the technology. So, in the face of many similar products which might be more innovative, we take a more stable and user-friendly approach. Our goal is to ensure users find value in their first interaction, and through comparison, begin to recognize what sets Alipay apart.

As you mentioned earlier, AR navigation is another example. Many triage tools offer a brief Q&A interaction and then leave it at that. In contrast, Alipay's service extends well beyond

basic interaction. After selecting a department, users can register and book appointments, receive department and doctor recommendations, get guidance on finding available doctors nearby, or be accompanied throughout the in-hospital experience – from check-in to medication pickup and lab testing. So we don't just provide a "chatbot." We offer end-to-end problem-solving, which is where I believe our second competitive advantage is: comprehensiveness and continuity of service.

The third competitive advantage relates to our agent ecosystem strategy. While many companies are exploring AI agents, much of that work remains under the radar. Alipay can use its own technological foundation to amplify these efforts through partnerships. Some companies have already developed highly specialized "Level 5" expert agents for certain medical domains. Alipay can help integrate and enhance these capabilities and bring them to a broader audience; essentially acting as both an enabler and a bridge to users. Another key strength is our deep integration with China's public healthcare system. Over the past decade, we've built connections with over 3,600 public hospitals and a wide range of healthcare institutions across the country. This gives our AI Health Butler a reliable operational foundation, allowing it to plug directly into hospital systems and support full-process experiences like triage, navigation, appointment booking, and payments. For example, in Zhejiang province alone, we've already supported over 10 million user visits. We're also integrated with local medical insurance bureaus, allowing users to get reliable answers to a wide range of insurance-related questions.

Q9: Since you mentioned agent ecosystem strategy, I see some are already on the platform, including partnerships with hospitals and authorities. Do you have a plan for scaling this?

P1: Another good one. As a user, I also found the agent page a bit unclear at first. It might seem like a ranking of widgets based on clicks, but healthcare isn't really about tools in that way. It is only needed at certain moments. So we're working on better planning and clearer categorization of agents by user case.

Institutionally, we aim to include as many partners as possible, such as public hospitals, healthcare commissions, and institutions, so users can access more reliable services. For individual doctors and specialists, we have two goals. First, help people see doctors as real, approachable professionals by using AI to close the gap between them. Second, addressing basic health concerns, such as menstrual tracking or body type insights with engaging tools. These are the two main directions we're pursuing.

Q10: Where does Alipay get its data from? And what's your strategy when it comes to medical data and ensuring AI accuracy?

P1: Alipay acts mainly as a platform in the digital health space. Honestly, we don't store or own a lot of medical data – that's both a limitation and a strength. From day one of building the assistant agent, we started the data system from scratch, because credibility in medical data is critical.

A lot of online health content is still inconsistent in quality, so we focused early on finding authoritative sources from research papers and peer-reviewed literature and partnered with publishers like People's Medical Publishing House co-create trustworthy, anonymized datasets.

We also collaborate with top hospitals and institutions on joint research projects to access non-sensitive clinical practice guidelines, reviewed and validated by real doctors. For traditional online Q&A or consultation data, we apply strict filtering and annotation by our in-house medical team.

Every dataset is evaluated for its impact on model quality – we optimize for positive feedback and cut what doesn't work. Our assistant agent has nearly 10 billion parameters, which all aligned to real clinical standards. For specialty areas, we use a co-development model, working with doctors and industry partners to build tailored diagnostic modules. We started by selecting relatively straightforward specialties – those involving more isolated systems, like the breast, thyroid, or urinary tract. These areas tend to have clearer, more standardized diagnostic and treatment protocols. Ironically, even something that seems simple, like a common cold or cough, can have a wide range of causes and a more complex diagnostic process. That's why we prioritized these simpler specialties for initial development.

Q11: In your view, what qualifies as "high-quality data"?

P1: High-quality data isn't just complete or well-labeled; it needs to capture how a doctor thinks – the process of thinking and decision-making. We want data that reflects the full reasoning process: from symptom interpretation to differential diagnosis to treatment decisions. A note that says "diagnosed pneumonia" is just an answer key. But if it includes how the doctor ruled out other conditions and why they chose a specific treatment, that's valuable training material. Agents learn better when the data teaches them to think, not just memorize. Unfortunately, much of today's clinical data is too compressed to capture that kind of logic.

Q12: Some people worry that users may not know how to effectively use AI tools, especially for health questions. Is that something you've run into?

P1: Definitely. This touches on a dilemma in healthcare, especially for consumer-facing services. Health isn't something most people interact with daily – until they need it. We have conducted a survey among the users, and it turned out that the average individual rarely seeks out health information proactively when they are well. As a result, most users lack the daily interaction or training needed to build effective search habits or health information literacy. In short, they often don't know what to ask or where to go. And that is why so many people fall into traps with ads or misinformation.

AI-powered chat interfaces are more intuitive, such as those inspired by ChatGPT, and people are more willing to use them. But many still struggle to describe their symptoms clearly. That's not a tech problem, it's a user education gap, or what we call gap in AI literacy. In my view, the issue is rooted in the asymmetry of healthcare information. For someone without a medical background, articulating symptoms accurately can be difficult unless the case is extremely straightforward. Therefore, even if we develop sophisticated AI systems on the supply side, their effectiveness may be limited unless we address usability on the demand side.

That said, even the best AI needs to be matched with smart design and good onboarding. We need to develop an assistant agent capable of natural, human-like dialogue. But as of now, most AI tools still require structured input. We recognize this as a key challenge, and we are working hard on making our tools feel natural and helpful, not just clever.

Q13: You mentioned Q&A-based products like the AI health butler and the concept of an AI-powered doctor. From what I've seen in the market, it feels like there're two distinct approaches rising. One aims AI functioning as an assistive tool, the other aims for a fully autonomous AI doctor. Where does Alipay stand?

P1: The core philosophy we've consistently upheld is a collaboration between a real doctor and an AI doctor that can achieve far greater results than even three human doctors working separately. Solving complex healthcare challenges requires collective intelligence and cooperation, that's why we fully respect and recognize the irreplaceable value of physicians.

We don't believe healthcare can be fully standardized or reduced to pure data-driven experience, at least at this stage of AI development. Every medical case is shaped by a mix of personalized diagnostic strategies, physicians' varying levels of expertise, local epidemiological trends, habitual practice patterns, and even emotional nuances. That's not something AI capabilities can replicate. Therefore, we see the best results when real doctors are supported by IAs. Incorporating AI as a supporter in the existing healthcare service chain is our primary goal. Think of it like going from Level 1 to Level 3 in autonomous driving – we're improving assistance, not handing over the wheel. What we want to achieve is to help doctors be more effective by assistive agents, especially with repetitive tasks and expand access to care without compromising quality.

Q14: Let's shift to business models. OpenAI offers a paid subscription to both consumers and enterprises. Would that model work for healthcare AI? What's your commercialization approach?

P1: OpenAI's subscription model works well for general productivity tools, but healthcare AI operates in a very different context. At least we don't think subscription models are viable in China's environment.

Right now, we're not selling AI as a standalone product. Our primary goal is to improve medical efficiency, not to create a product that individual doctors or patients would directly pay for—because realistically, they won't. Doctors already face heavy workloads and limited incentives for adopting new tools, and patients are accustomed to public healthcare services at low or no cost.

So we focus on indirect monetization. Our approach is to bundle AI into products and services. For instance, we integrate smart diagnostic agents and digital transformation services into hospital settings. These are part of the digital transformation packages we offer to public institutions. On the consumer side, we use AI as an engine behind broader services, integrating it into our ecosystem of online consultations, drug delivery, and health management. We also see value in precision matching – connecting doctors to the right patients for clinical trials or specialist care. In these cases, funding may come from research institutions or healthcare partners, not from users. To be honest, we're still in the early stages. Commercialization isn't our current focus. But if our ecosystem partners see real value in the project, they become our indirect clients. Long term, our business model is built around platform integration, not subscriptions.

Q15: Looking ahead, where do you see potential for agents in healthcare?

P1: Predicting the future is difficult, let alone the healthcare landscape. China's system is led by public institutions, but individual doctors still hold influences. We're open-minded about where AI will go, which depends on many evolving factors, including industry shifts and how AI is adopted.

Our focus is identifying the most effective care pathways for users. While deeper or serious treatments such as medication and surgeries still require offline care, there's a flexible space where AI can play a bigger role. Our approach is to capture and apply doctors' expertise to users, then build the right tools around it. For example, we partnered with Renji Hospital's urology team to co-build a specialty knowledge base and model, making the AI more precise in consultations and report analysis. Long-term, we hope to bring this expertise to primary care, helping standardize specialty treatment.

In short, the key is solving the right problems; the format can be shaped around that.

Q16: This has been an insightful conversation. Any final thoughts on what agents really mean for the future of care?

P1: In our view, agents in healthcare is not to replace the human touch, but to make it more effective. From medical Q&A to specialist referrals, from imaging analysis to report generation, IAs should let us build healthcare systems that are smarter, faster, and more patient centered. At Ali, we are focused on real partnerships, real data, and real-world impact. As regulations mature and ecosystems deepen, IAs may become the digital companions of tomorrow's care.

Appendix 2: Interview Transcript – Participant from Baichuan (P2)

Q1: What has Baichuan been focusing on recently?

P2: Over the past year, we have launched several generations of large models in multimodal learning, healthcare performance, and Chinese language capabilities. This year, we've been busy with building our first full-scenario inference model – Baichuan-M1-preview. We will also release an open-sourced multimodal model, Omini-1.5. We aim for both general capabilities and healthcare enhancement. Every generation we release must strengthen our healthcare AI.

Q2: Compared to other verticals, why has Baichuan chosen to prioritize healthcare as a key focus area? And why didn't you emphasize healthcare at the beginning?

P2: Two words: doctor shortages. Healthcare's bottleneck isn't information matching; it's the shortage of doctors. The internet healthcare era tried to solve this issue by optimizing scheduling, remote consultations, and resource allocation. But it didn't address the root issue: the total volume of available expertise. That's why the arrival of AI and their potential for general intelligence has reignited industry excitement.

If you are really interested in these problems, you have to know that there are currently two prevailing strategies in the AI healthcare landscape. The first, aims to build fully autonomous "AI doctors" like what Baichuan is focusing on – akin to Level 4 autonomous vehicles. The second, more conservative path involves creating smart copilots for human doctors – more like Level 1 or 2 autonomy. Although most companies shy away from overtly discussing the "AI doctor" concept due to regulatory sensitivities, it remains a central ambition. This is especially relevant in countries like China, where high-quality medical and healthcare resources remain scarce and unevenly distributed. Policymakers have long emphasized the need to expand and decentralize care where AI could be a powerful tool in achieving that. This year, especially in the second half, we've been intensively expanding into healthcare, fulfilling our vision of super-applications. We brought on a wide array of professionals from the medical field and began signaling our pivot – moving from a team "rich in models" to one "rich in medical expertise."

And to answer your last question – because back then, everyone's attention was on OpenAI. When communicating with investors and teams, overemphasizing healthcare might have led them to think it's just another niche market. Also, healthcare apps only blow up after general models mature. So, we built the basic models first. Actually, we've already started working on products in partnership with hospitals and government bodies, such as regulators and health commissions. This year has been about reconnecting with original passion. Now, we are working to go from enthusiastic amateur to a recognized professional in healthcare.

Q3: Let's get into what Baichuan's actually accomplished. Can you describe Baichuan's current product, particularly the newly launched AI pediatrician?

P2: In September 2024, we partnered with Beijing Children's Hospital to build AI pediatrician. This is a smart wedge. Parents know kids get sick constantly. And it's stressful; you have to take time off work, go to crowded hospitals, risk cross-infection, and face long waits. But most cases aren't complex or severe. Therefore, pediatrics is a strategic entry point because it addresses a massive pain point, and it lets us build from the ground up.

Initially, we wanted the tech to speak for itself. But in the end, positioning came first, then tech followed. We've rolled out what we call "one major, four minor" systems with Beijing Children's Hospital: four scenario-based models – home, community, hospital, and intelligent monitoring. In testing, we evaluated 11 professional medical questions across several online platforms where you can "pay to ask a doctor." At the cheap tier, accuracy was around 30%. For mid-tier doctors, accuracy climbed to about 70%. Our AI doctors slightly beat that, but they're not as accurate as more top-tier doctors, such as Bei'er Doctor, which can achieve 90% accuracy with an online service of nearly 200 RMB. So, Baichuan's AI doctor isn't as good as elite professionals, but it's better than most platforms. That's where we stand today.

Q4: How do you balance rapid innovation in healthcare AI with strict regulatory requirements?

P2: One axis is professional level: chief, deputy chief, etc. Another is price tier across platforms. Eventually, we'll categorize by condition, specialty, and pricing. The government and healthcare institutions are embracing this kind of model training.

In fact, our governments are desperate for solutions. Public hospitals are actively seeking collaboration with us. Of course, navigating regulatory hurdles requires quite a lot of effort. But overall, the momentum is positive. We are building pilots in key regions to integrate AI family doctors deeply into community. At the same time, real doctors are curious about how AI evaluates and grades them. This raises a bunch of new policy questions. For example, with traditional medical devices, if you change a part, you need months of re-approval. That can't scale with AI, which improves every day. If you freeze AI in place, it becomes useless. We'll have to rethink how doctors are evaluated, licensed, and held accountable. There's a ton of groundwork ahead.

Q5: What made Beijing Children's Hospital jump on board? How do you ensure the AI pediatrician is accurate and trustworthy for the diagnoses?

P2: Good questions. The hospital's leadership was visionary. Some institutions are still hesitant about AI doctors, but Beijing Children's wanted to "train a million AI pediatricians." This perfectly aligns with our urgent national needs. China is in a crazy shortage of pediatricians. Kids pack into hospitals, risk infections and overburden families. The future will be bringing care home. Early prevention, early diagnosis, and full lifecycle disease management at home, taking the pressure off brick-and-mortar hospitals.

And regarding how we ensure reliable diagnostics from the AI pediatrician, here are the steps. First, we train models using both public medical data and expert knowledge from partners like how Beijing Children's Hospital's top doctors are basically teaching our AI by turning their diagnostic thought processes into training data. It's like creating textbooks for AI to study. And we've jointly established a key laboratory with Beijing's healthcare authorities, where clinical data will be shared through the lab.

Second, we constantly update models with the latest research papers to make sure our agent stays on top of all the newest medical breakthroughs. Lastly, we take all that knowledge and build actual tools such as helping parents decide if their kid's fever needs an ER visit or flagging easy-to-miss symptoms of rare diseases for doctors. Plus, doctors get a "Why?" button – agents explain every decision, like "Here's why we flagged that tumor."

The next step is to get certification for our AI pediatrician – passing CFDA, achieving consistent performance, and getting licensed. We're piloting with a mini-app, and we aim to release full professional applications next year. But unlike the internet world, where you can launch fast and go viral, healthcare demands regulatory approval. We're dealing with health economists, policy compliance, and patient safety. This isn't "move fast and break things." This is "move carefully and prove everything."

Q6: What key shifts are needed in healthcare to make personalized AI doctors widely adopted?

P2: There are two core shifts to keep in mind. First, from hospital-centered to home-centered care. Think about how people used to bathe in communal bathhouses. Now every home has its own bathroom. That's the vision here: healthcare in your own space. And it might be where we see real breakthroughs in the near future. Second, from doctor-centered to patient-centered systems. If an AI doctor is doing your online consultation, most people can accept that. The mental shift is already happening. So how do we actually create AI doctors, and change how they work? We've spent a year exploring this. One thing we found is that AI can't operate in isolation; it still has to coordinate with in-hospital systems. Right now, going to the hospital is like opening a mystery box. You don't know what doctor you'll get. But with AI pre-screening, asking you detailed questions in advance, it can even assign you to the right physician ahead of time. The doctor doesn't have to guess, and you don't have to gamble. Everyone wins. Moving from hospital to home isn't a straight line. There's a long transition. And along the way, we have to constantly align our ideals with reality.

Q7: Despite significant investments in healthcare reform, there're still many unresolved fundamental issues. What's Baichuan's strategy in addressing these persistent healthcare challenges?

P2: We spent quite a lot of money before, and investors are still recovering emotionally. And that doesn't even count pharmaceutical R&D. Most of that money went toward internet healthcare: helping doctors make better decisions, AI-assisted medical imaging, and so on.

But here's the thing: all that effort didn't solve the supply problem. Healthcare isn't like ride-hailing or food delivery. You can't just throw more "drivers" into the system. Before AI, we simply didn't have the tools to scale the number of doctors. Now, with AI agents, we can create doctors. This is a brand-new starting point. For Baichuan, it might still take another 100 billion. But this time, the money won't go into digitizing workflows. It'll go into building the AI brain itself – training models, supporting medical researchers, and developing the tech backbone that can actually increase supply. This era gives us the opportunity to reinvent the supply side of healthcare from the ground up by using large models as the foundation.

Q8: What are the monetization strategies of Baichuan? Charging families, hospitals, or the government?

P2: Honestly, all three. For Government-end: Public health budgets will fund AI-assisted family doctor plans. Hospitals-end: Reimbursement through national medical insurance programs. Consumers-end: Down the line, maybe private insurance bundles. But we are still figuring how to do it.

Most companies couldn't touch core healthcare problems, they just build bandaids – apps for scheduling, billing, whatever. But we are creating AI doctors, not tools. Our agent model is the core value. In the beginning, deployments might be a little heavy, but once we finish pilots with the first few clients, we can outsource non-core deployment tasks to independent software vendors.

Q9: How will Baichuan approach fundraising? Is it easier or harder to raise money for AI healthcare?

P2: It's a common question people, particularly the investor, are asking. Healthcare doesn't bring in huge revenue. Who's going to pay for this? Just few weeks ago, I saw a news article about China's medical devices and how drug prices are being forced lower. But the good news is AI services are now being welcomed into public insurance. Previously, you had drugs, devices, and doctors. Now, even AI-powered image analysis can be reimbursed. That cracks open the door. AI doesn't have to be hidden inside a machine or a pill; it can be a service. More importantly, you need to have something new to tell. Our team realized that without a clear application focus, competition in general models was turning into a bloody red ocean. This year, we're leaning in. We're saying clearly: Baichuan isn't just a large model company; we are the leading healthcare large model startup in China. Next year, I think we'll have a whole new story to tell, one with massive market potential.

Q10: How do you navigate the tension between investor expectations and staying true to your long-term vision?

P2: From day one, when we started working on AI, we had firm faith that healthcare is sexier than large models. But we got pushback from investors: "No one's going to buy that. You're

not a doctor. You're not an academic. You can't just walk in and talk healthcare. You've got to talk about AGI." – which was fair. AGI is a legit pathway to the future; but they need to be top tier to power healthcare.

Last year, we told a "large model" story on the strategic meeting with investors. Healthcare was just one theme among others – "Health, Creativity, and Joy." That was the pitch. So, with each model release, we barely mentioned healthcare. And we still got investor recognition.

Q11: In a fast-moving field where many companies rush to match every new breakthrough, like OpenAI is launching video AGI and multimodal models every five minutes. Baichuan has taken a more deliberate path. What drives your decision not to follow industry trends?

P2: First, let me ask you a question: what's the point of "chasing"? To prove you're not falling behind? That's a treadmill. From day one, Baichuan's goal was to be in the first tier of model companies, and to make breakthroughs in healthcare via super-apps. So, while others were playing catch-up, we stayed focused.

There are two traps here: Frist is that you need a clear technical roadmap. I've said this over and over, that we focus on language intelligence. When Sora launched, companies that were supposedly language-first suddenly pivoted to video. But Sora and language models are fundamentally different tracks. Therefore, we skipped it. We're not copying ChatGPT just to copy it. There's no need to build a replica three months later. Be idealistically one step behind, but practically three steps ahead. Meaning, we don't chase blindly; we make smart choices.

Another one is beware of the application trap. Big companies dominate super-apps. ByteDance, for example, is unbeatable here. So, if you can't out-tech OpenAI and can't out-app ByteDance, you get crushed from both sides. That's why we picked a lane outside their artillery range. We're not playing their game. Giants do "one-size-fits-all" AI. We're hyper-local, like integrating with China's health records or targeting regional diseases. Our agents are lightweight, so hospitals don't need fancy cloud setups. And we're built for China's privacy laws. Try that, Silicon Valley. For example, just a few days ago, a stroke patient at Inner Mongolia People's Hospital was given a critical condition notice. His family brought him to Beijing Union Hospital and sent his records to us. Our model M1 generated three diagnostic suggestions, matching the four proposed later by expert consultation at Union Hospital. M1's ability has surpassed city-level hospitals.

Q12: So, the strategy is to figure out whose artillery can't reach Baichuan?

P2: Exactly. For example, people said, "You should be first to do AI search!" Sure, it could beat Baidu (a search engine company like Google) but you're still in ByteDance's firing line.

At Baichuan, we asked: In big tech's territory (daily active users), we're weak; in niche needs, small startups often win; but in healthcare, neither big companies nor small startups have a clear edge. It's considered "vertical," but the barrier to entry is high. Even giants like Baidu and Tencent have scaled back healthcare investment. Baidu just downgraded its medical division. Innovation demands more than what professional managers can offer. That gives us room to breathe and innovate.

Plus, healthcare isn't just one "vertical." Think of it like an evolutionary tree – some branches stop growing. Ad copywriting? Cool, easy money, but it doesn't require cutting-edge AI. It doesn't evolve. Healthcare, on the other hand, pushes us toward multimodal agents, analysing pathology slides, communicating empathetically, reasoning logically, handling long-term memory. It demands top tier AI and keeps pushing us forward. So no, we're not here for short-term monetization. We're building toward a new "era of life" after the "era of intelligence." Healthcare is our launchpad, but the horizon is vast.

Q13: Among Chinese tech companies, who do you think will survive, and who might collapse?

P2: Those who continue to burn cash without clear applications and markets will struggle. There are already companies that no longer need heavy financing because they've secured their niche. Also, companies trying to compete head-to-head with tech giants in general applications are at major risk.

Q14: Alipay recently released a similar product calls AI Health Butler, aiming to expand in AI healthcare services. Does that concern you?

P2: Their focus seems more on payments and business model innovation, rather than truly advancing AI capabilities or solving the doctor supply problem.

Q15: What do you see as the future direction for I, and what are the biggest opportunities and challenges ahead?

P2: 2025 will be the year of the AI Agent, and we'll see the rise of entirely new digital "species."

The longstanding issue of data silos in healthcare may finally meet its match through the rise of multimodal agents. In rare disease diagnosis, for example, an agent capable of synthesizing imaging, pathology, and genetic data could offer interdisciplinary insights, leaving the final decision to the physician. This shift from siloed systems to collaborative, AI-powered synthesis is one of the clearest signs of where agents are headed.

Another positive shift comes from the interaction models. Most current medical AI systems only provide passive suggestions. Future agents may demonstrate "executive intelligence", such as actively controlling equipment, adjusting parameters, and becoming embedded

participants in the clinical workflow. Of course, this vision faces significant barriers (data privacy, safety, trust, and regulation), but the direction is clear: from passive advisors to active collaborators.

Regarding challenges, while the success of IAs depends on on usable, user-centered design. That's where the next wave of opportunity lies. There are quite a number of partnership models between AI developers and hospitals or government ongoing, mirror similar efforts like us and Beijing Children's Hospital, which is very good momentum. But the true test will be whether the resulting products are actually usable. If the product is good, users, doctors, and even regulators will respond. Problems once thought unsolvable may turn out to be addressable – if the value is truly felt.

So, the future of IAs in healthcare is ambitious, uncertain, and extremely promising. The industry doesn't need more demos or hype; instead, it needs well-designed, trusted, effective tools. The next few years will decide whether IAs are just a fancy add-on or a fundamental shift in how we deliver care.

Q16: What kind of long-term transformation do you envision as AI begins to reshape the nature of work and human productivity?

P2: By 20 years out, when these new digital species truly arrive, they'll be able to conduct large-scale scientific experiments. With that, we'll gather more data and begin unlocking life's deepest mysteries. Humanity will reach a whole new level of understanding about the natural world—and about our place in it. It'll be a scientific explosion. And yes, robots will be running around everywhere. But we will be walking side by side with machine counterparts. That's the next evolution.

Appendix 3: Interview Transcript – Participant from hospital (P3)

Q1: Recently, many public hospitals have rapidly deployed large language models or IAs. Why are hospitals investing in AI now?

P3: Honestly, it all comes down to the right timing, favorable conditions, and the right people in charge. There has been a significant national push in healthcare innovation: from policy frameworks to funding, and even talent development. Overall, we have seen clear progress across the board. But we also need to be realistic. Foundational innovation remains relatively weak, and the coordination and efficiency of research resources still leave room for improvement. These are areas that require further breakthroughs. But overall, the momentum is positive. And there are three main drivers:

First, policy direction. Like I just said, at the national level, there's a clear push to distribute high-quality healthcare resources more evenly. But the lack of doctors and experience at grassroots facilities has been a long-standing issue. With the rise of AI in healthcare, there's a widespread hope among the government, industry, and the public that AI can help scale healthcare expertise. AI is like a "universal assistant," capable of replicating the knowledge of top hospitals and bringing it to underserved areas. It's a key tool for boosting grassroots capabilities.

Second, technical maturity. Whether it's a large model or an AI agent, it has lowered the barrier for tech validation. The applications in medical fields, like automatic record summarization and triage suggestions, have already been pilot-tested at top hospitals like Peking Union Medical College and West China Hospital. It works.

Third, operational demand. Hospitals are seeing more patients than ever. AI helps lighten the load for healthcare workers by automating routine tasks, such as medical record input and basic diagnostic support. It frees up doctors to focus on real clinical decisions instead of paperwork.

Q2: Public hospitals are typically cautious with new tech. What made this rollout move so quickly?

P3: You are not wrong. Healthcare usually moves slower than a snail on vacation. Public hospitals are risk-averse by nature. Without strong international case studies or domestic success stories from comparable institutions, most just sit on the fence. Even when some hospital leaders do take the plunge, they usually avoid announcing it publicly at the start. This recent acceleration comes down to collaboration. Tech companies took on most of the hardware and system integration costs, allowing hospitals to participate with minimal upfront investment. Besides, leading hospitals provided replicable case studies that others could benchmark, cutting down trial-and-error time. But let's not sugarcoat it, the core factor is people. It's about leadership awareness and vision. The bolder hospital directors genuinely believe in AI+ as a gamechanger, not just a task for the IT department. Plus, let's be real,

things like smart hospital ratings and electronic medical record scores are now tied to leadership evaluations; so having a timeline for implementation is no longer optional.

Q3: In your view, where does AI healthcare stand right now? What are the key focus areas for the next three years?

P3: In my opinion, we're still in the phase of scenario testing and tool development. AI mainly supports clinicians and healthcare staff for now – things like auto-generating medical record drafts for doctors to edit or serving as a triage assistant in apps to help patients describe symptoms and get routed to the right department.

Looking ahead, the focus will be on clinical decision support. That means applications like tumor margin mapping and surgery simulations high-value areas where AI needs validation through multicenter clinical trials before it can take part in core decision-making. Another direction is ecosystem integration – linking AI with insurance billing and healthcare supply chains. For example, AI-based DRG tools can assess whether a treatment path is reasonable in real time, helping hospitals optimize costs.

Q4: What role should public hospitals play in adopting AI technologies?

P3: Based on our experience, the key lies in building a systematic innovation model within the hospital.

First, you need a strong governance structure. Just take our hospital for example, we've set up a dedicated leadership group responsible for top-level planning, policy-making, and strategic coordination. This avoids fragmented efforts and ensures innovation is organized, not improvised. One reality is, hospitals, particularly public ones, don't have the luxury to chase every AI gold rush, there're a lot of under-resourced areas. So, we also maintain three critical "priority lists" during adoption – one for unresolved clinical challenges, one for first-of-its-kind exploratory projects, and one for tracking major trends and risks in the field. These are converted into focused R&D efforts that generate tangible results.

Operationally, we rely on five key support mechanisms – funding, platform, talent, translation, and incentives. For funding, we diversify our sources, such as combining national-level research projects with social capital. For platforms, we integrate resources like labs and databases. In terms of talent, we're breaking away from seniority-based systems and giving young scientists more opportunities. For results transformation, we've established clear, streamlined pathways to turn research into real-world applications. We're also reforming incentive structures so that researchers are fairly rewarded.

Q5: What's the cooperation model with tech companies? How do you ensure data security and compliance?

P3: Good question. Data security is non-negotiable. We've set up a three-layer protection system. First rule, data localization. All patient data is processed on in-house servers. Nothing goes to the public cloud. Partners can only access anonymized sample data. Second rule, dynamic anonymization. Sensitive data like ID numbers and addresses are automatically replaced before training. So, names become "Patient 001," and addresses are turned into "Some District." Even birthdates are generalized to age ranges—all while preserving essential medical traits. Third rule, tiered access control. Clinical departments can only access their own data. Every operation is logged and audited in real time. Any abnormal behavior triggers alerts.

Q6: What do you believe are the key areas the healthcare should focus on to build new competitive technologic strengths?

P3: To truly raise the standard of care and drive meaningful progress, the healthcare sector needs to advance on multiple fronts.

First and foremost, we have to strengthen technological innovation. That means connecting the full R&D chain – from lab research to clinical application, and enabling closer collaboration between hospitals, universities, and tech companies. Digital transformation is another critical direction. Hospitals need to start making full use of their data. At present, many systems operate in isolation, department by department, hospital by hospital, creating data silos. We need to integrate these fragmented sources and turn them into assets that drive clinical value.

From an industrial perspective, we need to focus on frontier technologies, areas like mRNA vaccines and 3D-printed medical devices. We must build complete innovative ecosystems so that lab discoveries can be translated quickly into real-world products. At the same time, we should actively support innovations especially in drugs and medical devices.

We also need to address regional imbalances in healthcare development. One approach is to build cross-regional innovation centers that combine the strengths of the more developed regions with the unique advantages of central and western areas. For instance, one region may specialize in traditional Chinese medicine, while another leads in medical device manufacturing. Creating collaborative clusters is far more effective than isolated efforts.

Of course, we can't ignore policy and talent. There's still work to be done in areas like drug approval and insurance reimbursement. We need a more flexible and adaptive regulatory framework. As for talent, the future depends on interdisciplinary professionals – people who understand medicine, engineering, and data science. These hybrid talents are highly sought after, and we must do more to cultivate them.

Q7: As the leader of hospital informatics, how do you view the continuous evolution of AI agents in clinical settings?

P3: You have to look at this problem from the ground up. During adoption, we've always followed a "two-engine" approach: one engine is technology, the other is real clinical practice. Tech can't just float in theory, it has to be grounded in what actually happens in the hospital, and it needs to learn from it constantly.

Let me share an example. Last year, we worked with United Imaging to bring in an AI-powered imaging system. On the surface, it looks like it's just using large models to interpret scans. But what's really interesting is the system we built behind it: a kind of feedback loop where the AI makes a preliminary diagnosis, and then our senior radiologists review it. When they make changes or corrections, the system tracks those differences and uses them to retrain itself—kind of like how a doctor builds experience by reviewing cases over time. So the model keeps learning, not just from data, but from real clinical judgment. For us, that's what real AI evolution looks like in healthcare – continuous improvement based on collaboration with experts and real patients.

Q8: How can we ensure that AI agents remain efficient and continuously learn in healthcare landscape?

P3: That's a big question. Generally speaking, there are two main levers. The first is the core capability of the AI agents itself. The second is the real-world data feedback loop that comes from actual clinical deployment. Think of these as the agent's left and right brain. One side provides the basic cognitive capabilities, such as understanding, reasoning, decision-making, and task execution in a clinical context. The other ensures the system can keep learning and iteratively improve to solve real clinical problems over time.

I'm not sure if you are familiar with the "impossible triangle" in health economics: traditionally, you can't achieve high quality, low cost, and high efficiency all at once. But what's exciting now is that through the co-evolution of AI, we're starting to break that limitation. For example, in Shanghai, radiologists at top-tier hospitals can now remotely support imaging consultations at lower-tier clinics. And family doctors are managing three times as many chronic disease patients as before, thanks to intelligent assistants. That kind of "fast, good, and cost-effective" breakthrough is a clear sign that technology is helping reshape the value equation in healthcare.

Q9: The tech sounds promising, but real-world adoption must be tough. Can you share any specific challenges you've faced?

P3: Oh, you hit the sore spot. This question goes to the core of our healthcare system – there're no shortage of challenges. Honestly, the biggest one is that the fully adoption isn't going to take off overnight – it is a long game. Take clinical validation for example. Say you've got an AI for cancer screening. It might catch more early-stage cases, but unless long-term, prospective trials show that this actually reduces mortality, it's not considered proven. And these trials can take a decade and there's no guarantee of success.

Then, there's the infrastructure limitations. Take our hospital for example, some departments have outdated terminal devices that can't support the agents. Imagine peak hours at outpatient clinics, and your doctor tries to use the DeepSeek R1 (deep think) – it could be a disaster for both them and the patients in line. Also, strict privacy rules make cloud deployment nearly impossible since most of our data can't leave the premises. But setting up enough local computing power? Most hospitals just can't afford the hardware. We've spoken with top-tier hospitals, and even they admit they can't run large models on-site. So you end up with this mismatch: hospitals can't send data out but also can't process it in-house. What else concerns me most is the quality of the data itself. While we often say that the healthcare industry has a massive volume of structured data such as electronic medical records and hospitalization cases, the question is whether the data is truly reliable enough to support more advanced AI work or keep pace with how quickly we're consuming it to train models. I have my doubts. Even outside of healthcare, other industries are reporting that data quality has become a major bottleneck.

Besides that, you also have to think about the people. Developing healthcare AI isn't just about feeding data to models. You need deep collaboration between clinical experts and IT engineers, or some hybrid experts who understand both healthcare and IT. These labors are incredibly rare. Even top-tier healthcare data depends on seasoned experts. And even if we did, good AI still depends on quality training data. That means experienced clinicians doing data labeling. And honestly, they barely have time to treat patients, let alone annotate datasets. So we're always short on time, tools, and people.

Next one is ecosystem issue. Right now, every hospital is all doing its own thing – partnering with different vendors, training different models or platforms. Everyone is burning resources duplicating work. If we had a shared hub or some kind of unified platform like what Alipay is starting to build, it would make a huge difference. But we are not there yet.

Ethics is another concern. AI improves efficiency, yes, but then comes the question of responsibility. If AI writes a medical report, who's legally responsible? Who signs off? Hospitals can't solve these issues alone. We need national policies that give clear guidance or framework that tell us where the boundaries are.

And of course we can't ignore economic viability. Many AI tools out there aren't cost-effective. They haven't reached a level of maturity where hospitals see a strong reason to keep using them. Especially if we're using enterprise-grade AI just to direct patients to the right floor. That's not a smart use of resources. We need funding, support, and frankly, more maturity in the tech before it becomes sustainable.

So yeah, adoption isn't just about plugging in some software. It's a full-scale transformation. You need upstream support – strong tech, shared infrastructure – and downstream buy-in from users. Also, a safety net of policies and regulations. Only then can healthcare AI actually break through the pilot stage and move forward.

Q10: How do hospitals approach the commercialization of AI?

P3: Personally, I think healthcare is well-suited for scalable AI adoption since it has both massive data volumes and a significant unmet demand compared with other industries. From our hospital's perspective, we're not looking for one-off procurement, we care more about co-building ecosystems. I believe that's where commercial success will be decided. Since the launch of GPT and other large models, the hope put on AI is not just reducing costs anymore; we're using it to create new forms of digital productivity for the industry. Great technology isn't enough on its own, in contrast, we need more participants, more real-world use cases, and constant iteration. A powerful system is like good soil, but it takes a variety of AI models or agents – each tailored to a specific need or environment – to grow a healthy, thriving ecosystem.

Q11: What makes a commercialization path meaningful for hospitals?

P3: When we evaluate AI products, we treat it like triage for an emergency patient, where safety and regulatory compliance always come first. We heard lots of vendors boast about "98% accuracy" or "40% time savings," but for us, that remaining 2% error margin could cost a life. In our hospital, we have three non-negotiable rules:

Data sovereignty must be absolute – 100% of patient data stays within our medical data hubs. Even electronic medical records can't leave our premise.

System integration costs must not exceed 15% of our HIS annual maintenance budget, and the system has to be intuitive enough for our doctors to use without heavy retraining.

Full traceability – such as who ran the AI, which version of the algorithm was used, and a digital fingerprint of the original data. If anything goes wrong, we need to be able to track it back immediately.

Being overly cautious might mean missing out on some of AI's potential benefits, but rushing in would expose us to risks we simply can't afford. At the end of the day, we're not looking for flashy "black box" tech. What we need is an AI system that can embed seamlessly into our existing workflows, like a digital resident-in-training – someone who reduces repetitive work in radiology but never crosses the safety lines that, frankly, are stricter than even the financial industry's risk controls. If an AI tool can solve real problems without stepping over the line, that's the kind of innovation public hospitals are actually willing to pay for.

Q12: What advice would you give to other hospitals planning to deploy AI?

P3: Maybe three tips. First, prioritize use cases. Start with low-risk, repetitive tasks like auto-reading lab reports. Prove the value early before moving into complex areas. Then, build

cross-functional teams. IT must work closely with clinical and research departments. Hold regular interdisciplinary meetings to ensure the tech meets real needs. The last one is planning for upgrades. Use modular designs so you're not stuck with one vendor. This keeps things flexible for future algorithm improvements and data expansion.

Q13: From your perspective, what new roles do you think AI could take on within healthcare?

P3: While we are open to future possibilities, and we do hope technologies will bring revolutionary change in healthcare, the reality is more complex. Dramatic reform needs the right timing and a well-planned strategy. In my view, a more practical approach is introducing new roles or service scenarios within the existing hospital systems.

Personally, I don't believe AI will easily replace doctors because the knowledge base keeps changing. From the hospital side, we're more looking for the valuable roles AI can play in care coordination assistant, like a care planner. This is about connecting the dots for both doctors and patients – where to go, who to see, how to schedule it, and what to expect.

You have to understand the current bottleneck of the public hospital system in big cities like Beijing and Shanghai – strained medical resources, overwhelmed hospitals, overcrowded wards, and depleted staff push care to the brink. Especially under the tiered healthcare policy promoted by the government, matching patients with the right level of service is a constant operational challenge. Most people don't need a top specialist; they need a well-informed, local provider. If someone lives in a remote area, it's simply unrealistic to expect the country's top expert to treat them at minimal cost. Ideally, they'd find a qualified local specialist to help. That's where an AI assistant could step in: triaging based on need, location, availability, and even cost, then helping align patients with the right level of care at the right time.

Another area where AI could make a huge impact is in mental health support and emotional counseling. This is a massive unmet need in our healthcare system. Psychological issues are widespread, but we have very limited resources to address them. For example, while over 1.5 million people were once registered as counselors in China, fewer than 50,000 are currently practicing. Hotlines, like the one operated by the Guangzhou government, only handle about 30,000 calls per year, with a connection rate below 20%. The supply of quality mental health support is nowhere near the scale of the demand.

Personally, I believe AI is at its best when solving large-scale social challenges, especially those that can't be met through human capacity alone. Emotional distress may not always require clinical intervention, but people still need to talk, to vent, to be heard. We envision AI playing the role of a compassionate listener, someone users can turn to anytime, without judgment or pressure. That kind of low-barrier, always-available support could prevent stress from snowballing into deeper social or medical crises. So whether it is helping users navigate the healthcare system more efficiently or simply being there for them during emotionally

difficult moments, AI has the potential to serve as both an intelligent coordinator and a social stabilizer.