Warsaw Forum of Economic Sociology vol. 16, no. 32 Autumn 2025
SGH Warsaw School of Economics; Collegium of Socio-Economics; Department of Sociology

From Clinical Hype to Operational Value: Assessing AI Use Cases in Polish Hospitals

Piotr Markowski[*], Anna Kowalczyk[**], Hubert Łazarczyk[***], Tomasz Berdyga[****]

Abstract

Despite intense policy attention, artificial intelligence (AI) adoption in the Polish healthcare system remains far more visible in rhetoric than in formal procurement requirements. This study provides a managerial and policy-oriented mapping of AI deployment intent by analysing public procurement as an empirical proxy for organisational investment decisions. Using official repositories the Polish Public Procurement Bulletin (BZP) and EU Tenders Electronic Daily (TED), the study compiled a corpus of 85,501 healthcare-related notices published in 2022–2024 and applied a deliberately conservative, multi-stage identification pipeline combining large-language-model screening, Terms of Reference (TOR) verification, and expert oversight. Only 210 procedures (approximately 0.25% of the corpus) contained explicit, verifiable AI functionality at TOR level, highlighting a substantial gap between "AI hype" and procurement-grade specification. The confirmed AI procurements are strongly concentrated in clinical applications (64.2%), while administrative processes account for 15% and research/scientific applications for 14%. The findings suggest a structural

[*] SGH Warsaw School of Economics and National Institute of Cardiology, Poland, e-mail: piotrmarkowski83@gmail.com, ORCID: 0000-0002-6201-7137
[**] National Institute of Cardiology, Poland; e-mail: anna.elzbieta.kowalczyk@gmail.com, ORCID: 0009-0009-5389-6174.
[***] e-mail: hlazarczyk@gmail.com, ORCID: 0009-0003-7199-8699
[****] e-mail: tberdyga@gmail.com, ORCID: 0009-0001-6954-1337

imbalance that may reflect not only technology maturity and clinical prestige, but also asymmetric funding incentives, including the National Recovery Plan (KPO), which has explicitly linked parts of digital-health modernisation to clinically framed AI implementations more strongly than to back-office automation. The article discusses governance implications under the EU AI Act, and argues that administrative AI can provide a lower-risk, faster-to-benefit pathway to measurable operational value, if procurement strategies and incentive structures are rebalanced.

Keywords: artificial intelligence, healthcare system, public procurement, hospital management, operational efficiency; EU AI Act; National Recovery Plan

JEL Classification: I1; H57; O33; L86

1. Introduction

1.1. Background: AI in healthcare beyond the clinical hype

The rapid expansion of artificial intelligence (AI) in healthcare has attracted sustained attention from policymakers, clinicians and technology vendors across Europe. Large-scale public programmes and optimistic productivity narratives have helped position AI as a strategic lever for modernising health systems and addressing persistent capacity constraints (Ng Kok Wah, 2025, pp. 55-56). Economic modelling has also been used to argue that automation and decision support could yield sizable efficiency gains at system level (Ng Kok Wah, 2025, pp. 55-56).

Despite this breadth of interest, the dominant policy and commercial narrative remains disproportionately centred on clinical applications, particularly in data-rich domains such as diagnostic imaging, radiology, pathology and clinical decision support. Medical imaging has effectively become the canonical use case, with a growing ecosystem of solutions targeting triage, reconstruction and lesion detection (Martin et al., 2025, pp. 2-3; Bukowski et al., 2020). Professional societies have also begun to issue specialty guidance, for example in endoscopy (Messmann et al., 2022, pp. 1211-1212). In parallel, health technology assessment (HTA) scholarship has expanded to question how existing HTA domains should be adapted for AI-based medical devices (Boverhof et al., 2024, p. 4; Farah et al., 2024, p. 4). At national level, mapping studies

similarly emphasise clinical and patient-facing digitisation trajectories, including telemedicine and eHealth initiatives (Glinkowski et al., 2025, p. 1).

By contrast, non-clinical AI applications, particularly those aimed at hospital administration and management, remain comparatively under-examined. Early contributions point to potential value in areas such as operational forecasting, resource allocation and process automation (Urbi & Tiva, 2025, p. 34), as well as decision support for management functions (Alves et al., 2024, p. 11) and broader administrative performance improvement agendas (Almagadi et al., 2025). Evidence on organisational readiness and adoption dynamics is also emerging, including descriptive insights from hospital IT decision-makers (Weinert et al., 2022).

This imbalance matters because the implementation risks, evidentiary requirements and governance challenges of clinical AI differ markedly from those of administrative AI. Clinical tools often require rigorous validation, continuous monitoring and careful consideration of workflow integration and safety (Boverhof et al., 2024, p. 4; Coiera & Liu, 2022; Chomutare et al., 2022, pp. 10-14). Administrative and managerial applications may be lower-risk in clinical terms, yet they can still reshape organisational processes, accountability structures and resource distribution (Alves et al., 2024, p. 11; Almagadi et al., 2025; Sciarretta et al., 2022, p. 2). Procurement and implementation frameworks therefore need to be assessed not only for their technical adequacy but also for their ability to support integration, monitoring and evaluation across diverse hospital contexts (Urbi & Tiva, 2025, p. 34; Khan et al., 2024).

In this article, the term "deployment inequalities" refers to systematic differences in the adoption and operationalisation of AI across settings, functions and organisational capacities. Such inequalities can be reinforced by uneven access to infrastructure, specialist staff and implementation know-how, which in turn shapes where AI is piloted, scaled and sustained in routine practice (Martin et al., 2025, pp. 2-3; Bukowski et al., 2020).

1.2. Problem statement and research gap

Although several strands of literature address clinical AI, HTA, and high-level digital transformation, there is limited empirical work that systematically maps where AI is being adopted in hospitals and whether adoption concentrates in particular functions or settings. Existing procurement-oriented syntheses indicate that most implementation guidance remains focused on planning and evaluation, with much

less attention to the operational phases of 'doing' and 'acting' (Khan et al., 2024). At the same time, national and sector reports emphasise the perceived need for modern technologies and AI, but often lack granular evidence on real-world uptake across organisational processes (Bartusek & Kulawik, 2021, p. 133).

The motivation for this study is therefore twofold. First, it seeks to clarify whether AI adoption in healthcare is skewed towards clinical use cases at the expense of administrative innovation, despite the potential of AI to relieve operational pressure and improve management capabilities. Second, it assesses what these patterns imply for evidence generation and funding decisions, given that economic evaluations of AI remain relatively scarce and heterogeneous.

Notably, a recent scoping review of economic evaluations identified only a small number of robust cost-effectiveness studies relative to the overall volume of AI publications, highlighting the methodological and reporting gaps that limit comparability across settings and tools (Gomez Rossi et al., 2022).

This evidence gap is particularly salient in Poland, where digital transformation has progressed rapidly in selected domains, yet implementation remains uneven across organisations and regions (Glinkowski et al., 2025, p. 1). Moreover, the adoption of AI intersects with public-sector governance and administrative law, raising questions about accountability, transparency and legal remedies in the event of algorithmic error or contested decisions (Jakubek-Lalik, 2024).

1.3. Aim and research questions

To address these issues, this article investigates how AI adoption is distributed between clinical and administrative domains, and whether observable patterns suggest structural inequalities in deployment across hospitals and healthcare systems.

Methodologically, the study draws on publicly available procurement data as a pragmatic proxy for adoption signals, building on procurement frameworks research that demonstrates how purchasing and contracting artefacts can reveal implementation priorities and bottlenecks (Khan et al., 2024; Ramsay et al., 2025, p. 1).

RQ1: What types of AI applications (clinical vs administrative) are most frequently procured and implemented in hospitals?

RQ2: How do adoption patterns vary across hospital functions and organisational characteristics?

RQ3: What factors appear to contribute to observed deployment inequalities, and what are the implications for policy and governance?

The analysis is intentionally confined to procurement as an observable adoption pathway. It does not aim to measure clinical effectiveness or algorithmic performance, but rather to characterise the distribution of investments and intended uses as reflected in procurement records.

By combining procurement signals with an implementation-oriented interpretation, the study aims to contribute evidence that is relevant to both AI governance debates and practical decision-making in hospital management.

1.4. Contribution and structure of the paper

The article contributes to the literature by offering an empirically grounded view of where AI-related investments are directed, thereby illustrating how clinical and administrative trajectories may diverge in practice. It also links these patterns to the broader discussion on AI opportunity costs, governance and evidence gaps (Ng Kok Wah, 2025, pp. 55-56; Urbi & Tiva, 2025, p. 34; Khan et al., 2024; Gomez Rossi et al., 2022).

The paper is structured to progress from problem framing to a transparent empirical assessment based on public procurement evidence. Section 1 introduces the policy and research context for AI in hospitals, articulates the gap between clinical prominence and organisational back-office innovation, and formulates the study aim alongside a set of research questions that guide the remainder of the analysis.

Section 2 develops the conceptual and theoretical foundations required to interpret procurement signals, including the distinction between clinical and non-clinical AI domains, the related risk and value profiles, and the concept of "deployment inequalities". It also justifies why procurement records can be treated as a credible proxy for real investment decisions, while acknowledging the methodological trade-offs inherent in this lens.

Section 3 specifies the methodological approach, including the data sources and corpus construction, the multi-stage AI identification pipeline (from exploratory screening to strict confirmation), and the taxonomy used to classify AI use cases.

Section 4 then reports results on the overall prevalence of confirmed AI procurement, domain-level distributions, and functional profiles across clinical, patient-facing non-clinical, administrative, and research applications, which are subsequently interpreted in the discussion (Section 5) and synthesised into concluding implications (Section 6).

2. Conceptual and Theoretical Background

The integration of artificial intelligence (AI) into the healthcare system constitutes one of the most consequential technological shifts in contemporary medicine. Yet the trajectory of AI adoption reveals a persistent paradox. Clinical AI applications, particularly in diagnostic imaging and clinical decision support, have attracted intense research attention and prominent media coverage, while real-world implementation remains limited and concentrated within a narrow set of use cases. By contrast, non-clinical and administrative AI applications, which in principle face fewer safety and regulatory barriers, remain comparatively underexamined in the academic literature and are sparsely documented in empirical deployment studies. This chapter establishes the conceptual and theoretical foundations needed to explain this asymmetry. It distinguishes clinical and non-clinical AI domains, reviews frameworks for analysing risk, value, and deployment inequalities, and outlines how public procurement data can serve as a practical lens on actual investment patterns in healthcare AI.

2.1. AI in Healthcare Systems: Clinical and Non-Clinical Domains

AI applications in healthcare span a wide range of functions, from patient-facing diagnostic tools to back-end administrative systems designed to optimise organisational processes. Distinguishing these domains is essential for analysing adoption patterns because clinical and non-clinical AI systems differ materially in risk profiles, regulatory requirements, implementation complexity, and their potential to influence healthcare delivery.

2.1.1. Clinical AI Applications

Clinical AI refers to technologies that directly support or automate elements of clinical decision-making, including diagnosis, treatment planning, prognosis, and patient monitoring. The defining feature of clinical AI is its direct bearing on patient care pathways and outcomes, which subjects these systems to stringent oversight and heightened expectations regarding safety, effectiveness, and accountability.

Medical Imaging and Diagnostic Support

The most prominent domain of clinical AI research and development is medical imaging, notably in radiology, pathology, and dermatology. Deep learning methods, particularly convolutional neural networks, have demonstrated strong performance in tasks such as lesion detection, tumour classification, and quantitative image analysis (Celi et al., 2022; Horgan et al., 2019). European research activity in this area has been extensive, with numerous proof-of-concept studies reporting high diagnostic accuracy under controlled conditions. AI systems for detecting diabetic retinopathy, identifying pulmonary nodules in chest radiographs, and classifying skin lesions have at times achieved performance comparable to, or occasionally exceeding, specialist clinicians in experimental evaluations (Wolff et al., 2021).

However, a substantial gap persists between experimental performance and routine clinical deployment. Systematic reviews indicate that, despite impressive sensitivity and specificity reported in academic publications, sustained large-scale adoption across European healthcare systems remains uncommon (Popescu et al., 2022; Wolff et al., 2021). Moving from a research prototype to a clinically embedded tool requires more than technical validation: it also requires workflow integration, credible evidence of clinical utility beyond accuracy metrics, and navigation of complex regulatory pathways (Alami et al., 2020).

Clinical Decision Support Systems

Beyond imaging, AI-enabled clinical decision support systems (CDSS) represent another major category of clinical applications. These systems analyse patient data, including laboratory results, vital signs, electronic health records, and genomic information, to provide recommendations for diagnosis, treatment selection, or risk stratification. Most contemporary CDSS are intended to augment rather than replace clinical judgement, positioning AI as a collaborative tool supporting clinician decision-making (McKee & Wouters, 2022).

Across Europe, development and implementation trajectories remain heterogeneous, reflecting differences in national priorities and infrastructure. For example, Italy's national programme to develop an AI platform for public healthcare illustrates efforts at coordinated implementation, emphasising semantic interoperability, data governance, and ethical guidance as prerequisites for effective deployment (Horgan et al., 2019). Such initiatives highlight that clinical AI adoption depends not only on

technical capability but also on systemic readiness across organisational, regulatory, and infrastructural dimensions.

Integrated Diagnostics and Multimodal Analysis

An emerging frontier involves integrated diagnostic systems that combine multiple data modalities, including imaging, histopathology, laboratory results, clinical notes, and patient-reported outcomes, to produce more comprehensive diagnostic assessments. These multimodal approaches may improve diagnostic precision by leveraging complementary sources of information (Bukowski et al., 2020). At the same time, they intensify challenges around interoperability, standardisation, and the validation of systems synthesising diverse information streams.

The Implementation Gap

A recurring finding in the clinical AI literature is the scale of the implementation gap. Despite thousands of studies reporting high experimental accuracy, relatively few systems have achieved sustained deployment in routine clinical practice within European settings (Popescu et al., 2022; Wolff et al., 2021). Barriers include regulatory uncertainty, liability and accountability concerns, difficulties integrating tools into established workflows, limited evidence of cost-effectiveness, and professional resistance where AI is perceived as threatening autonomy or increasing workload (Alami et al., 2020; McKee & Wouters, 2022).

In addition, clinical AI research remains concentrated in particular specialties, notably radiology, oncology, and ophthalmology, contributing to uneven innovation across clinical domains (Celi et al., 2022). This clustering reflects the availability of large, well-curated datasets in image-intensive fields and the relative tractability of bounded tasks suited to supervised learning.

2.1.2. Non-Clinical and Administrative AI Applications in Healthcare

While clinical AI dominates research attention and public discourse, non-clinical and administrative applications constitute a substantial but underdocumented domain of AI use in healthcare. These systems do not directly intervene in clinical decision-making; instead, they target operational, administrative, and managerial processes that underpin service delivery.

Operational Automation and Resource Optimisation

Non-clinical AI includes patient scheduling and appointment optimisation, resource allocation (e.g., theatre scheduling, bed management), supply chain and inventory optimisation, workforce planning, and epidemiological surveillance (Alami et al., 2020). Techniques such as predictive modelling, optimisation algorithms, and natural language processing (NLP) can reduce administrative burden, improve utilisation, and support more responsive operational planning.

Predictive models may forecast admission volumes to inform staffing and capacity decisions. NLP can extract structured information from unstructured clinical text for billing, quality reporting, or research data extraction. Optimisation methods can improve surgical scheduling by predicting procedure durations and reducing idle time, thereby increasing throughput and reducing waiting lists.

Organisational and Economic Functions

Beyond day-to-day operations, non-clinical AI can influence organisational governance, cost structures, and strategic decision-making. AI-driven analytics may identify inefficiencies, reveal patterns in resource use, and inform investment priorities. In this sense, AI functions not only as a technical tool but also as a potential lever of organisational transformation, with implications for workforce composition, skill requirements, and institutional governance (Alami et al., 2020).

The Evidence Gap in Non-Clinical AI Deployment

Despite its theoretical potential, empirical evidence on large-scale deployment of administrative AI in European healthcare remains limited. The academic literature contains relatively few detailed case studies or evaluations of AI in scheduling, billing, supply chain management, or human resources within European contexts. This paucity contrasts sharply with the extensive body of work on clinical AI, indicating a substantial research gap.

Several factors may contribute. First, non-clinical implementations may be perceived as less scientifically novel, reducing publication incentives. Second, administrative AI is often deployed by commercial vendors under proprietary arrangements, limiting visibility and access. Third, outcomes such as cost savings, efficiency gains, or reduced administrative burden can be diffuse and context-dependent, making at-

tribution to a specific AI intervention more challenging than reporting discrete clinical outcomes such as diagnostic accuracy.

Distinguishing Clinical from Non-Clinical Domains

Multiple taxonomies have been proposed. One common approach differentiates systems by primary function: decision support that augments clinical judgement versus automation that executes operational or administrative processes (Shaw et al., 2019). Another treats applications as lying on a continuum from direct patient impact (typically subject to medical device regulation) to system-level organisational interventions with economic, legal, and workforce implications (Alami et al., 2020). This latter view emphasises that non-clinical AI can still shape care delivery indirectly by restructuring organisations and influencing resource allocation.

2.2. Risk, Value, and Inequality in AI Deployment

AI adoption in healthcare is shaped not only by technical capability and organisational readiness but also by perceived risk, assessed value, and structural factors that produce uneven access and uptake. Conceptual clarity therefore requires engagement with frameworks that explain interactions between technology, institutions, and social contexts.

2.2.1. Risk Profiles: Clinical Versus Non-Clinical AI

Clinical and non-clinical AI entail distinct risk profiles, with implications for regulation, implementation pathways, and stakeholder acceptance. In the European Union, these profiles are increasingly shaped by a layered governance environment, combining sectoral rules for medical devices (MDR/IVDR), horizontal data protection (GDPR), and the risk-based obligations introduced by the EU Artificial Intelligence Act (AI Act) (European Parliament & Council of the European Union, 2024; Aboy et al., 2024).

Clinical AI Risks

Clinical AI risks primarily relate to patient safety and clinical outcomes. Errors in diagnostic or decision-support systems can contribute to misdiagnosis, delayed treatment, or inappropriate interventions (McKee & Wouters, 2022). Concerns abo-

ut bias further complicate risk, as systems trained on non-representative datasets may perform unevenly across demographic groups, potentially exacerbating health inequalities (Celi et al., 2022).

Further issues include explainability and accountability. Many high-performing models operate as "black boxes", limiting transparency for clinicians and patients and complicating error detection and correction (McKee & Wouters, 2022). Questions of legal and professional responsibility also remain contested: when AI contributes to harm, assigning accountability among developers, institutions, and clinicians is ethically and legally complex (Pashkov et al., 2020). Ethical analysis therefore extends beyond technical performance to include fairness (justice), non-maleficence, and respect for patient autonomy, especially where AI meaningfully influences clinical judgement or patient pathways (Char et al., 2018; World Health Organization, 2021).

In the European Union, many clinical AI systems fall under medical device regulation, requiring clinical validation, conformity assessment, and post-market surveillance. The MDR and IVDR impose substantial compliance burdens on developers and deploying organisations (McKee & Wouters, 2022). The AI Act adds a complementary, horizontal layer: many healthcare AI systems are treated as high-risk, triggering lifecycle obligations such as a continuous risk management system, data governance requirements, technical documentation, record-keeping (including logging), transparency to deployers, human oversight, and requirements on accuracy, robustness, and cybersecurity (European Parliament & Council of the European Union, 2024; European Commission, n.d.-a, n.d.-b). In practice, this pushes clinical AI governance towards more explicit "assurance" work: documenting intended purpose and limitations, monitoring performance drift, and ensuring that human oversight is operationally meaningful rather than nominal (Aboy et al., 2024; van Kolfschooten & van Oirschot, 2024).

Non-Clinical AI Risks

Non-clinical systems entail a different constellation of risks, centred on privacy, data governance, organisational disruption, and economic implications. Administrative AI often processes sensitive personal data, including identifiers, financial information, and workforce records, raising concerns about security, unauthorised access, and compliance with GDPR (Alami et al., 2020). Even where non-clinical systems do not directly affect diagnosis or treatment, they can shape patients' experiences and institutional practices through prioritisation, resource allocation, and process

automation, which introduces ethically salient risks related to fairness, transparency, and the distribution of burdens and benefits across patient groups and staff (World Health Organization, 2021).

Organisational risks include workforce displacement, deskilling, and disruption to established workflows. Poorly designed systems may introduce bottlenecks or unintended consequences that degrade performance (Shaw et al., 2019). Economic risks include misallocation of resources, overinvestment in technologies that do not deliver expected returns, and vendor lock-in that constrains future flexibility. Compared with direct patient harm, these impacts can be more diffuse and harder to quantify (Alami et al., 2020). From an AI Act perspective, many non-clinical tools may fall outside the strictest "high-risk" category; however, where systems are high--risk in context of use, deployers may face additional governance duties, including a fundamental rights impact assessment for certain deployers and settings (European Commission, n.d.-c; van Kolfschooten & van Oirschot, 2024). This reinforces the ethical need to treat "back-office" AI not as value-neutral infrastructure, but as socio-technical systems with real distributional consequences.

2.2.2. Comparative Framework: Clinical Versus Non-Clinical AI

Comparative evaluation requires frameworks extending beyond technical performance metrics to include organisational, economic, ethical, and social dimensions.

Health Technology Assessment Approaches

Health Technology Assessment (HTA) offers a multidimensional framework encompassing clinical effectiveness, safety, cost-effectiveness, ethical implications, organisational impact, and social consequences (Alami et al., 2020; Leo et al., 2022). HTA has increasingly been applied to AI, reflecting the view that accuracy metrics alone are insufficient: workflow integration, user acceptance, and system-level effects are often decisive for value (Alami et al., 2020).

HTA highlights key differences between domains. Clinical AI evaluation prioritises efficacy, safety, and evidence of benefit for patient outcomes, often requiring robust trials or high-quality observational studies. Non-clinical evaluation emphasises operational indicators such as efficiency, throughput, cost savings, and user satisfaction, which are typically more heterogeneous and context-dependent (Leo et al., 2022).

Implementation Science Frameworks

Implementation science frameworks, including NASSS (Non-adoption, Abandonment, Scale-up, Spread, and Sustainability), structure analysis of barriers and facilitators in complex healthcare settings (Shaw et al., 2019). NASSS examines the condition, the technology, the value proposition, the adopter system, the organisation, the wider institutional context, and adaptation over time.

Applied comparatively, clinical AI often exhibits high complexity in regulatory compliance, clinical validation, and professional acceptance, reflecting the high-stakes nature of clinical decision-making. Non-clinical AI may face greater complexity in organisational integration, workflow redesign, and change management because administrative processes cut across departments and stakeholder groups (Shaw et al., 2019).

Operational Versus Systemic Adoption

Some frameworks distinguish operational adoption, local uptake within a setting, from systemic adoption, sustained scaled deployment supported by governance, funding, and policy (Li, 2022). This distinction is particularly relevant in comparing clinical and non-clinical AI. Clinical AI frequently requires systemic conditions such as regulatory approval, guideline integration, reimbursement mechanisms, and training. Non-clinical AI may be adopted operationally more readily via managerial decisions, given lower regulatory and safety barriers (Alami et al., 2020).

2.2.3. The Concept of Deployment Inequalities

Deployment inequality refers to systematic disparities in the availability, adoption, and quality of AI technologies across populations, institutions, and regions. Such inequalities can emerge through multiple mechanisms and carry implications for the distribution of AI's benefits and burdens.

Geographical and Institutional Disparities

Research activity is geographically concentrated, with a large share of datasets, publications, and commercial developments originating from the United States and China (Celi et al., 2022). Within Europe, adoption patterns are heterogeneous, shaped by healthcare system structures, digital infrastructure maturity, funding, and regulatory environ-

ments (Bukowski et al., 2020). Countries with advanced digital health ecosystems and strong data governance, such as Nordic states, Germany, and the Netherlands, may be better positioned for adoption than those with less developed infrastructures.

Institutional disparities reinforce these patterns. Large academic medical centres and private hospitals may have greater capacity to invest in AI, recruit specialist staff, and manage regulatory requirements than smaller, resource-constrained public hospitals or rural facilities (Horgan et al., 2019). This risks creating a two-tier system where AI-enabled services cluster in well-resourced settings.

Technological Readiness and Digital Maturity

Deployment inequalities are also mediated by digital maturity, including EHR capability, interoperability, IT infrastructure, and workforce digital literacy (Bukowski et al., 2020). AI adoption typically presupposes standardised data formats, secure data storage and transfer, and adequate computational resources. Institutions lacking these prerequisites face substantial structural barriers regardless of the availability of AI solutions (Horgan et al., 2019).

The Polish healthcare system illustrates these constraints. While Poland has advanced digitisation initiatives, interoperability remains limited, data quality varies, and many providers, especially in rural or underserved areas, lack technical infrastructure and workforce capacity needed for advanced AI implementation (Bukowski et al., 2020). These conditions may widen disparities between well-resourced urban centres and peripheral regions.

Socioeconomic Dimensions of Access

Budget constraints and uncertain returns can limit adoption in public systems, particularly where upfront costs are high and benefits accrue over longer horizons. Reimbursement structures that do not recognise AI-enabled services may further discourage adoption (Alami et al., 2020). Workforce capacity is another critical determinant: AI deployment requires specialised human capital, including data scientists, informaticians, and clinicians with AI literacy. Shortages of such professionals can amplify inequality, particularly in resource-limited settings (Horgan et al., 2019).

The Hype-Implementation Gap as a Deployment Inequality

A further form of deployment inequality concerns the mismatch between the visibility of clinical AI and the limited extent of routine implementation. Clinical AI attracts disproportionate attention and investment despite persistent barriers, while non-clinical AI, which may be easier to deploy from a regulatory and safety perspective, remains underexplored (Popescu et al., 2022; Wolff et al., 2021). This suggests that deployment patterns are shaped not only by feasibility or value but also by prestige, funding incentives, and commercial dynamics. Explaining this asymmetry requires attention to institutional and cultural drivers alongside technical evaluation.

2.3. Public Procurement as a Window on AI Adoption

Public procurement is an underutilised yet potentially informative source of evidence on real-world AI investment decisions. Unlike research publications that reflect proof-of-concept studies, procurement records document purchasing decisions by healthcare institutions and thus reflect concrete financial commitments and organisational priorities.

2.3.1. Why Procurement Data Reflect Real Investment Decisions

Procurement records offer several advantages as indicators of adoption. First, they capture actual transactions and budget allocations, providing objective evidence of institutional commitment rather than aspirational or self-reported adoption (Bukowski et al., 2020). Second, public procurement is typically governed by transparency and competitive tendering rules that generate structured data on contract values, suppliers, and contracting authorities. In the European Union, above-threshold procurements are commonly published through standardised platforms such as Tenders Electronic Daily (TED), enabling comparative analysis.

Third, procurement datasets facilitate longitudinal analysis, allowing researchers to track trends, identify emerging technologies, and examine diffusion across regions and institution types. Such analysis can reveal disparities in investment and the evolution of priorities over time.

2.3.2. Completeness and Formalisation of Procurement Records

The research utility of procurement data depends on the detail and standardisation of records. Well-functioning e-procurement systems typically provide fields for contracting authority, contract value, contract type, Common Procurement Vocabulary (CPV) codes, and a brief description of the procurement subject (Bukowski et al., 2020).

However, detail varies substantially. Some notices only indicate broad categories (e.g., "IT services" or "medical equipment") without identifying whether AI components are included. Richer information often appears in technical specifications, frequently provided in separate tender documents, such as the "Detailed description of the Subject of Procurement". These materials are crucial for distinguishing AI-related purchases from conventional IT procurement and for classifying procurements by clinical versus non-clinical domain.

2.3.3. Traceability of Decision Pathways

Procurement records can also support analysis of decision pathways by linking purchasing choices to institutional characteristics, policy contexts, and external drivers such as dedicated funding programmes or regulatory changes. By examining which institutions procure AI systems, researchers can assess associations with size, ownership, geography, teaching status, or specialisation (Bukowski et al., 2020).

Procurement data may also be linked with other datasets, including performance indicators or workforce information, to investigate organisational impacts. While such linkage does not guarantee causal inference, it can support quasi-experimental strategies and more robust evaluation designs.

2.3.4. Examples of Procurement-Based AI Adoption Research in Europe

Despite its promise, procurement-based research on healthcare AI adoption in Europe remains limited. Systematic searches indicate relatively few studies that treat public procurement as a primary data source for assessing investment patterns, deployment inequalities, or diffusion dynamics. This scarcity is notable given the increasing availability of e-procurement platforms and open data initiatives.

Potential explanations include the burden of collecting full tender documentation (often decentralised or access-restricted), the need for extensive cleaning and classifi-

cation, and the interdisciplinary expertise required to interpret procurement data appropriately, spanning health services research, public administration, and data science.

2.3.5. Advantages and Limitations of Procurement Data

Procurement data provide objective, transaction-level evidence and can offer broad coverage of public-sector purchases across clinical and non-clinical applications, including hardware, software, and services. They enable longitudinal analysis and can support examination of diffusion patterns and disparities.

However, important limitations apply. Procurement records typically exclude private-sector purchases, in-house developments, and research-funded pilots, potentially underestimating overall AI adoption. The granularity and completeness of notices vary across jurisdictions and time. Identifying AI-related procurements within large datasets often requires domain expertise and sophisticated text analysis because AI may be embedded within broader systems or described using inconsistent terminology.

A further limitation concerns access to detailed documents. While summary notices are public, full technical specifications are frequently absent from central repositories and may require formal requests to contracting authorities or navigation of decentralised repositories. This raises collection burdens, constrains scalability, and may introduce selection bias if only easily accessible procurements are analysed.
Finally, procurement data primarily capture inputs (investment decisions) rather than outcomes (performance) or impacts (health outcomes, efficiency, equity). Accordingly, procurement analysis should ideally be complemented by methods such as case studies, surveys, or linkage to performance data.

Despite these constraints, procurement data remain a valuable and underused resource for understanding AI deployment in healthcare. The present study addresses this gap by analysing public procurement records from the Polish healthcare system to characterise AI investment patterns, compare clinical and non-clinical adoption, and examine deployment inequalities across institutions and regions. By grounding analysis in procurement decisions, the study aims to move beyond clinical AI hype towards documented patterns of institutional prioritisation and the structural factors shaping equitable AI deployment.

3. Methodology

3.1. Study design and analytical approach

This study adopts a multi-stage, procurement-based analytical design to examine the distribution and functional profile of artificial intelligence (AI) use cases in the Polish healthcare system. Public procurement notices are used as an empirical proxy for organisational adoption intent, reflecting concrete investment decisions rather than aspirational strategies or pilot declarations.

The methodological framework is explicitly conservative and filtration-oriented. Its primary objective is to minimise false-positive identification of AI systems by progressively narrowing the dataset through increasingly stringent analytical stages. AI identification is therefore treated not as a single classification task, but as a sequential verification pipeline, combining large language model (LLM) screening with document-level validation and expert oversight.

The analysis focuses on intended functionality as specified in procurement documentation. It does not attempt to assess post-procurement implementation success, clinical effectiveness, or real-world performance of AI systems.

3.2. Data sources and corpus construction

The initial dataset comprised all healthcare-related procurement notices published between **1 January 2022 and 31 December 2024** in two official repositories:

- the **Polish Public Procurement Bulletin (BZP)**, and
- the **EU Tenders Electronic Daily (TED)** database.

The combined corpus included **85,501 procurement notices**, covering supplies, services, and works procured by public healthcare providers at national, regional, and local levels. Notices were retrieved together with associated metadata (title, description, CPV codes, contracting authority, procedure identifiers) and, where available, full **Terms of Reference (TOR)** documentation, including technical specifications and functional requirements.

Polish Public Procurement Bulletin (BZP)

1. Scope: national and below-EU-threshold public procurement procedures in Poland.

2. Format: XML documents retrieved via the ezamowienia.gov.pl export functions.
3. Fields used: notice title (*title*), short description (*short_description*), notice content (*content*), main and additional CPV codes (*cpv_main*, *cpv_additional*), type of contracting authority, date of publication.

EU Tenders Electronic Daily (TED).

1. Scope: above-threshold public procurement procedures in the European Union.
2. Format: XML files compliant with TED standards.
3. Fields used: analogous to BZP, including *title*, *short_description*, *content*, CPV codes, type of contracting authority, date of publication.

In total, **85,501** notices were harvested from both sources (49,281 from BZP and 36,220 from TED). A multi-step sampling and classification procedure was then applied to identify and code notices related to AI/ML.

Sampling strategy and inclusion criteria

The sampling strategy was designed to approximate, as closely as possible, the universe of AI-related procurements in the Polish healthcare sector in the period 2023-2025.

BZP sample: hospitals and national health institutes

For BZP, inclusion criteria were defined to capture hospital-type entities and key national institutes providing healthcare services or performing national health functions. The following filters were applied:

1. **Type of notice –** Only "contract notice" records were included (initial procurement announcements), excluding e.g. contract award notices.
2. **Timeframe –** Date of publication between **1 January 2023** and **10 December 2025** (inclusive).
3. **Type of contracting authority**

Two complementary filters were used to capture the relevant group of entities:

- notices where the contracting authority was classified as a "Samodzielny publiczny zakład opieki zdrowotnej" (independent public healthcare provider);
- notices where the contracting authority was a "Zamawiający publiczny" (pu-

blic contracting authority) and the name of the institution contained at least one of the following keywords or matched one of the explicitly listed national institutes:
  - keyword filters (substring match in the contracting authority's name): „szpital" (hospital), „Opieki Zdrowotnej" (healthcare), „ZOZ" (healthcare institution), „*lecznictwa*" (treatment / medical care);
- explicitly included national institutes and medical research centres:
  - Państwowy Instytut Medyczny
  - Instytut Gruźlicy i Chorób Płuc
  - Narodowy Instytut Zdrowia Publicznego PZH
  - Narodowy Instytut Kardiologii
  - Narodowy Instytut Geriatrii
  - Narodowy Instytut Onkologii
  - Instytut Psychiatrii i Neurologii
  - Wojskowy Instytut Medyczny
  - Wojskowy Instytut Medycyny Lotniczej
  - Instytut „Centrum Zdrowia Matki Polki"
  - Instytut Pomnik-Centrum Zdrowia Dziecka
  - Instytut Fizjologii i Patologii Słuchu
  - Instytut Matki i Dziecka
  - Instytut Immunologii i Terapii Doświadczalnej.

These criteria produced a BZP sub-sample corresponding to hospitals and national health institutes that are likely to engage in clinically or organisationally significant AI deployments.

TED sample: health sector procurement in Poland

For TED, sampling was performed using the platform's structured search fields, with the following parameters:

- **Scope of search:** all notices.
- Type of notice:
  - "Contract notice – simplified procedure", or
  - "Contract notice" or "Concession notice – standard procedure".
- **Main CPV classification:** all CPV classification
- **Nature of contract:** Supplies or services.

- **Place of performance:** Poland.
- **Publication date:** between **1 January 2023** and **10 December 2025**.
- **Sector of activity of the contracting authority:** health.

In the exported dataset, a substantial number of records did not correspond to distinct procurement procedures. Specifically, the raw export included 10,644 notices relating solely to modifications or changes to previously published procedures, as well as 52 duplicate procedures resulting from repeated publication across stages or notice types. These records were programmatically identified and removed during the Documentation collection and aggregation stage, described later in the methodology. This cleaning step was necessary to avoid artificial inflation of procurement counts and to ensure that each observation in the final dataset represented a unique procurement procedure rather than an administrative update or repetition.

The resulting TED sub-sample includes above-threshold procurement procedures launched by Polish health-sector contracting authorities in AI-relevant CPV categories.

Both sub-samples (BZP and TED) were then pooled into a single analytical dataset representing procurement activity by the Polish healthcare sector. Subsequent AI identification and categorisation were applied uniformly across the combined dataset.

3.3. Multi-stage AI identification pipeline

3.3.1. Stage 1 – Exploratory screening (AI potential)

In the first stage, all procurement notices were subjected to exploratory semantic screening to identify *potential relevance to AI*. This stage was designed for high recall and intentionally permissive classification.

Screening was performed using GPT-4.5, instructed to classify notices as:

- *AI Potential = Yes*, or
- *AI Potential = No*,

based on explicit or implicit references to AI-related concepts, including (but not limited to): artificial intelligence, machine learning, neural networks, predictive analytics, intelligent automation, or adaptive algorithms.

At this stage, no attempt was made to verify whether AI functionality was genuinely required or technically specified.

3.3.2. Stage 2 – Indicative analysis (AI likeness)

Notices flagged as AI potential underwent indicative analysis, again using GPT-4.5, but with stricter criteria. Each notice was assigned one of four AI likeness levels:

- **Certain** – AI explicitly stated as a required component,
- **High probability** – AI strongly implied through functional descriptions,
- **Low probability** – AI mentioned ambiguously or as an optional feature,
- **Unlikely** – AI references judged superficial or marketing-oriented.

This stage served to prioritise records for document-level verification, while retaining borderline cases for transparency and auditability.

3.3.3. Extended analysis set construction

Procurement notices classified as *certain*, *high probability*, or *low probability* were retained for extended analysis. Notices classified as *unlikely* were excluded from further processing but preserved in the dataset for methodological traceability.

At this stage, **842 notices** were retained across both repositories.

3.4. Documentation collection and aggregation

3.4.1. TOR acquisition

For each notice in the extended analysis set, the availability of **Terms of Reference (TOR)** or equivalent technical documentation was verified. Notices lacking accessible TOR documentation were excluded from confirmatory analysis, as AI functionality could not be validated at the specification level.

3.4.2. Aggregation of multi-lot procedures

Procurement procedures comprising multiple lots or repeated technical descriptions were aggregated into single analytical records, provided that AI functionality was consistent across lots. Aggregation was performed to avoid artificial inflation of AI adoption counts due to procedural fragmentation.

3.5. Stage 3 – Criteria-based assessment (AI confirmation)

3.5.1. TOR-level verification

Aggregated records with available TOR documentation were subjected to criteria--based assessment using GPT-5, operating under a strictly defined prompt requiring:

- explicit description of AI or machine learning methods,
- AI-driven functionality embedded in system logic (not optional add-ons), and
- operational relevance beyond generic automation or rule-based systems.

Each record was classified as:

- AI certain = Yes, or
- AI certain = No.

Records failing to meet all criteria were excluded from the confirmed AI dataset.

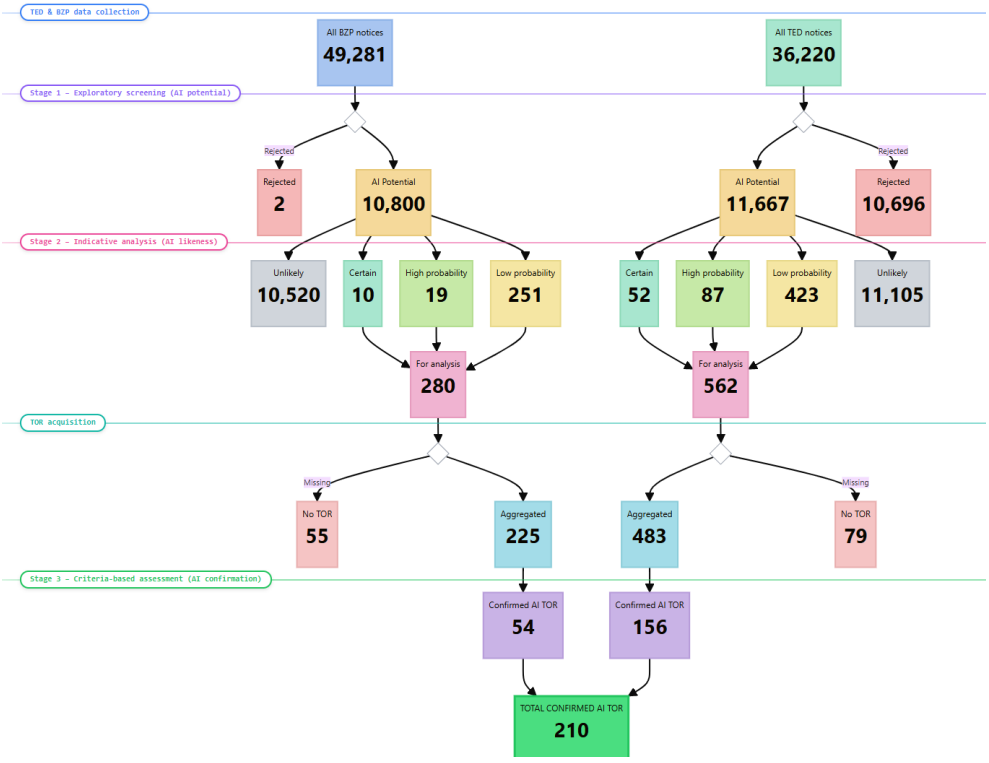

Figure 1. Multi-stage AI identification pipeline diagram

3.5.2. Stage 4 – Strict confirmatory evaluation and expert revision

In the final stage, all records classified as *AI certain = Yes* were subjected to strict confirmatory evaluation, combining:

- **GPT-5.2** analysis at category level, and
- **manual expert review** for ambiguous, high-impact, or borderline cases.

Expert revision focused on preventing misclassification of:

- conventional IT systems marketed as "intelligent",
- rule-based automation lacking learning capability, and
- infrastructure-only procurements without AI logic.

Only records surviving this final verification were included in the confirmed AI procurement dataset.

3.6. Taxonomy of AI application domains

Confirmed AI records were classified using a **two-level functional taxonomy** derived from the functional intent explicitly specified in Terms of Reference (TOR) documentation.

Level 1: Main application domains

At the first level, each confirmed AI procurement was assigned to one of four mutually exclusive domains, reflecting its dominant functional intent:

- **Clinical applications** – systems directly supporting diagnosis, treatment, monitoring, or clinical decision-making.
- **Non-clinical patient services** – AI supporting patient interaction, navigation, registration, or communication without direct clinical decision impact.
- **Non-clinical administrative processes** – AI applied to back-office, managerial, logistical, governance, or IT operational functions.
- **Research and scientific applications** – AI systems primarily supporting research activities, analytics, data preparation, and research-related data governance.

Where a procurement spanned more than one domain, Level 1 assignment followed the primary intended use stated in the TOR, rather than optional or ancillary functionalities.

Level 2: Detailed functional categories

At the second level, procurements could be assigned to one or more detailed functional categories within the Level 1 domain(s), capturing specific AI use cases (e.g., imaging-related AI, document automation, anonymisation workflows, predictive modelling). Detailed categories were applied only when explicitly supported by TOR evidence, and were not treated as mutually exclusive, allowing multiple category assignments where justified.

The definitions, scope, and interpretative guidance for all detailed categories, including their intended meaning and boundary cases, are provided in Appendix A (Detailed taxonomy). This appendix also documents the operational rules used to ensure consistent category assignment across the corpus.

3.7. Analytical strategy

The analytical strategy adopted in this study is descriptive and comparative, with the primary objective of mapping the structural patterns of AI-related procurement in the healthcare sector rather than testing causal hypotheses or estimating effect sizes.

The analysis focuses on three core dimensions:

1. **Prevalence of confirmed AI procurement**, measured as the proportion of procurement procedures containing explicitly specified AI functionality relative to the total volume of healthcare procurement notices.
2. **Distribution of AI use cases across application domains**, capturing how confirmed AI procurements are allocated between clinical applications, non-clinical patient services, administrative processes, and research and scientific activities.
3. **Functional concentration within and between domains**, assessed through detailed category-level classification to identify dominant, marginal, and absent AI use cases across organisational functions.

Results are reported using absolute counts and proportional shares, complemented by structured visualisations (e.g. radar charts) to illustrate internal functional profiles. No inferential statistical techniques were applied, as the study does not aim to establish statistical associations, causal relationships, or predictive models.

Analytical scope and exclusions

The unit of analysis in this study is the procurement procedure as an expression of organisational intent, rather than its contractual or economic outcome. Consequently, several procurement attributes were intentionally excluded from the analytical scope.

First, the analysis does not consider the final status of procurement procedures, such as whether a procedure was awarded, annulled, or cancelled. From an analytical perspective, these outcomes are not essential to the study's objective, which is to capture interest in and intended deployment of AI systems, as articulated at the stage of formal procurement documentation.

Second, the study does not analyse the number of bidders, competitive intensity, or awarded contract values. While economically relevant, these variables primarily reflect market structure and pricing dynamics, rather than organisational adoption priorities or functional expectations associated with AI use. Including them would shift the focus from deployment intent to procurement performance, which lies outside the scope of this analysis.

Third, the analysis does not differentiate between procurement procedures based on funding source, contract duration, or supplier characteristics. The emphasis remains on the functional specification of AI systems rather than procurement modality.

Temporal considerations and contextual exclusions

The study period coincides with a phase of heightened public and institutional attention to artificial intelligence, particularly during the last three years of the analysed timeframe. This period is characterised by intensified policy discourse, strategic roadmaps, and media narratives promoting AI as a transformative technology in healthcare.

However, the analytical approach deliberately avoids trend extrapolation or temporal causality claims. The study does not seek to measure growth rates, year-on-year change, or the direct impact of AI-related policy initiatives. Instead, it treats the analysed period as a contextual backdrop against which procurement behaviour is observed, focusing on structural allocation patterns rather than temporal dynamics.

As a result, fluctuations potentially attributable to short-term "AI hype", policy announcements, or vendor-driven marketing cycles are not isolated or modelled. The analysis prioritises stable functional patterns emerging from confirmed procurement specifications, rather than transient signals.

4. Results

4.1. Overall prevalence of AI-related procurement in the healthcare sector

The analysis covered a total of 85,501 healthcare-related public procurement notices published between 2022 and 2024 across the Polish Public Procurement Bulletin (BZP) and the EU Tenders Electronic Daily (TED). Following the multi-stage identification and verification pipeline described in Section 3, 210 procurement procedures were confirmed as containing explicit and verifiable AI functionality at the level of Terms of Reference.

In relative terms, confirmed AI procurements represent approximately 0.25% of all analysed healthcare procurement notices. This proportion remains marginal despite extensive policy attention, funding initiatives, and public discourse positioning artificial intelligence as a strategic driver of healthcare system modernisation.

The sharp reduction from initial "AI potential" signals to confirmed AI cases illustrates the importance of TOR-level verification. A substantial share of procurement notices initially flagged as AI-related did not specify AI functionality in a manner sufficient to meet confirmatory criteria, underscoring the gap between AI-related rhetoric and formal procurement requirements.

Overall, the findings indicate that AI adoption, as reflected in public procurement decisions, remains selective, limited in scale, and highly concentrated, rather than widespread or systemic.

4.2. Distribution of AI use cases across main domains

Table 2 presents the distribution of unique confirmed AI procurement records across four main application domains: clinical applications, non-clinical patient services, administrative processes, and research and scientific applications. The table reports both coverage (the proportion of analysed records in which a given domain appears) and share of total records, reflecting the dominant classification assigned to each procurement.

Table 1. Distribution of confirmed AI procurements by application domain

Domain	Unique Records	Coverage (%)	Share of Total
Clinical Applications	124	79%	64.2%
Non-clinical Patient Services	13	8.3%	6.7%
Administrative Processes	29	18.5%	15%
Research & Scientific	27	17.2%	14%
Total	157	100%	—

Clinical applications clearly dominate the confirmed AI procurement landscape. 124 records (79% coverage) involved clinical AI functionalities, accounting for 64.2% of all confirmed AI procurements. This dominance reflects a strong institutional focus on patient-facing, clinically embedded AI systems, including diagnostic support, medical imaging, and AI-assisted clinical documentation. The high coverage value indicates that clinical AI frequently co-occurs with other AI functionalities within the same procurement procedures.

In contrast, non-clinical patient services represent the smallest segment of the dataset. Only 13 records (8.3% coverage) were classified in this domain, corresponding to 6.7% of the total confirmed AI procurements. These systems primarily address patient registration, navigation, and communication. Despite their comparatively low clinical risk and limited regulatory burden, such applications remain marginal in procurement practice.

AI applied to administrative processes constitutes a more substantial, though still secondary, category. Twenty-nine records (18.5% coverage) were classified as administrative AI, representing 15% of the total dataset. Typical use cases include document workflow automation, intelligent document processing (OCR/IDP), reporting and compliance support, and AI-assisted IT operations. The divergence between coverage and share suggests that administrative AI is often embedded within broader procurement initiatives rather than procured as a standalone strategic investment.

Finally, research and scientific applications account for 27 records (17.2% coverage) and 14% of total confirmed AI procurements. This category includes AI systems supporting clinical research, data harmonisation, predictive modelling, and advanced analytics. Compared with administrative AI, research-oriented procurements display a more balanced relationship between coverage and share, indicating more focused and purpose-specific acquisition patterns.

Taken together, the distribution reveals a pronounced structural imbalance in AI procurement priorities, with clinical applications strongly prioritised over non-clinical domains.

4.3. Functional profile of clinical and non-clinical AI

Beyond the high-level domain distribution presented in Section 4.2, a category-level analysis reveals substantial differences in the internal functional structure of AI procurements across the four main application groups: clinical applications, non-clinical patient services, non-clinical administrative processes, and scientific research applications. This section reports descriptive results only, based strictly on confirmed categories and their proportional shares.

4.3.1. Clinical applications

Clinical AI exhibits a highly concentrated functional profile, dominated by image-based and diagnostic use cases. As shown in Figure 2, the radar chart for clinical applications is characterised by a single dominant axis corresponding to *Imaging and auxiliary diagnostic studies*.

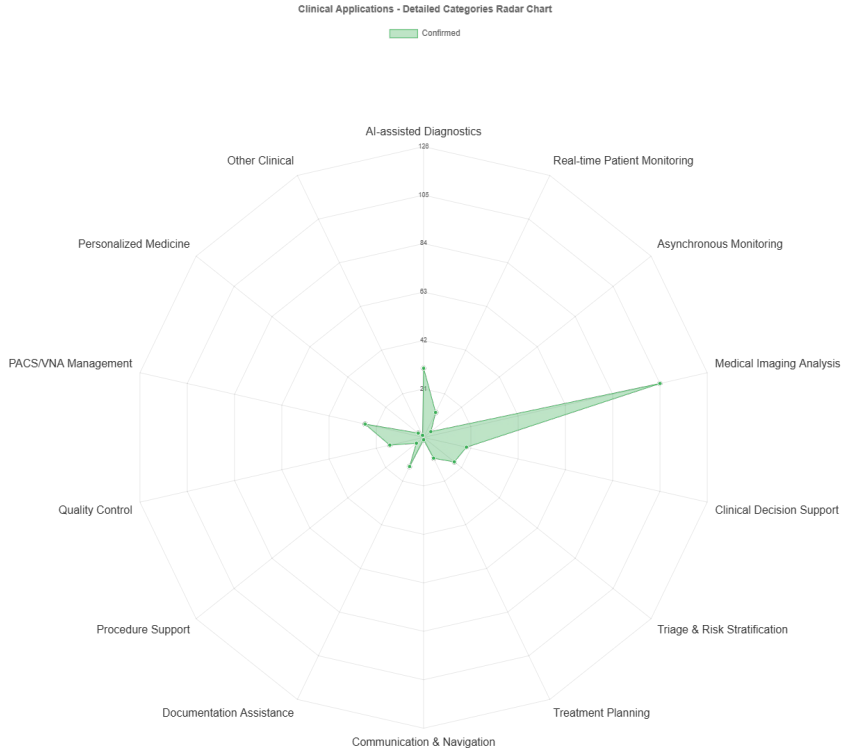

Figure 2. Clinical applications – detailed categories radar chart (confirmed AI cases)

The most prevalent category is Imaging and auxiliary diagnostic studies, with 105 confirmed cases, accounting for 49.5% of all confirmations. This single category represents nearly half of all confirmed AI deployments, indicating a strong concentration of procurement activity in radiology and related diagnostic domains.

Secondary categories include Diagnostics (14.2%), PACS/VNA AI-based image data management (12.3%), and Clinical decision support (9%). Monitoring-related applications are present at lower levels, with Live monitoring (5.7%) and Asynchronous monitoring (1.9%). Advanced clinical use cases such as Personalised medicine and omics analysis (1.4%) and Intraoperative support and navigation (1.9%) remain marginal.

Overall, the clinical AI profile reflects a narrow, image-centric adoption pattern, with limited representation of therapy optimisation, communication support, or system-wide clinical intelligence.

Table 2. Clinical AI – detailed category statistics

Category	Group	Confirmed	% confirmations
Imaging and auxiliary diagnostic studies	Clinical	105	49.5%
Diagnostics	Clinical	30	14.2%
PACS/VNA AI image data management	Clinical	26	12.3%
Clinical decision support	Clinical	19	9%
Triage and risk stratification	Clinical	17	8%
Quality control and error detection	Clinical	15	7.1%
Clinical documentation assistance	Clinical	14	6.6%
Live monitoring	Clinical	12	5.7%
Treatment planning and optimisation	Clinical	10	4.7%
Asynchronous monitoring	Clinical	4	1.9%
Intraoperative support and navigation	Clinical	4	1.9%
Personalised medicine and omics analysis	Clinical	3	1.4%
Clinical communication and patient navigation	Clinical	1	0.5%
Other clinical (unspecified)	Clinical	1	0.5%

4.3.2. Non-clinical patient services

Non-clinical patient-facing AI represents the smallest and least diversified functional group. The radar chart for this domain shows a sharply polarised structure focused almost exclusively on omnichannel communication tools.

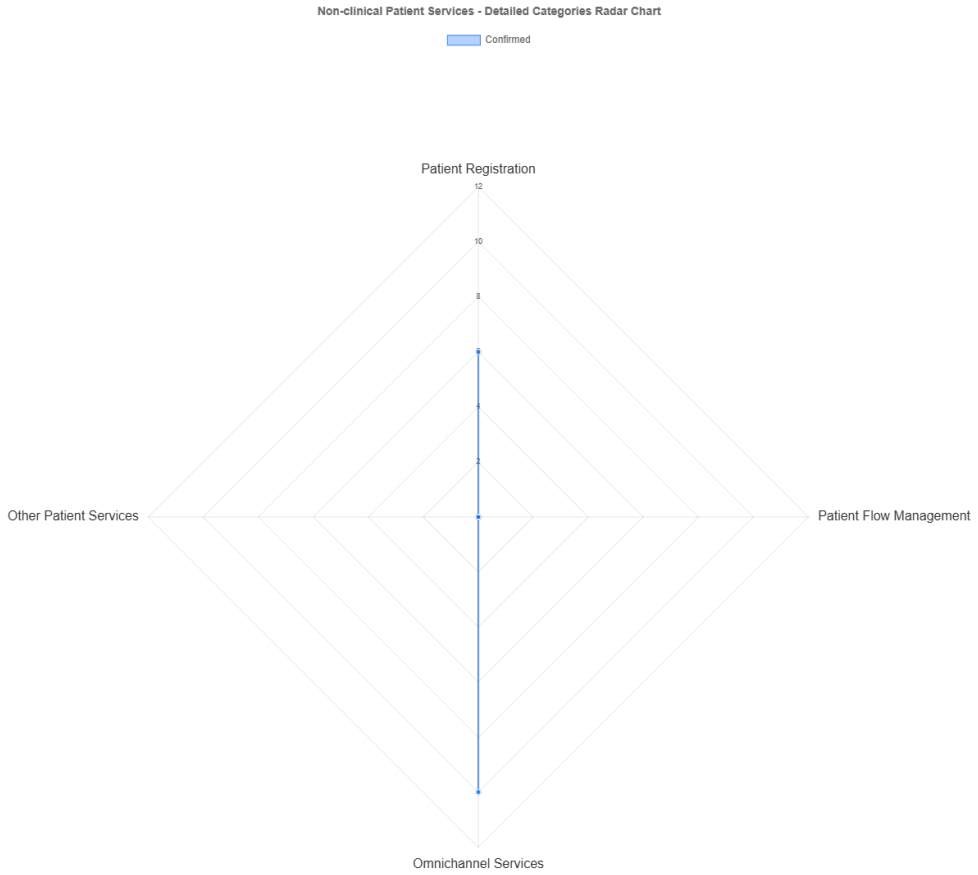

Figure 3. Non-clinical patient services – detailed categories radar chart (confirmed AI cases)

The dominant category is Omnichannel patient services, with 10 confirmed cases representing 4.7% of all confirmations. Patient registration systems appear in 6 cases (2.8%), while Patient flow management and other patient-service categories show no confirmed AI deployments.

This distribution indicates that AI in patient services is primarily used to mediate communication channels, rather than to support patient logistics, flow optimisation, or service coordination.

Table 3. Non-clinical patient services – detailed category statistics

Category	Group	Confirmed	% confirmations
Omnichannel patient services	Non-clinical patient services	10	4.7%
Patient registration	Non-clinical patient services	6	2.8%
Patient flow management	Non-clinical patient services	0	0%
Other patient services (unspecified)	Non-clinical patient services	0	0%

4.3.3. Non-clinical administrative processes

Non-clinical administrative AI displays the widest functional dispersion among all analysed groups. The corresponding radar chart shows multiple moderate peaks rather than a single dominant category.

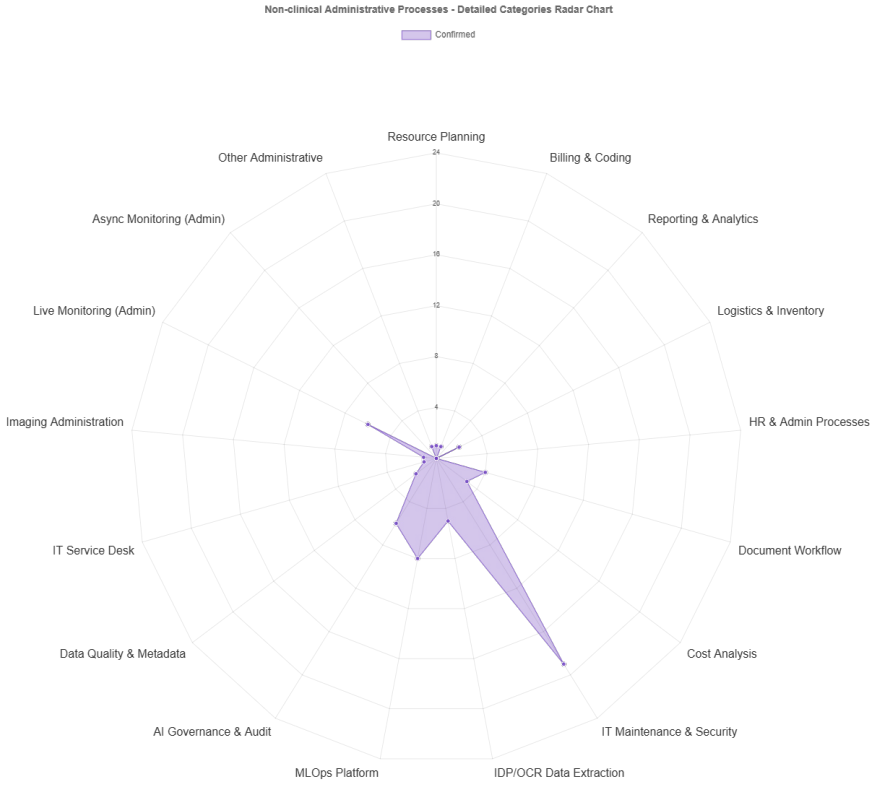

Figure 4. Non-clinical administrative processes – detailed categories radar chart (confirmed AI cases)

The most prominent category is IT infrastructure maintenance and AI-supported security, with 19 confirmed cases, accounting for 9% of all confirmations. Other recurring categories include AI/MLOps platforms and model operationalisation (3.8%), Governance, audit and AI monitoring (2.8%), Live administrative monitoring (2.8%), and IDP/OCR-based document data extraction (2.4%).

In contrast, categories associated with strategic management and organisational decision support—such as cost analysis, resource planning, or billing and coding—are either marginal (≤1.4%) or absent. This suggests that administrative AI is predominantly deployed as technical and infrastructural support, rather than as a tool for managerial analytics.

Table 4. Non-clinical administrative processes – detailed category statistics

Category	Group	Confirmed	% confirmations
IT infrastructure maintenance / AI security	Non-clinical administrative	19	9%
AI/MLOps platforms and model operationalisation	Non-clinical administrative	8	3.8%
Governance, audit and AI monitoring	Non-clinical administrative	6	2.8%
Live monitoring (administrative)	Non-clinical administrative	6	2.8%
IDP/OCR document data extraction	Non-clinical administrative	5	2.4%
Document workflow and correspondence	Non-clinical administrative	4	1.9%
Cost analysis and decision support	Non-clinical administrative	3	1.4%
Logistics and inventory management	Non-clinical administrative	2	0.9%
Data quality and metadata management	Non-clinical administrative	2	0.9%
Resource planning	Non-clinical administrative	1	0.5%
Billing and coding	Non-clinical administrative	1	0.5%
IT service desk support	Non-clinical administrative	1	0.5%
Image processing (administrative)	Non-clinical administrative	1	0.5%
Other administrative (unspecified)	Non-clinical administrative	1	0.5%

4.3.4. Scientific and research applications

Scientific AI applications show a moderately diversified profile, with a strong emphasis on data governance, compliance, and research infrastructure rather than experimental clinical AI.

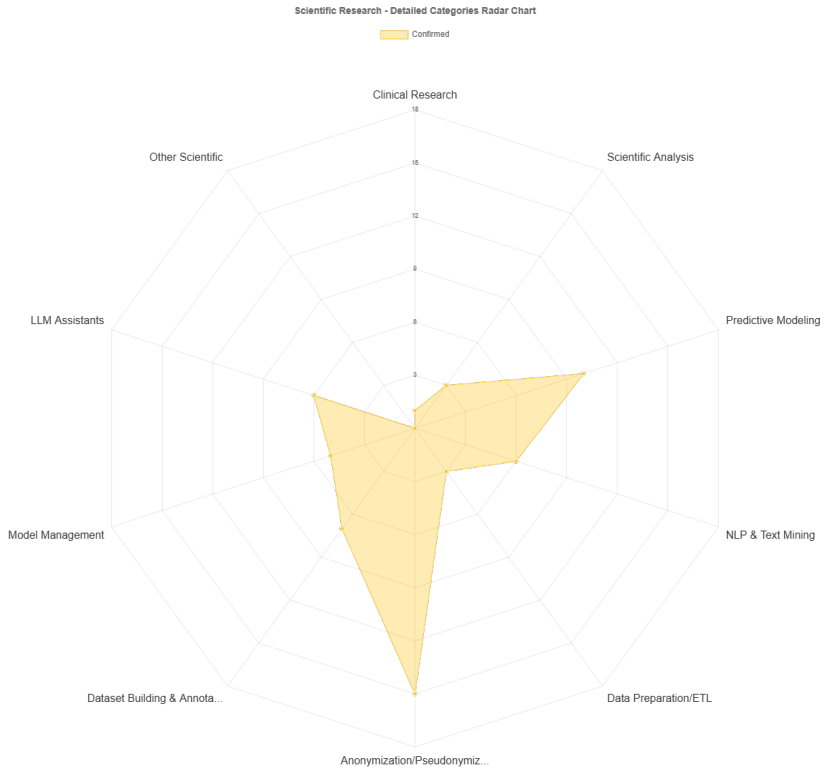

Figure 5. Scientific and research applications – detailed categories radar chart (confirmed AI cases)

The largest category is Anonymisation, pseudonymisation and secure data sharing, with 15 confirmed cases (7.1%), reflecting the regulatory requirements associated with secondary use of health data. Other relevant categories include Predictive modelling (4.7%), Dataset building and annotation (3.3%), LLM-based research assistants (2.8%), and NLP and text mining (2.8%).

This profile indicates that AI in the scientific domain is primarily deployed as enabling infrastructure for compliant research workflows, rather than as a direct driver of AI-led discovery.

Table 5. Scientific and research applications – detailed category statistics

Category	Group	Confirmed	% confirmations
Anonymisation / pseudonymisation and secure sharing	Scientific	15	7.1%
Predictive modelling	Scientific	10	4.7%
Dataset building and annotation	Scientific	7	3.3%
LLM research assistants	Scientific	6	2.8%
NLP and text mining	Scientific	6	2.8%
Model management and reuse	Scientific	5	2.4%
Scientific data analysis	Scientific	3	1.4%
Data preparation and ETL	Scientific	3	1.4%
Clinical research support	Scientific	1	0.5%
Other scientific (unspecified)	Scientific	0	0%

5. Discussion

The central finding emerging from the analysis of public procurement is a marked disproportion in implementation priorities favouring clinical applications at the expense of administrative innovation. Clinical applications accounted for 64.2% of all confirmed AI-related public procurements in the studied sample, whilst administrative and operational processes represented merely 15%. This dominance confirms the thesis that political and commercial narratives are disproportionately focused on clinical domains.

Interpreting the Imbalance Between Clinical and Non-Clinical AI

The observed disparity between clinical and non-clinical AI adoption in Polish healthcare procurement reveals a pattern consistent with broader international trends, yet amplified by specific national policy mechanisms. The 64.2% concentration in clinical applications, particularly the 49.5% devoted to imaging and auxiliary diagnostic studies, mirrors the global healthcare AI landscape where radiology and diagnostic imaging dominate research, regulatory approvals, and commercial attention (George et al., 2023; Sathya, 2024). This alignment suggests that Polish procurement patterns are shaped not only by local factors but also by international market dynamics and the maturity gradient of AI technologies across different healthcare domains.

However, the magnitude of this imbalance in Poland may have been significantly influenced by targeted funding mechanisms, most notably the Krajowy Plan Odbudowy (KPO; National Recovery Plan). Within the KPO framework, substantial resources were earmarked for digital health transformation, including initiatives that explicitly reference the implementation of certified AI tools alongside the expansion and integration of electronic medical documentation in hospitals (Ministerstwo Funduszy i Polityki Regionalnej, 2022; Ministerstwo Zdrowia, 2025). More broadly, the KPO governance and implementation architecture underscores how eligibility criteria and programme design can steer institutional priorities, even where multiple domains compete for attention (European Parliamentary Research Service, 2022).

Notably, whilst the KPO also allocated resources for back-office functions, these funding streams did not impose equivalent AI implementation requirements. This asymmetry in policy design may have inadvertently reinforced the clinical-administrative divide observed in procurement data. Healthcare institutions, responding rationally to funding incentives, directed their AI investments towards domains where both financial support and compliance expectations converged, namely, clinically visible applications that can be readily framed in terms of patient benefit and system modernisation (European Parliamentary Research Service, 2022; Ministerstwo Zdrowia, 2025). The absence of similar AI-specific mandates for administrative functions, despite available funding, likely contributed to the underrepresentation of back-office AI solutions in the procurement landscape.

This funding-driven explanation is further supported by the temporal concentration of procurements within the study period (2022–2024), which coincides with the active implementation phase of KPO-funded healthcare modernisation initiatives. The policy architecture of the KPO thus emerges as a critical contextual factor that may have amplified the natural tendency towards clinical AI prioritisation, creating a structural bias in the procurement ecosystem that extends beyond purely technical or clinical considerations (European Parliamentary Research Service, 2022; Ministerstwo Funduszy i Polityki Regionalnej, 2022).

The dominance of imaging applications within the clinical domain reflects the technological maturity and regulatory acceptance of AI in radiology. Deep learning for image analysis, automated detection, and workflow optimisation represents the most mature clinical AI domain, with established validation frameworks and clear clinical utility (European Society of Radiology, 2025; Maleki Varnosfaderani & Forouzanfar, 2024). This maturity translates into lower perceived risk for procurement officers and clinical decision-makers, creating a self-reinforcing cycle where proven

technologies attract further investment whilst less mature applications struggle to gain traction. The concentration of 105 cases (49.5% of all AI procurements) in imaging and auxiliary diagnostic studies demonstrates how technological readiness interacts with policy incentives to shape adoption patterns.

In contrast, the 15% allocation to administrative processes represents a substantial underutilisation of AI potential in operational domains. The literature consistently identifies administrative and back-office applications, including patient demand forecasting, workforce planning, scheduling, billing automation, and documentation assistance, as high-impact, comparatively lower-risk use cases that offer clear return on investment and face fewer regulatory hurdles than clinical applications (Chowdhury Urbi & Gazi Tiva, 2025; Sachdeva & Jain, 2025). The observed procurement patterns suggest that despite these advantages, administrative AI remains overshadowed by the prestige and policy attention accorded to clinical innovations. This represents a significant opportunity cost, as administrative AI can deliver efficiency gains with shorter implementation timelines and lower validation requirements than many clinical applications (Davenport & Glaser, 2022).

The research and scientific applications domain, accounting for 14% of procurements, reveals an interesting intermediate category that bridges clinical and operational concerns. The prominence of anonymisation and pseudonymisation solutions (7.1% of total procurements) reflects growing awareness of data governance requirements, particularly in anticipation of stricter regulatory frameworks. However, the relatively modest investment in predictive modelling (4.7%) and dataset building (3.3%) suggests that Polish healthcare institutions are still developing the infrastructure and capabilities necessary for advanced analytics and continuous learning systems (Ali et al., 2023; Esmaeilzadeh, 2024).

Implications for Healthcare AI Governance and Policy

The procurement patterns documented in this study illuminate several critical governance challenges that extend beyond the Polish context. The 0.25% prevalence of AI-related procurements among all healthcare tenders indicates that despite considerable policy rhetoric surrounding AI as a strategic priority, actual adoption remains marginal and highly selective. This gap between discourse and implementation is consistent with reviews that emphasise a persistent research-to-deployment chasm, where extensive research activity translates into relatively limited routine integration (Ali et al., 2023; Esmaeilzadeh, 2024).

The concentration of AI investment in specific domains raises questions about the alignment between procurement decisions and healthcare system priorities. Whilst diagnostic imaging undoubtedly benefits from AI augmentation, the neglect of administrative automation may perpetuate operational inefficiencies that ultimately constrain clinical capacity. The literature emphasises that administrative AI applications often offer clearer operational business cases than many clinical applications, yet they may suffer from lower visibility and prestige within institutional innovation narratives (Davenport & Glaser, 2022; Sachdeva & Jain, 2025). This suggests that procurement strategies driven primarily by clinical prestige and external funding incentives may not optimally address the multidimensional challenges facing healthcare systems.

The forthcoming implementation of the EU Artificial Intelligence Act introduces a new layer of complexity that will disproportionately affect clinical AI applications. The Act classifies many healthcare AI systems as high-risk, imposing lifecycle obligations for development, validation, market placement, and post-market surveillance (European Parliament & Council of the European Union, 2024; van Kolfschooten & van Oirschot, 2024). These requirements will likely increase the compliance burden and time-to-deployment for clinical AI, potentially exacerbating existing implementation barriers related to regulatory fragmentation, technical integration, and evidence generation (Shah et al., 2025; van Kolfschooten & van Oirschot, 2024). In contrast, many administrative AI applications will fall into lower-risk categories, suggesting that regulatory dynamics may eventually help rebalance the clinical-administrative divide, though only if procurement strategies and funding mechanisms adapt accordingly (European Society of Radiology, 2025).

The role of funding architecture in shaping AI adoption patterns cannot be overstated. The KPO example demonstrates how policy design, specifically, the coupling of funding with explicit implementation expectations in certain domains but not others, can create structural incentives that override purely clinical or operational considerations. Future policy interventions aimed at promoting more balanced AI adoption should consider symmetric incentive structures that encourage innovation across both clinical and administrative domains. This might include targeted grants for operational AI, procurement rules favouring demonstrable return on investment in back-office automation, and reimbursement schemes that reward efficiency improvements alongside clinical outcomes (Davenport & Glaser, 2022; Taheri Hosseinkhani, 2025).

Barriers to Implementation and the Clinical-Administrative Divide

The observed procurement patterns reflect not only funding incentives but also fundamental differences in the implementation barriers facing clinical versus administrative AI applications. Clinical AI deployment confronts a convergence of regulatory complexity, technical integration challenges, and validation requirements that collectively explain why many promising models fail to reach routine care (Ali et al., 2023; Esmaeilzadeh, 2024; Shah et al., 2025).

Regulatory fragmentation remains a persistent obstacle for clinical AI. The patchwork of national, regional, and sectoral rules creates uncertainty around evidence thresholds, liability expectations, and post-market monitoring obligations (European Parliament & Council of the European Union, 2024; Shah et al., 2025). This regulatory complexity increases the perceived risk of clinical AI procurements, particularly for hospital administrators who must navigate unclear compliance pathways. The literature documents widespread calls for health-specific guidance under the EU AI framework to operationalise legal obligations and reduce regulatory uncertainty (European Society of Radiology, 2025; van Kolfschooten & van Oirschot, 2024).

Technical integration represents an equally significant barrier. Hospital information technology infrastructure is often fragmented, with siloed electronic health record systems, incompatible data formats, and insufficient platform architecture to support seamless AI deployment (Ali et al., 2023; Esmaeilzadeh, 2024; Maimaitiaili et al., 2025). These constraints impede data flows and workflow integration, making many AI models impractical to embed in everyday care despite their demonstrated performance in controlled research settings. The integration challenge is particularly acute for clinical AI, which must interface with complex care pathways, real-time clinical workflows, and multiple legacy systems simultaneously.

Validation and evidence generation requirements further differentiate clinical from administrative AI. Clinical applications demand robust external validation, prospective clinical trials, and ongoing post-market surveillance to establish safety and efficacy, requirements that are less stringent for operational tools (European Society of Radiology, 2025; Maleki Varnosfaderani & Forouzanfar, 2024). Reviews consistently emphasise the need for standardised validation frameworks and continuous evidence generation before broad clinical adoption, reflecting appropriate caution given the direct patient safety implications. However, these rigorous validation requirements extend development timelines and increase costs, creating additional barriers to clinical AI procurement and deployment.

In contrast, administrative AI applications face a different risk profile centred on privacy, organisational disruption, and economic implications rather than direct patient safety (Chowdhury Urbi & Gazi Tiva, 2025; Sachdeva & Jain, 2025). This distinction in risk characteristics suggests that administrative AI could be deployed more rapidly and with lower validation costs, yet the procurement data reveal that this theoretical advantage has not translated into proportionate adoption. The underutilisation of administrative AI thus appears to reflect not technical or regulatory constraints but rather organisational priorities, prestige considerations, and, as discussed above, the asymmetric structure of funding incentives.

Economic and incentive misalignments further explain the clinical-administrative divide. Payment models and reimbursement mechanisms in many healthcare systems reward clinical interventions more directly than operational efficiencies, reducing provider willingness to invest in administrative automation even when return on investment is demonstrable (Davenport & Glaser, 2022; Taheri Hosseinkhani, 2025).

Positioning Within the Broader Healthcare AI Landscape

The findings of this study contribute to an emerging body of evidence documenting the gap between AI research activity and real-world healthcare deployment. Whilst the volume of published AI research continues to expand exponentially, reviews consistently reveal that few models progress beyond pilot studies or vendor demonstrations to achieve routine clinical integration (Ali et al., 2023; Esmaeilzadeh, 2024). The 0.25% prevalence of AI-related procurements observed in this study provides empirical confirmation of this implementation gap from a procurement perspective, complementing existing evidence from clinical audits and deployment surveys.

The dominance of diagnostic imaging in healthcare AI adoption is well-established internationally, driven by the convergence of technological maturity, clear clinical utility, regulatory acceptance, and commercial investment (George et al., 2023; Maleki Varnosfaderani & Forouzanfar, 2024; Sathya, 2024). However, the magnitude of imaging dominance observed in Polish procurement data, with imaging and auxiliary diagnostics accounting for nearly half of all AI procurements, suggests that local factors, including KPO funding structure and associated implementation expectations, may have amplified a global trend to create an even more pronounced concentration in Poland (European Parliamentary Research Service, 2022; Ministerstwo Funduszy i Polityki Regionalnej, 2022).

The relative neglect of administrative AI applications represents a missed opportunity that is increasingly recognised in the literature. Operational domains such as scheduling, resource allocation, billing automation, and workforce planning offer tractable problems with clear metrics, shorter feedback loops, and lower regulatory barriers than many clinical applications (Chowdhury Urbi & Gazi Tiva, 2025; Davenport & Glaser, 2022; Sachdeva & Jain, 2025). Several analyses argue that reallocating attention and resources towards these domains could yield substantial efficiency gains and free clinical capacity for higher-value activities. The procurement patterns documented here suggest that such reallocation has not yet occurred in Polish healthcare, despite the theoretical advantages of administrative AI.

The research and scientific applications domain, whilst representing only 14% of procurements, merits particular attention as an enabler of future AI capabilities. Investment in data infrastructure, anonymisation technologies, predictive modelling platforms, and research support tools creates the foundation for more sophisticated AI applications across both clinical and administrative domains (Ali et al., 2023; Esmaeilzadeh, 2024; Maimaitiaili et al., 2025). The modest but consistent presence of these procurements suggests emerging recognition of the importance of data governance and analytical infrastructure, though the scale of investment remains limited relative to the transformative potential of advanced analytics in healthcare.

6. Conclusions

This study shows that AI-related public procurement in the Polish healthcare system is both rare and strongly skewed towards clinical applications, especially diagnostic imaging. The observed structure suggests that adoption is shaped less by a balanced assessment of system-wide needs and more by the combined force of technological maturity, visibility and prestige of clinical tools, and the design of public funding incentives. As a result, administrative and operational AI, despite offering faster implementation pathways and potentially high efficiency gains, remains underrepresented.

The findings point to a governance challenge: if procurement continues to prioritise clinically "showcase" solutions while back-office automation is treated as secondary, system-level constraints such as staffing, throughput, scheduling, and documentation burdens may persist and continue to limit clinical capacity. A more effective trajectory would treat AI as a portfolio of interventions across the full value chain of

care, with intentional balancing between clinical innovation, operational efficiency, and investments in data infrastructure that enable long-term learning and analytics.

Ultimately, the results underline that procurement and funding architecture can create structural bias in technology adoption. If policy instruments tie resources to specific types of AI while leaving other domains without comparable incentives, institutions will rationally follow the strongest signals. Achieving more sustainable and system-relevant AI uptake will require symmetric incentives, clearer implementation pathways, and procurement strategies that reward measurable operational impact alongside clinical benefits.

References

Aboy, M., Minssen, T., & Vayena, E. (2024). Navigating the EU AI Act: Implications for regulated digital medical products. npj Digital Medicine, 7(1), Article 237.https://doi.org/10.1038/s41746-024-01232-3

Alami, H., Lehoux, P., Auclair, Y., de Guise, M., Gagnon, M.-P., Shaw, J., Fortin, J.-P., Fleet, R., Ag Ahmed, M. A., Zinszer, K., & Savoldelli, M. (2020). Artificial intelligence and health technology assessment: Anticipating a new level of complexity. *Journal of Medical Internet Research*, *22*(7), e17707. https://doi.org/10.2196/17707

Ali, O. A., Abdelbaki, W., Shrestha, A., Elbasi, E., & Alryalat, M. A. A. (2023). A systematic literature review of artificial intelligence in the healthcare sector: Benefits, challenges, methodologies, and functionalities. *Journal of Innovation & Knowledge, 8*(1), 100333. https://doi.org/10.1016/j.jik.2023.100333

Almagadi, M. A. S., Banday, A. H., Almagadi, R. A., Adnan, N. M., & Alharbi, N. N. (2025). Enhancing administrative performance in the health sector through artificial intelligence: Pathways to operational efficiency and service quality. Journal of International Crisis and Risk Communication Research, 8(2), 194-201. https://doi.org/10.63278/jicrcr.vi.3137

Alves, M., Seringa, J., Silvestre, T., & Magalhães, T. (2024). Use of artificial intelligence tools in supporting decision-making in hospital management (Version 1) [Preprint]. Research Square. https://doi.org/10.21203/rs.3.rs-4491119/v1

Bartusek, M., & Kulawik, A. (2021). Analiza potrzeb zastosowania nowoczesnej technologii i sztucznej inteligencji w sektorze ochrony zdrowia [Analysis of the needs for modern technology and artificial intelligence application in the healthcare sector]. Forum Rozwoju Przedsiębiorczości. https://doi.org/10.56583/frp.633

Boverhof, B.-J., Redekop, W. K., Visser, J. J., de Feijter, J. M., & Majoie, C. B. L. M. (2024). Broadening the HTA of medical AI: A review of the literature to inform a tailored approach. Health Policy and Technology, 13(3), 100868. https://doi.org/10.1016/j.hlpt.2024.100868

Bukowski, M., Farkas, R., Beyan, O., et al. (2020). Implementation of eHealth and AI integrated diagnostics with multidisciplinary digitized data: Are we ready from an international perspective? European Radiology, 30(12), 6893-6902. https://doi.org/10.1007/s00330-020-06874-x

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., & Yao, S. (2022). Sources of bias in artificial intelligence that perpetuate healthcare disparities—A global review. *PLOS Digital Health*, *1*(3), e0000022. https://doi.org/10.1371/journal.pdig.0000022

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

Chomutare, T., Alonso Tejedor, M. A., Olsen Svenning, T., et al. (2022). Artificial intelligence implementation in healthcare: A theory-based scoping review of barriers and facilitators. International Journal of Environmental Research and Public Health, 19(23), 16359. https://doi.org/10.3390/ijerph192316359

Chowdhury Urbi, S. R., & Gazi Tiva, M. (2025). Technology and innovation in healthcare: Adoption of AI and predictive analytics in hospital management. https://doi.org/10.69937/pf.por.3.2.52

Coiera, E., & Liu, S. (2022). Evidence synthesis, digital scribes, and translational challenges for artificial intelligence in healthcare. Cell Reports Medicine, 3(12), 100860. https://doi.org/10.1016/j.xcrm.2022.100860

Davenport, T. H., & Glaser, J. (2022). Factors governing the adoption of artificial intelligence in healthcare providers. *Discover Health Systems, 1*(1), 13. https://doi.org/10.1007/s44250-022-00004-8

Esmaeilzadeh, P. (2024). Challenges and strategies for wide-scale artificial intelligence deployment in healthcare practices: A perspective for healthcare organizations. *Artificial Intelligence in Medicine, 156*, 102861. https://doi.org/10.1016/j.artmed.2024.102861

European Commission. (n.d.). *AI Act Service Desk: Article 9 (Risk management system).* Retrieved December 16, 2025, from https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-9

European Commission. (n.d.). *AI Act Service Desk: Article 12 (Record-keeping).* Retrieved December 16, 2025, from https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-12

European Commission. (n.d.). *AI Act Service Desk: Article 13 (Transparency and provision of information to deployers).* Retrieved December 16, 2025, from https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-13

European Commission. (n.d.). *AI Act Service Desk: Article 14 (Human oversight).* Retrieved December 16, 2025, from https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-14

European Commission. (n.d.). *AI Act Service Desk: Article 15 (Accuracy, robustness and cybersecurity).* Retrieved December 16, 2025, from https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-15

European Commission. (n.d.). *AI Act Service Desk: Article 27 (Fundamental rights impact assessment).* Retrieved December 16, 2025, from https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-27

European Parliament and the Council of the European Union. (2024). *Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union*, L 2024/1689. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401689

European Parliamentary Research Service. (2022). *Krajowy Plan Odbudowy i Zwiększania Odporności Polski* [Briefing]. European Parliament. https://www.europarl.europa.eu/RegData/etudes/BRIE/2022/733665/EPRS-Briefing-733665-NRRP-Poland-PL.pdf

European Society of Radiology. (2025). Guiding AI in radiology: ESR's recommendations for effective implementation of the European AI Act. *Insights into Imaging*. https://doi.org/10.1186/s13244-025-01905-x

Farah, L., Borget, I., Martelli, N., et al. (2024). Suitability of the current health technology assessment of innovative artificial intelligence-based medical devices: Scoping literature review. Journal of Medical Internet Research, 26, e51514. https://doi.org/10.2196/51514

George, A. S. H., Shahul, A., & George, A. S. (2023). *Artificial intelligence in medicine: A new way to diagnose and treat disease* [Preprint]. Zenodo. https://doi.org/10.5281/zenodo.8374066

Glinkowski, W., Cedro, T., Wołk, A., et al. (2025). Telemedicine, eHealth, and digital transformation in Poland (2014-2024): Trends, specializations, and systemic implications. Applied Sciences, 15(16), 8793. https://doi.org/10.3390/app15168793

Gomez Rossi, J., et al. (2022). Approaches to the economic evaluation of artificial intelligence in health care: A scoping review. Journal of Medical Internet Research, 24(5), e35603. https://doi.org/10.2196/35603

Horgan, D., Romão, M., Morré, S. A., & Kalra, D. (2019). Artificial intelligence: Power for civilisation—and for better healthcare. *Public Health Genomics*, *22*(5-6), 145-161. https://doi.org/10.1159/000504785

Jakubek-Lalik, J. (2024). The challenges of AI in administrative law and the need for specific legal remedies: Analysis of Polish regulations and practice. Central European Public Administration Review, 22(2), 85-106. https://doi.org/10.17573/cepar.2024.2.05

Khan, S. D., Hoodbhoy, Z., Raja, M. H. R., Kim, J. Y., Hogg, H. D. J., Manji, A. A. A., Gulamali, F., Hasan, A., Shaikh, A., Tajuddin, S., Khan, N. S., Patel, M. R., Balu, S., Samad, Z., & Sendak, M. P. (2024). Frameworks for procurement, integration, monitoring, and evaluation of artificial intelligence tools in clinical settings: A systematic review. PLOS Digital Health, 3(5), e0000514. https://doi.org/10.1371/journal.pdig.0000514

Kolfschooten, H., & van Oirschot, J. (2024). The EU Artificial Intelligence Act (2024): Implications for healthcare. Health Policy, 149, 105152. https://doi.org/10.1016/j.healthpol.2024.105152

Leo, C. G., Tumolo, M. R., Sabina, S., Bianco, A., Bianchi, F. P., Nobile, C. G. A., & Pavia, M. (2022). Health technology assessment for in silico medicine: Social, ethical and legal aspects. *Interna-*

tional Journal of Environmental Research and Public Health, *19*(3), 1510. https://doi.org/10.3390/ijerph19031510

Li, Q. (2022). *The regulation of medical AI: Policy approaches, data, and innovation incentives* (NBER Working Paper No. 30639). National Bureau of Economic Research. https://doi.org/10.3386/w30639

Maimaitiaili, M., Jiamaliding, J., Dai, W., Xiao, Y., & Kuerbanjiang, K. (2025). Artificial intelligence platform architecture for hospital systems: Systematic review. *Journal of Medical Internet Research, 27*, e79788. https://doi.org/10.2196/79788

Maleki Varnosfaderani, S., & Forouzanfar, M. (2024). The role of AI in hospitals and clinics: Transforming healthcare in the 21st century. *Bioengineering, 11*(4), 337. https://doi.org/10.3390/bioengineering11040337

Martin, J., Hurcum, Z., Cross, S., Sivarajah, R., Papoutsaki, M., Adams, E., Peplinski, J., Newby, T., & Clark, J. (2025). Increasing MRI capacity at a clinical diagnostic centre and a trauma hospital using artificial intelligence-based image reconstruction (AI-IR): A quality improvement project using the Model for Improvement framework. BMJ Open Quality, 14(4), e003470. https://doi.org/10.1136/bmjoq-2025-003470

McKee, M., & Wouters, O. J. (2022). The challenges of regulating artificial intelligence in healthcare. *International Journal of Health Policy and Management*, *11*(9), 1966-1968. https://doi.org/10.34172/ijhpm.2022.7261

Messmann, H., Bisschops, R., Antonelli, G., Libânio, D., Sinonquel, P., Abdelrahim, M., Ahmad, O. F., Areia, M., Bergman, J. J. G. H. M., Bhandari, P., Boskoski, I., Dekker, E., Domagk, D., Ebigbo, A., Eelbode, T., Eliakim, R., Häfner, M., Haidry, R. J., Jover, R., … Dinis-Ribeiro, M. (2022). Expected value of artificial intelligence in gastrointestinal endoscopy: European Society of Gastrointestinal Endoscopy (ESGE) position statement. Endoscopy, 54(12), 1211-1231. https://doi.org/10.1055/a-1950-5694

Ministerstwo Funduszy i Polityki Regionalnej. (2022). *Krajowy Plan Odbudowy i Zwiększania Odporności*. Government of Poland. https://www.funduszeeuropejskie.gov.pl/media/109762/KPO.pdf

Ministerstwo Zdrowia. (2025, April 11). *Ruszył konkurs na cyfrową transformację ochrony zdrowia*. Gov.pl. https://www.gov.pl/web/zdrowie/ruszyl-konkurs-na-cyfrowa-transformacje-ochrony-zdrowia

Ng Kok Wah, J. (2025). AI-driven eHealth technologies revolution: A novel review of emerging digital healthcare innovations and their transformative impact on global healthcare systems. Global Health Synapse, 2(1), 55-79. https://doi.org/10.63456/ghs-2-1-14

Parikh, R. B., & Helmchen, L. A. (2022). Paying for artificial intelligence in medicine. npj Digital Medicine, 5, 63. https://doi.org/10.1038/s41746-022-00609-6

Pashkov, V., Harkusha, A. O., & Harkusha, Y. O. (2020). Artificial intelligence in medical practice: Regulative issues and perspectives. *Wiadomości Lekarskie*, *73*(12), 2722-2727. https://doi.org/10.36740/WLEK202012204

Popescu, E.-R., Geanta, M., & Brand, A. (2022). Mapping of clinical research on artificial intelligence in the treatment of cancer and the challenges and opportunities underpinning its integration in the European Union health sector. *European Journal of Public Health*, *32*(3), 389-395. https://doi.org/10.1093/eurpub/ckac016

Ramsay, A. I. G., Crellin, N., Lawrence, R., Walton, H., Bagri, S., Dodsworth, E., Elphinstone, J., Gleeson, J., Halliday, L., Herbert, H., Lloyd, A., Massou, E., Mehta, N., Morris, S., Ng, A., O'Regan, C., Sherlaw-Johnson, C., & Fulop, N. J. (2025). Procurement and early deployment of artificial intelligence-based chest diagnostics in the English National Health Service: A retrospective, multicentre analysis. EClinicalMedicine, 89, 103481. https://doi.org/10.1016/j.eclinm.2025.103481

Sachdeva, C., & Jain, P. K. (2025). AI-driven innovations in healthcare administration: Streamlining processes for improved operational efficiency. *International Journal for Multidisciplinary Research, 7*(3). https://doi.org/10.36948/ijfmr.2025.v07i03.37788

Sathya, M. (2024). Shaping the future of healthcare with artificial intelligence: Current trends and beyond. *African Journal of Biomedical Research, 27*(4s). https://doi.org/10.53555/ajbr.v27i4s.5590

Sciarretta, E., Mancini, R., & Greco, E. (2022). Artificial intelligence for healthcare and social services: Optimizing resources and promoting sustainability. Sustainability, 14(24), 16464. https://doi.org/10.3390/su142416464

Shah, A., Mitchell, S. F., Coad, G., & Michels, D. (2025). Safe and responsible use of artificial intelligence in health care: Current regulatory landscape and considerations for regulatory policy. *JCO Clinical Cancer Informatics, 9*, e2500123. https://doi.org/10.1200/cci-25-00123

Shaw, J., Rudzicz, F., Jamieson, T., & Goldfarb, A. (2019). Artificial intelligence and the implementation challenge. *Journal of Medical Internet Research*, *21*(7), e13659. https://doi.org/10.2196/13659

Taheri Hosseinkhani, N. (2025). Economic evaluation of artificial intelligence integration in global healthcare: Balancing costs, outcomes, and investment value. *OSF Preprints*. https://doi.org/10.31219/osf.io/6k3bx_v1

Urbi, S. R. C., & Tiva, M. G. (2025). Technology and innovation in healthcare: Adoption of AI and predictive analytics in hospital management. Pathfinder of Research, 3(2), 22-45. https://doi.org/10.69937/pf.por.3.2.52

van Kolfschooten, H., & van Oirschot, J. (2024). The EU Artificial Intelligence Act (2024): Implications for healthcare. *Health Policy, 149*, 105152. https://doi.org/10.1016/j.healthpol.2024.105152

Weinert, A., et al. (2022). Perspective of information technology decision makers on factors influencing adoption and implementation of artificial intelligence technologies in 40 German hospitals: Descriptive analysis. JMIR Medical Informatics, 10(6), e34678. https://doi.org/10.2196/34678

Wolff, J., Pauling, J. K., Keck, A., & Baumbach, J. (2021). Success factors of artificial intelligence implementation in healthcare. *Frontiers in Digital Health*, *3*, 594971. https://doi.org/10.3389/fdgth.2021.594971

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance.* World Health Organization. https://www.who.int/publications/i/item/9789240029200