

Immune System Intelligence as a Blueprint for AI-Based Disease Detection

Cynthia Silvia¹, Karolina Kopczynski², Zoe Auezov³

¹American Public University System, School of Business American Public University, Charlestown, WV, USA

²American Public University System, School of Arts, Humanities and Education, American Public University, Charlestown, WV, USA

³Laurel Springs School, West Chester, PA, USA

Email: cynthia.silvia@mycampus.apus.edu, karolina.kopczynski@mycampus.apus.edu, zoeauezov@gmail.com

How to cite this paper: Silvia, C., Kopczynski, K., & Auezov, Z. (2026). Immune System Intelligence as a Blueprint for AI-Based Disease Detection. *Creative Education, 17*, 1109-1129.
<https://doi.org/10.4236/ce.2026.176067>

Received: April 25, 2026

Accepted: June 26, 2026

Published: June 29, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

The human immune system can be viewed as a conceptual model of pattern recognition, adaptive learning, and anomaly detection, providing a rich theoretical basis for the future development of AI in disease detection. As a framework, not for empirical evidence; instead, it sheds light on how principles of self/non-self discrimination, continuous surveillance, and adaptive responses to emerging threats contribute to the design of AI-based healthcare systems. In this conceptualization, AI is presented not simply as a tool but as an extension of biologically inspired intelligence—that is able to aid in better diagnosis, early detection, and personalized treatment. This perspective emphasizes AI systems' potential to embody traits of adaptability, precision, and resilience through an analogy between biological immune processes and computational architectures. Thus, framing immune system intelligence as a conceptual perspective provides a more holistic view of the direction in which bio-inspired models will guide the development of effective, ethical, and human-centered strategies for disease detection in modern healthcare.

Keywords

Artificial Intelligence, Human Immune System, Disease Detection, Bio-Inspired Models, Pattern Recognition, Adaptive Learning, Anomaly Detection, Personalized Medicine, Human-Centered AI, Healthcare Innovation

1. Introduction

The human immune system is one of the most sophisticated natural models of pattern recognition, anomaly detection, and adaptive learning, making it a valua-

ble blueprint for developing artificial intelligence in healthcare. Through constant monitoring of the body, separating self from non-self, and responding to dynamic threats, the immune system exhibits characteristics similar to those of contemporary AI-based disease classification systems. Advances in artificial intelligence particularly in deep learning have enabled the analysis of complex medical data, improving diagnostic accuracy and early disease identification (Esteva et al., 2019; Yu et al., 2018).

As outlined by Jiang et al. (2017) and Topol (2019), AI technologies are increasingly transforming healthcare by augmenting clinical decision-making and enhancing personalized medicine. Drawing inspiration from immune system intelligence, AI systems can be designed to better recognize subtle patterns, adapt to new data, and respond to anomalies with greater precision. This biological analogy highlights the potential for AI to emulate the efficiency, adaptability, and resilience of immune processes, ultimately advancing more accurate, timely, and human-centered approaches to disease detection. This paper is a narrative and the scope follows a conceptual progression with biological theory, computational abstraction (AIS) and modern AI and continual learning applications.

2. Definitions

Artificial immune systems (AIS) are a class of computational models within biologically inspired computing that draw on principles, mechanisms, and metaphors of the biological immune system to develop algorithms for problem solving, particularly in areas such as pattern recognition, optimization, and anomaly detection (Timmis et al., 2004).

Continual learning refers to the capability of an intelligent system to incrementally acquire, update, and retain knowledge over time from sequentially presented data, allowing it to adapt to changing environments while accumulating experience across tasks (Wang et al., 2015).

Explainable artificial intelligence (XAI) refers to a set of methods and approaches designed to make the decision-making processes of artificial intelligence systems transparent, interpretable, and understandable to human users, thereby enhancing trust and accountability in AI applications (Hamida et al., 2024).

Explainable artificial intelligence (XAI) refers to a set of methods and approaches designed to make the decision-making processes of artificial intelligence systems transparent, interpretable, and understandable to human users, thereby enhancing trust and accountability in AI applications (Hamida et al., 2024).

3. Innate vs. Adaptive Immunity as Models for Layered AI Detection Systems

This section delves into how the immune system's two-tier defense has prompted the development of AI frameworks. The innate and adaptive components of the immune system together provide a logical foundation for multi-layered detection systems. Its innate immunity, coupled with rapid, generalized responses, allows AI to detect more than aberrant appearance. Adaptive immunity provides speci-

ficiency and memory in much the same way that deep learning models do, as prediction becomes more accurate over time. [Medzhitov and Janeway \(2000\)](#) pioneering work demonstrates the influence of innate immunity on adaptive practices and supports a hierarchical AI framework. This dual-layer architectural model provides stability and versatility in both artificial and biological systems.

The immune system comprises a number of highly coordinated cellular, tissue, and molecular systems that collaborate to shield the body from infection by bacteria and other agents. [Chaplin \(2010\)](#) emphasizes that this response paradigm comprises four phases: pathogen recognition, activation, regulation, and resolution. These will help determine whether the immune response is effective and not overly reactive, thereby preventing inflammation of the host's tissues. And the intricacy of these processes highlights the twin imperatives of the immune system: speed and precision.

Innate immunity is the body's most immediate and nonspecific defense and provides immediate protection from invading microorganisms. [Medzhitov and Janeway \(2000\)](#) explain innate immunity as evolutionarily ancient and necessary for recognizing conserved molecular patterns found across wide classes of pathogens. This system depends on physical barriers, phagocytic cells, and pattern recognition receptors that trigger inflammation. Although innate mechanisms act quickly, they are also central to the adaptive immune response that follows, presenting antigens that elicit cytokines that direct lymphocyte activation.

Adaptive immunity, on the other hand, provides high specificity and durable protection. [Murphy et al. \(2016\)](#) claim that adaptive responses are coordinated by B and T lymphocytes, whose clonal growth and differentiation occur upon encountering their respective antigens. This system not only eliminates pathogens efficiently but also creates immunological memory for faster and more resolute responses upon re-exposure. Innate and adaptive immunity, combined, enable the body to respond to both familiar and novel threats.

Collectively, they exemplify the immune system's astonishing ability to integrate fast, unpredictable defenses with targeted, memory-dependent responses. As [Chaplin \(2010\)](#) writes, the coordination between these layers is critical for homeostasis in the face of infection. The pioneering study of [Medzhitov and Janeway \(2000\)](#) and the exhaustive framework that [Murphy et al. \(2016\)](#) emphasize the intertwined role of innate and adaptive systems, rather than as two separate immune responses. The immune system's synchronized fusion of rapid, nonspecific screening and memory-driven specificity in a disease-detection environment can be mapped onto multi-modal AI systems such as those combining real-time biosensor monitoring with longitudinal EHR analysis to enable both early anomaly detection and progressively refined, patient-specific diagnosis. Within a disease-detection setting, the immune system's integrated application of rapid, nonspecific screening and memory-driven specificity can be mapped onto multi-modal AI systems, such as those that combine real-time biosensor monitoring with longitudinal EHR analysis, both to enable early anomaly detection and to progres-

sively refine the diagnosis based on patient-specific factors.

Innate vs. Adaptive Immunity

Innate and adaptive immunity are two distinct but complementary branches of the human immune response. Both of which carry out complementary and yet distinct functions. Innate immunity provides the body's preemptive shield against invading pathogens by means preceding the onset of infection. This primary response is mediated through physical barriers, inflammatory mediators, and cellular defenses that work in tandem to halt threats (Chaplin 2010). Although innate mechanisms are efficient, they do not offer long-term memory and cannot learn from encounters.

Innate immunity is driven by evolutionary-based recognition systems that identify common molecular patterns present across broad classes of microbes. Medzhitov and Janeway (2000) describe these pattern recognition receptors as crucial for initiating inflammation and activating phagocytic cells that engulf and kill pathogens. Early responses shape the later adaptive immune response through antigen presentation and cytokine signaling. This way not only does innate immunity offer immediate refuge, but it also provides a necessary environment for more specific defenses to emerge. Adaptive immunity differs from innate immunity in its capacity for specificity and memory. According to Murphy et al. (2016), adaptive responses are mediated by B and T lymphocytes, which select, activate, and clonally expand upon encountering an antigen. As a result, highly specific responses are produced that more precisely target the pathogen.

Interestingly, adaptive immunity generates memory cells, which enables rapid and more effective responses to re-exposure and hence long-lasting immunity that innate mechanisms themselves cannot generate. Innate and adaptive immunity operate in tandem within a system that balances speed and specificity. Innate immunity serves as the initial barrier, containing the infection and promoting the generation of adaptive responses, while adaptive immunity provides tailored and durable protection. As Chaplin (2010) also notes, the cooperation between these two branches is a critical part of the management of immune homeostasis and the effective protection against diverse pathogens. The original work of Medzhitov and Janeway (2000) is the one to follow, together with the overall framework presented by Murphy et al. (2016): these systems are interdependent in maintaining human health. This interplay between rapid, pattern-based innate screening and memory-driven adaptive specificity can be mapped onto medical imaging systems, where initial anomaly detection algorithms flag potential abnormalities that are then refined by deep learning models trained on large annotated datasets to support accurate and context-aware diagnosis.

4. Pattern Recognition Receptors (PRRs) as Inspiration for Feature Extraction in AI

Pattern recognition receptors (PRRs) are essential in the innate immune re-

response; they detect conserved molecular patterns called pathogen-associated molecular patterns (PAMPs). These receptors are biological detectors that recognize shared elements among pathogens, facilitating dynamic immune responses to potential risks. According to Akira, Uematsu, and Takeuchi (2006), PRRs trigger immune responses by identifying conserved patterns and initiating signaling pathways that evoke defense responses. Similar to how artificial intelligence (AI) systems use large datasets to identify meaningful features, this illustration shows how efficiently biological systems can extract meaningful information from complex environments.

The conceptual similarity between PRRs and feature extraction in AI is particularly evident in machine learning models such as convolutional neural networks (CNNs). In AI, feature extraction involves identifying important patterns such as edges, textures, and shapes from raw input data to enable classification and decision-making. In the same way, PRRs help to obtain molecular signatures that distinguish pathogens from host cells. Janeway (1989) highlighted the evolutionary significance of the immune system, recognizing conserved features of microbes and arguing that biological systems also depend on pattern-based recognition strategies, comparable to computational modeling. This analogy reinforces the conception of PRRs as natural feature extractors, with hierarchical features for detection, similar to the architecture of neural networks used in AI.

Moreover, the hierarchical structure of immune recognition parallels the multi-layered architecture found in deep learning systems. PRRs act as the front line of detectors that identify broad molecular patterns and initiate downstream responses, which may refine and amplify the immune response. Kawai and Akira (2010) describe how the various classes of PRRs (e.g., Toll-like and NOD-like receptors) can detect distinct molecular patterns, thereby eliciting a coordinated, multi-level immune response. The layered detection mechanism is comparable to how AI systems incrementally optimize each feature in stages, using several layers of features to improve classification performance. Both systems depend on blending basic pattern identification with more advanced interpretations.

The study of PRRs provides valuable insights for developing AI systems, particularly for feature extraction. By understanding how biological systems efficiently detect and respond to patterns, researchers can design more robust and adaptive machine learning models. The parallel between PRRs and AI feature extraction not only highlights the interdisciplinary connection between immunology and computer science but also underscores the potential for bio-inspired approaches to enhance artificial intelligence technologies. This comparison demonstrates that both natural and artificial systems rely on the fundamental principle of pattern recognition to interpret and respond to complex information. The analogy of pattern recognition receptors to hierarchical feature extraction can be applied to pathology image analysis, where AI systems identify low-level cellular patterns and progressively refine these signals through layered processing to detect disease-specific morphological abnormalities.

5. Immune Memory as a Model for Continual Learning in AI

Immune memory is an integral mechanism of the human body, often unified into technical systems due to its highly sophisticated and intelligent nature (Greensmith et al., 2010). Specifically, within artificial intelligence (AI), computational paradigms frequently parallel that of immunological systems, incorporating concepts such as clonal and negative selection to construct adaptive, continually learning networks (Greensmith et al., 2010; Al-Enezi et al., 2010). These immune-based frameworks enable long-term learning without catastrophic forgetting—a limitation that immune memory’s biological logic is uniquely positioned to resolve (Timmis et al., 2004; Souza et al., 2020; Rodríguez et al., 2020). Due to this positioning, examining how the immune system’s memory and recall processes parallel that of continual learning strategies in AI architectures clarifies how these dynamics enhance pathological detection within diagnostic models (Nunes de Castro & Von Zuben, 1999; Sithungu & Ehlers, 2023).

Central to this analysis is the biological memory index mediated by B- and T-cells within the immune system: threats are marked, stored, and made retrievable while discriminatory agents determine appropriate responses to both familiar and novel foreign bodies (Nunes de Castro & Von Zuben, 1999; Sithungu and Ehlers, 2023). These pathways work in tandem to produce a both stable and adaptive memory system eliminating continual learning difficulties such as the stability-plasticity dilemma where models fail to retrieve prior, grounding parameters in light of learning new ones (Harman, 2026; Bruno et al., 2025; Rodríguez et al., 2020). However, by modeling immune operations, AI is able to strike a balance within this computational paradox as evidenced by modern AIS detection tools (Greensmith et al., 2010; Park et al., 2021; Timmis, 2010). Yet the persistence of catastrophic forgetting across AI infrastructure—particularly within disease-detection applications—signals the remaining unresolved structural challenge (Harman, 2026; Bruno et al., 2025; Rodríguez et al., 2020). As models continue to acquire new information, prior parameters are repetitively forgotten, undermining the stability that diagnostic reliability demands (Harman, 2026; Bruno et al., 2025).

The tension derived by the competing demands of serving prior knowledge while simultaneously integrating new information produces asymmetry within the stability-plasticity framework, which immune-inspired algorithms naturally counteract (Harman, 2026; Bruno et al., 2025). Researchers have begun bridging this memory attrition gap via immune-inspired avenues—prioritizing continual, retained learning through both clonal and negative selection-based algorithms for proactive classification (Al-Enezi et al., 2010; Timmis, 2010; Souza et al., 2020; Sithungu and Ehlers, 2023). This shift from reactive to proactive is particularly significant regarding disease detection, where evolving diagnostic needs require that models maintain dynamic adaptability (Mukherjee et al., 2024; Salwa et al., 2025; Harman, 2026; Greensmith et al., 2010).

However, AIS programs are not without limitation, given they often require supervised learning or joint collaboration with neural networks, which possibly limits the potency of autonomous application (Al-Enezi et al., 2010; Park et al., 2021; Greensmith et al., 2010). Even so, collaboration between AIS programs and other frameworks mirrors a biological reality: innate and adaptive memory are not autonomous immune mechanisms (Greensmith et al., 2010). For that reason, dynamic interplay between structures is necessary for efficient anomaly detection within both organic processes and AI infrastructure (Greensmith et al., 2010; Park et al., 2021; Xie et al., 2026). This immune-inspired model of distributed pattern recognition and long-term memory can be mapped onto longitudinal EHR anomaly detection systems, where adaptive algorithms retain prior patient patterns while continuously integrating new data to detect subtle, evolving indicators of disease without losing previously learned clinical context.

6. Clonal Selection Theory and Its Influence on Evolutionary Algorithms

While immune memory describes how prior exposures are captured, clonal selection theory explains how the immune system refines and adapts its responses to pathogens over time (Nunes de Castro & Von Zuben, 1999; Al-Enezi et al., 2010). In general, B-cells produce antibodies which mark targeted antigens, while T-cells regulate the responses encoded within the thymus to either respond or suppress an immune reaction (Nunes de Castro & Von Zuben, 1999; O'Shea & Murray, 2008; Al-Enezi et al., 2010). More specifically, clonal selection refers to the process of these cells being chosen, copied, and optimized for prime threat recognition, forming improved affinity maturation over time (Nunes de Castro & Von Zuben, 1999; Souza et al., 2020). In the context of evolutionary algorithmic application, this biological process is often used as both a model and an outline for mapping high-performing solutions and under-explored areas simultaneously (Nunes de Castro & Von Zuben, 1999).

In practice, the optimization method labels all possible solutions as antibodies and treats the problem itself as an antigen (Al-Enezi et al., 2010; Greensmith et al., 2010). The algorithm then determines solutions of the highest affinity and clones them to the extent of their effectiveness, followed by the hyper-mutation process (Al-Enezi et al., 2010; Nunes de Castro & Von Zuben, 1999). Hyper-mutation describes the proportional inverse relationship between solution affinity and clone-rate, allowing the system to search underexplored areas and gradually refine top-performing solutions (Souza et al., 2020; Al-Enezi et al., 2010). This process fosters context recognition within the solution base, producing a diverse population of optima that retain anomalies while preserving established patterns (Nunes de Castro & Von Zuben, 1999). Consequently, discrimination between diverse inputs enhances detection reliability and supports refinement among learned representations.

Such discriminative capability is operationalized in models such as CLONALG

and AIRS (Artificial Immune Recognition System) (Al-Enezi et al., 2010; Timmis, 2010). In which CLONALG serves as a foundational clonal selection-inspired algorithm that leverages B-cell proliferation as an optimization technique for iterative solution calibration (Al-Enezi et al., 2010; Nunes de Castro & Von Zuben, 1999). Over successive iterations, CLONALG expands its receptiveness without disturbing established parameters, allowing continuous learning and smarter discernment (Greensmith et al., 2010; Nunes de Castro et al., 1999). Likewise, AIRS extends this framework to classification by retaining training data as a reference to compare against emerging threats (Timmis, 2010). While this reference system enables memory and affinity maturation, immune effectiveness ultimately relies on continuous surveillance to respond to such threats in real-time (Souza et al., 2020; Nunes De Castro & Von Zuben, 1999; Xie et al., 2026). This clonal selection-inspired process of iterative refinement and affinity optimization can be mapped onto pathology-based disease detection, where diagnostic models iteratively improve classification accuracy by amplifying high-confidence features while exploring ambiguous regions to enhance detection of subtle or emerging abnormalities.

7. Immune Surveillance as a Framework for Real-Time AI Health Monitoring

Immune surveillance functions as a primary defense mechanism, enabling the detection of cancerous cells (Swann & Smyth, 2007). It is a process of discrimination that persistently distinguishes between new information and potential threats, requiring constant vigilance rather than learning or optimization alone (Swann & Smyth, 2007). Similarly, health-monitoring algorithms must extend beyond memory and optimization to incorporate real-time detection for actionable insights (Salwa et al., 2025; Mukherjee et al., 2024). Conventionally, deep learning models rely on static, fully labeled datasets (Harman, 2026); however, data distributions in health monitoring are frequently unable to remain fixed due to the dynamics of clinical settings (Harman, 2026). These environments are often adaptive due to shifting demographics, protocols, and disease variants, which in turn cause static diagnostic model performance to deteriorate (Harman, 2026). Performance decay is particularly harmful for patients with time-sensitive diseases, where reactive medicine is not only ineffective but detrimental to treatment outcomes (Mukherjee et al., 2024).

Subsequently, it is not only the immune system, but cancerous or otherwise harmful cells that adapt as well (Tufail et al., 2025). Immunoediting depicts the phases tumor cells enact to avoid detection within a continually adapting system; the phases of evasion include elimination, equilibrium, and escape (Swann & Smyth, 2007). Beginning with elimination, the immune system targets and attempts to eradicate the tumor (Swann & Smyth, 2007). The tumor then enters a state of equilibrium, where it remains undetected while developing protective safeguards within its microenvironment, ultimately evading detection completely (Swann & Smyth, 2007; Tufail et al., 2025). This biological struggle parallels the challenges

found within diagnostic AI given that patient data distributions are continually evolving (Harman, 2026).

Detection is an ongoing requirement rather than a final state; accordingly, just as the biological immune system continuously surveils, health monitoring systems must ensure that detection continually adapts over time to support optimal prophylactic care (Swann & Smyth, 2007; Harman, 2026). One mechanism that supports this continuous detection is negative selection, where models such as NSMutation and Artificial Negative Selection Classifier (ANSC) utilize negative selection-inspired paradigms to establish a baseline of healthy data for robust anomaly filtration (Greensmith et al., 2010; Al-Enezi et al., 2010). This approach is increasingly useful for detecting unknown diseases or rare conditions because of its strong deviation identifiers that do not rely on predefined labels (Greensmith et al., 2010). Thus, by fusing computational paradigms with immune-inspired surveillance, adaptation, and anomaly filtering, diagnostic AI can achieve early detection of disease states before they become severe. However, in order to fully mitigate health repercussions, diagnostic systems require communication between agents a principal native to cytokine signaling. This dynamic of immune surveillance and adaptive evasion can be mapped onto continuous biosensor monitoring systems, where real-time data streams are persistently analyzed to detect emerging abnormalities while adapting to shifting physiological patterns and minimizing missed detections in evolving disease conditions.

8. Cytokine Signaling as a Model for Distributed AI Communication

Cytokine signaling, a sophisticated biological messaging system, is composed of small cell-secreted proteins that enact chemical communications between cells (Dinarello, 2007; Nunes De Castro & Von Zuben, 1999). Analogous to multi-agent communication protocols, cytokines function as signals within AI systems, while immune cells serve as agents which receive, send, and interpret information (Dinarello, 2007). The reliance on a decentralized, rather than centralized, communication method allows cytokines to facilitate sensitive, networked responses such as coordinated repair, defence, and pro-inflammatory processes (Dinarello, 2007). These signals bind to receptors, enabling autonomous systems to locally adapt to damage while learning resilience (Dinarello, 2007). Even if independent agents fail, the overall system is able to adapt, reallocate, and persevere avoiding bottlenecks (Greensmith et al., 2010).

For diagnostic systems, this model underscores the importance of distributed data interpretation rather than reliance on a single decision pathway. Effective diagnostics integrate diverse signals from a range of inputs to accurately quantify solutions, requiring both inter-agent communication and contextual awareness. While neural networks approximate this by mirroring structured reasoning processes, immune systems embed signal diversity, analysis, and retention through structural integration into the decision-making process rather than as a separate

component. Furthermore, by utilizing the structural cohesion inherent in AIS frameworks, diagnostic models gain adaptability based on context.

In practice, this may include key information such as environmental stress, health background, or daily life, particularly as emerging sensor technologies have revealed new biomarker data that provides in-vitro access to previously unobservable biological processes (Mukherjee et al, 2024; Salwa et al., 2025). This access to biological signals reframes diagnostics, allowing biomarkers to be interpreted not as individual data points, but as expressions of the broader environment from which they arise, enabling context to enter the diagnostic space and shape system-wide behavior.

This context-sensitive interpreting of in-vitro diagnostic signals is reflected in immune signaling architecture through a layered feedback structure. Cytokine feedback loops first operationalize context by producing responses that are contingent on receptor context rather than set parameters (O’Shea & Murray, 2008). At a second level, JAK-STAT-SOCS signaling forms an intracellular pathway in which information is transmitted, enabling decentralized signal propagation that is not uniformly distributed but directed toward concentrated regions where attention is most necessary for homeostasis (Murray, 2007). These network-like immune systems describe how signaling is generated and routed; however, interactions between components can regulate the system itself indirectly (Murray, 2007).

Idiotypic networks house antibody-to-antibody interactions that implicitly govern system-wide updating, where local interactions shape environmental shifts (Greensmith et al., 2010). Computationally, this hierarchy serves as operational inspiration for multi-agent AI systems in which context is encoded into the architecture through communicative agent adaptation, decentralized delegation, and indirect global state updates (Greensmith et al., 2010). Together, these mechanisms allow diagnostic systems to adapt to learned and shifting contexts, supporting scalable AI that inherently optimizes for informed detection (Harman, 2026). This cytokine-inspired model of decentralized, context-sensitive communication can be mapped onto biosensor monitoring systems, where distributed data streams—such as wearable and in vitro biomarker signals—are dynamically integrated and interpreted to enable real-time, context-aware disease detection.

9. Immune System Robustness as a Template for Fault-Tolerant AI

The immune system is a model system for AI (Artificial Intelligence) fault tolerance. Biological robustness is the ability of a system to perform functions in response to external and internal perturbations, including infection, mutations, or environmental change. Such robustness can be developed owing to redundancy, diversity, and modular organization (Kitano, 2004). These characteristics enable the immune system to remain effective when specific components fail, making them a key entry point into the study of how complex systems can resist stress.

Redundancy is a major concept in immune robustness; when many components perform similar tasks, breaking one does not necessarily weaken the whole. At the same time, the AI system needs redundancy to process noisy data, adversarial inputs, and hardware failures. According to [Csete and Doyle \(2002\)](#), biological systems are not meant solely for efficiency; they are designed for survivability, with pathways that overlap and include backup mechanisms. For an AI model where AI designers use redundant architectures, this insight is highly salient because redundancy in the underlying architecture can improve system stability and help mitigate failures or attacks.

Diversity, too, is critical because it enables the immune system to recognize a variety of pathogens and respond to them. This diversity allows the system to adapt, thus reinforcing its ability to address new challenges. The same goes for AI systems; they also benefit from varied models, training data, and algorithms to achieve better overall performance (generalization) and stability. [Alon \(2007\)](#) explains that the repetition of key recurring patterns in biological networks through motifs plays a crucial part in building stable (but more robust) systems. These motifs can inspire AI architectures that perform well regardless of conditions, underscoring the role of structural diversity in building resilient systems.

Therefore, the immune system's resilience to potential interruptions enables our design of intelligent or fault-tolerant AI systems. Combining concepts of redundancy, diversity, and modularity is the approach that makes AI more resilient to failures and unexpected situations. Biological robustness reveals useful practical approaches to strengthening the reliability and adaptive properties of artificial intelligence systems and deepens our appreciation of natural systems. This interdisciplinary link is important for encouraging the use of biological systems to generate more powerful computational models. This immune-inspired emphasis on robustness through redundancy, diversity, and modular organization can be mapped onto EHR anomaly detection systems, where ensemble models and multi-source data integration ensure stable and reliable disease detection despite noisy, incomplete, or evolving patient data.

10. Antigen Presentation as a Model for Explainable AI (XAI)

Antigen presentation provides a useful biological analogy to aid comprehension of explainable AI (XAI), as antigen-presenting cells (APCs) are innate “explainers” within the immune system. Dendritic cells, macrophages, and B cells, known as antigen-presenting cells (APCs), for example, work to process pathogenic proteins into peptide fragments for presentation on major histocompatibility complex (MHC) molecules. This process delivers a distinct and context-rich signal to T cells about potential threats. Antigen processing, as described by [Blum, Wearsch, and Cresswell \(2013\)](#), does not occur randomly but unfolds through a series of steps to inform T cells accurately, relevantly, and transparently. Indeed, APCs achieve XAI objectives by converting complex internal processes into interpretable outputs that serve as conduits for informing decisions.

Antigen Presentation

From an AI perspective, antigen presentation is an example of how explanations can enhance both the accuracy and credibility. Neefjes et al. (2011) define antigen presentation as a systems-level communication process where APCs integrate environmental signals, intracellular pathways, and the steps of molecular editing to elicit a significant “explanatory” signal. Likewise, XAI frameworks are intended to convert opaque model computations into human-readable rationale, maintaining contextual elements and highlighting salient features. Rock, Reits, and Neefjes (2016) point out that the immune system is successful when provided with reliable explanations; when it receives inaccurate or incomplete versions, immune failure may occur. This is similar to issues in AI systems where poor interpretability can lead to loss of reliability, user trust, and decision quality.

By investigating the interpretable pathways of natural APCs, XAI researchers can generate biologically intelligent models that provide comprehensible, useful explanations of outputs. Accordingly, antigen presentation offers a novel biological lens for understanding explainable AI (XAI) as a whole, especially for APCs, which are native “explainers” of the immune response. These cells, which comprise dendritic cells, macrophages, and B cells, convert pathogenic proteins into their constituent peptides, which are expressed on major histocompatibility complex (MHC) molecules. T cells will thus receive unambiguous signals, with pertinent context for hazards on their way.

Blum et al. (2013) note that antigen processing is not a random process but a multi-step, systematic one; the architecture provides for clear, accurate outlines that are clear to T cells. Thus, APCs demonstrate the basic goal of XAI: abstracting complex internal systems into constructs interpretable enough to facilitate well-informed decision-making. Understanding how AI applies to antigen presentation highlights how explanatory mechanisms bolster the accuracy and reliability of decision-making processes.

By exploring the interpretable pathways of natural APCs, XAI researchers have begun building biologically intelligent models that provide understandable, beneficial explanations for outputs. That’s why the presentation of antigen provides a new biological perspective on explainable AI (XAI), but especially for APCs, which act as native “explainers” in the immune response. These cells, which contain dendritic cells, macrophages, and B cells, process pathogenic proteins into their constituent peptides, which are expressed on major histocompatibility complex (MHC) molecules. As a result, T cells will receive clear signals, with pertinent context for hazards. Blum et al. (2013) also argue that antigen processing is not a random process, but a multi-step, systematic one; the architecture provides clear, accurate information outlines that are clear to T cells. So, APCs recognize the underlying goal of XAI, abstracting intricate internal systems into frameworks understandable enough to enable informed decision-making. It shows how AI can be applied and the benefits it brings when AI is applied to antigen presentation.

From the future, the investigation of APCs will furnish XAI researchers with

strong models grounded in biological assumptions, for example, the possibility of creating systems that can produce justifications consistent with the goals. Antigen Presentation as a Model for Explainable AI (XAI) shows how APCs “explain” threats to T-cells, paralleling interpretable AI pipelines. Antigen-presenting cells (APCs) process and present antigens to T-cells, thereby signaling threats. This biological transparency parallels explainable AI, where models must justify their decisions. Blum and Wearsch’s research on antigen presentation highlights the importance of clarity and context. XAI frameworks can draw inspiration from APCs to improve interpretability. The immune system demonstrates how explanation enhances decision-making accuracy. This antigen presentation–inspired emphasis on transparent, context-rich signaling can be mapped onto medical imaging systems, where explainable AI models highlight diagnostically relevant features within scans and provide interpretable rationales that support clinician trust and accurate disease detection.

11. Immune System Diversity as a Model for Ensemble Learning

In this context, the diversity of the immune system offers a great biological model for ensemble learning in artificial intelligence (AI). The many types of receptors generated by specialized genetic mechanisms in the adaptive immune system enable it to identify diverse forms of attack by a wide range of pathogens. For example, [Tonegawa \(1983\)](#) showed that antibody diversity arises from somatic recombination, enabling the immune system to generate millions of specific antigen receptors. This variety enables the immune system to sense and respond to numerous threats, underscoring how diversity enhances detection performance and flexibility in diverse environments. A T-cell receptor diversity strategy that [Davis and Bjorkman \(1988\)](#) argue is necessary for a properly functioning immune system to detect novel antigens.

This analogy reinforces the idea that diversification in biological and computational systems leads to more accurate and reliable detection. Immune diversity is then augmented by the various mechanisms involved in differentiation, making such comparisons clearer. [Alt et al. \(2013\)](#) describe how programmed DNA rearrangements, such as V(D)J recombination, generate a dynamic repertoire of immune receptors. The immune system must also be robust enough to change as these threats present themselves; it must adapt to new ones, which are likely to emerge. In the case of AI, ensemble learning accomplishes this as well.

Hence, we can gain insights into how the immune system leverages diversity as a core strategy for recognition and response in an ensemble AI. Both biological and artificial systems have found ways to achieve higher precision, stability, and adaptability by manipulating diversity and combining different perspectives. Similar to immune receptor diversity and ensemble learning, it suggests that a variety of models will be needed, or that one might combine multiple paths across systems to improve performance. As a result, this interdisciplinary connection illustrates

how principles from immunology can lead us toward more effective and resilient strategies in artificial intelligence. This immune-inspired use of receptor diversity and ensemble decision-making can be mapped onto medical imaging systems, where multiple complementary models integrate distinct feature representations to improve diagnostic accuracy and robustness in detecting complex or subtle disease patterns.

12. Autoimmunity as a Warning for AI Misclassification and Bias

Autoimmune diseases offer a compelling biological analogy for assessing risks to artificial intelligence systems. In autoimmunity, the immune system fails to maintain self-tolerance and incorrectly classifies the body's own cells as foreign and subsequently attacks healthy tissue. This process is akin to AI misclassification, in which systems misidentify benign or "normal" data as problematic or threatening, often leading to biased, dangerous outcomes (Goodnow, 2007).

Rose and Mackay (1998), as well as Davidson and Diamond (2001) and Goodnow (2007), found that immune recognition failures may have serious and systemic implications. These failures are not random; they frequently result from complex, multi-step breakdowns in regulatory mechanisms designed to ensure accuracy and balance. So, too, can AI systems become biased and generate false positives if those protections, diverse training data, fairness measures, and ever-present monitoring are inadequate or, at best, insufficiently designed.

Just as the immune system must separate self from non-self to remain safe, so do AI systems need to distinguish between correct and incorrect inputs without overcorrecting or introducing bias. Therefore, lessons from immunology can be useful for developing AI systems that prioritize fairness, minimize harm, and deploy multiple defenses against discrimination. Ultimately, autoimmunity serves as a warning when classification systems fail; the consequences can be harmful and pervasive throughout society. These lessons can also be applied to AI development to inform more responsible, trustworthy, equitable, and inclusive technology (Davidson & Diamond, 2001). This autoimmune misclassification dynamic can be mapped onto EHR-based disease detection systems, where poorly calibrated models may generate false positives by incorrectly flagging normal patient patterns as pathological, highlighting the need for robust safeguards to ensure accurate and equitable clinical decision-making.

13. Vaccination Principles as a Model for AI Pre-Training and Fine-Tuning

Vaccination provides a high-level of conceptual insight into how AI systems learn, adapt, and prevent harmful errors. Fundamentally, vaccination uses regulated exposure to antigens attenuated, inactivated, or molecularly modified copies of pathogens to hone the immune system without causing disease. Plotkin (2005) also points out that this type of controlled exposure is essential for establishing

broad, durable immunity. Similar to AI pretraining, the parallel is straightforward: large-scale supervised or self-supervised pretraining puts models through the wringer, providing them with enormous amounts of structured data so they can learn generalizable patterns from it and fine-tune them for specific downstream tasks. In the same way that vaccines train the immune system for real-world, pathogenic encounters, pre-training provides AI models with underlying architectures that enable elastic, context-sensitive performance across application contexts.

Vaccination

Vaccination provides a useful biological analogy to understand the pre-training and fine-tuning steps of AI models. In immunology, vaccines allow the immune system to encounter safe and controlled versions of a pathogen, leading to a variety of memory responses. As Plotkin (2005) states, this process of judicious exposure teaches the immune system to recognize threats efficiently without causing disease. This is akin to AI pre-training, in which a model is fed a massive, curated dataset that enables it to learn a broad set of concepts before being fine-tuned for its specific use case. Just as vaccines equip the immune system for the real world, pre-training equips AI models with knowledge applicable to target fields.

The same analogy is also made by the practices followed for vaccination. Puledran and Ahmed (2011) explain that vaccines activate both the innate and adaptive immune pathways, mechanisms that improve the quality, length, and specificity of immune memory. So in AI, fine-tuning works the same way: by using broad pre-training datasets and using specific ones, we optimize the model for increased accuracy and contextual relevance. Both rely on iterative exposure, feedback loops, and selective reinforcement. Thus, the ability of the immune system to modify its responses provides a biological basis for training an AI model that can be developed, maintaining a tradeoff between generalization and accuracy.

This comparison becomes more useful when comparing the consequences of misclassification. Autoimmune diseases occur when the immune system mistakenly attacks its own tissues as foreign invaders, leading to destructive attacks. Rose and Mackay's work has highlighted the serious impact that can occur when internal classification boundaries weaken within the system; the consequences of such mistakes can be catastrophic. Just as in the case of biased training data and wrong labeling of data in AI systems, the very data used in their training, bias, and labeling create opportunities for potentially damaging misclassifications in areas with implications, such as healthcare or criminal justice, that are life and death, and therefore the risks, especially in critical areas like medicine, healthcare, or crime, for example. Such parallels highlight the critical need for robust regulatory mechanisms, bias-minimization mechanisms, and continuous control measures to guard against "autoimmune-like" AI failures.

This includes recent developments in vaccine research, which are helping to develop more AI technology. Rappuoli (2018) highlights that systems biology, en-

gineered adjuvants, and precise targeting approaches are used in next-generation vaccines to amplify safety and effectiveness. Vaccination provides a high-level conceptual insight into how AI systems learn, adapt, and prevent harmful errors. Plotkin (2005) also notes that this type of controlled exposure is essential for establishing broad, durable immunity. Similar to AI pretraining, the parallel is straightforward: large-scale supervised or self-supervised pretraining puts models through the wringer, providing them with enormous amounts of structured data so they can learn generalizable patterns from it and fine-tune them for specific downstream tasks. In the same way that vaccines train the immune system for real-world, pathogenic encounters, pre-training provides AI models with underlying architectures that enable elastic, context-sensitive performance across application contexts. This vaccination-inspired process of controlled exposure, learning, and adaptive refinement can be mapped onto medical imaging systems, where AI models are pretrained on large, diverse datasets and fine-tuned on specific clinical images to enable accurate, early-stage disease detection while minimizing misclassification risk.

14. Immune Network Theory as a Blueprint for Self-Organizing AI Systems

Immune network theory can provide a useful model for both investigating and designing decentralized AI systems. Jerne (1974) proposed that the immune system functions as a dynamic, self-regulating network in which antibodies interact not only with external antigens but also with each other, forming a complex web of feedback and regulation. This decentralized environment helps the immune system keep up with dynamic changes and adaptations while avoiding a central control center. Likewise, self-organizing AI systems seek to distribute decision-making across interconnected parts to enable more flexible action in new situations.

At an earlier stage, such immunological principles were used in machine-learning experiments by Farmer, Packard, and Perelson (1986) to demonstrate that adaptive behavior can be generated by local interactions within a network. Perelson's (1989) further research suggests that such systems can reach equilibrium through constant internal regulation. Based on the framework of immune network theory, AI creators could design scalable, resilient, and autonomous adaptation systems, underscoring the importance of biological models in the development of distributed, self-organizing systems. This immune network-inspired model of decentralized, self-regulating interactions can be mapped onto biosensor monitoring systems, where distributed sensor inputs continuously communicate and locally adapt to detect emerging disease signals without reliance on a central processing unit.

15. Limitations

Although the immune system provides a compelling conceptual model for adap-

tive intelligence, its direct translation into artificial intelligence design and clinical deployment remains fundamentally limited. Biological immune responses arise from highly complex, nonlinear, and context-dependent interactions shaped by evolutionary pressures across millions of years (Murphy & Weaver, 2016). By contrast, artificial intelligence systems operate within computational architectures defined by discrete algorithms, constrained data inputs, and predefined optimization objectives (Russell & Norvig, 2021).

Because of these foundational differences, immunological constructs such as self versus non-self discrimination, clonal selection, or immune memory cannot be mapped onto computational systems without substantial abstraction. Such abstraction risks oversimplifying biological processes that depend on biochemical signaling networks, stochastic variability, and organism-level homeostasis. Clinical integration of AI introduces additional constraints that biological analogies do not inherently address. Healthcare AI systems must contend with heterogeneous data quality, regulatory oversight, and ethical considerations related to fairness, transparency, and patient safety (Topol, 2019).

These factors shape model development and deployment in ways that diverge significantly from the adaptive mechanisms observed in immunology. For example, while immune responses evolve through decentralized, self-organizing processes, clinical AI systems require explicit validation, monitoring, and governance frameworks to ensure reliability and mitigate harm (Amann et al., 2020). As a result, immune system intelligence should be treated as a conceptual guide rather than a literal blueprint for AI design. Maintaining this distinction ensures that bio-inspired approaches remain grounded in the practical realities of computational modeling and healthcare implementation while still benefiting from the conceptual richness of immunological theory.

16. Conclusion

The human immune system can be represented as a theoretical model for intelligence under adaptive, pattern-governed conditions and can thus serve as a useful way of viewing the applications of artificial intelligence to detect disease. Its ability for continuous tracking, pattern detection, self/non-self-discrimination, and dynamic adaptation to changing conditions offers a coherent framework that also corresponds to ongoing work in AI systems. To this conceptual synthesis, current research on AI in healthcare (e.g., Esteva et al., 2019; Jiang et al., 2017; Yu et al., 2018) is a symptom of evolving approaches towards a more proactive and anticipatory nature of diagnostic practice than discrete technological innovations. Immune system intelligence is therefore not meant to be adopted at a high level as a model, but as an overarching pattern of design that emphasizes flexibility and precision, and the capacity to build resilience as desirable qualities of the system.

The analogy between biological immune processes and computational systems suggests a shift from fixed process models of pattern recognition to a framework that is more context-sensitive and dynamic. Concurrently, other views on AI in

healthcare (i.e., those presented by Topol (2019)) consider technological development within the framework of human-centered care and argue that the intelligence of these systems must go beyond efficiency, but should also include ethical responsiveness and compatibility with the messiness of biosciences. Taken together, this interdisciplinary synthesis situates immune system intelligence as a conceptual model to facilitate understanding and evolution of AI systems, to inspire an integrated strategy to disease detection that is not only technically sound but also personalized, adaptive, and dynamic in terms of both healthcare environments.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- Akira, S., Uematsu, S., & Takeuchi, O. (2006). Pathogen Recognition and Innate Immunity. *Cell*, *124*, 783-801. <https://doi.org/10.1016/j.cell.2006.02.015>
- Al-Enezi, J. R., Abbod, M. F., & Alsharhan, S. (2010). Artificial Immune Systems-Models, Algorithms and Applications. *International Journal of Research and Reviews in Applied Sciences*, *3*, 118-131.
- Alon, U. (2007). Network Motifs: Theory and Experimental Approaches. *Nature Reviews Genetics*, *8*, 450-461. <https://doi.org/10.1038/nrg2102>
- Alt, F. W., Zhang, Y., Meng, F. L., Guo, C., & Schwer, B. (2013). Mechanisms of Programmed DNA Rearrangements. *Immunological Reviews*, *253*, 53-69.
- Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for Artificial Intelligence in Healthcare: A Multidisciplinary Perspective. *BMC Medical Informatics and Decision Making*, *20*, Article No. 310. <https://doi.org/10.1186/s12911-020-01332-6>
- Blum, J. S., Wearsch, P. A., & Cresswell, P. (2013). Pathways of Antigen Processing. *Annual Review of Immunology*, *31*, 443-473. <https://doi.org/10.1146/annurev-immunol-032712-095910>
- Bruno, P., Quarta, A., & Calimeri, F. (2025). Continual Learning in Medicine: A Systematic Literature Review. *Neural Processing Letters*, *57*, Article No. 2. <https://doi.org/10.1007/s11063-024-11709-7>
- Chaplin, D. D. (2010). Overview of the Immune Response. *Journal of Allergy and Clinical Immunology*, *125*, S3-S23. <https://doi.org/10.1016/j.jaci.2009.12.980>
- Csete, M. E., & Doyle, J. C. (2002). Reverse Engineering of Biological Complexity. *Science*, *295*, 1664-1669. <https://doi.org/10.1126/science.1069981>
- Davidson, A., & Diamond, B. (2001). Autoimmune Diseases. *New England Journal of Medicine*, *345*, 340-350. <https://doi.org/10.1056/nejm200108023450506>
- Davis, M. M., & Bjorkman, P. J. (1988). T-Cell Antigen Receptor Genes and T-Cell Recognition. *Nature*, *334*, 395-402. <https://doi.org/10.1038/334395a0>
- Dinarello, C. A. (2007). Historical Insights into Cytokines. *European Journal of Immunology*, *37*, S34-S45. <https://doi.org/10.1002/eji.200737772>
- Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K. et al. (2019). A Guide to Deep Learning in Healthcare. *Nature Medicine*, *25*, 24-29.

- <https://doi.org/10.1038/s41591-018-0316-z>
- Farmer, J. D., Packard, N. H., & Perelson, A. S. (1986). The Immune System, Adaptation, and Machine Learning. *Physica D: Nonlinear Phenomena*, *22*, 187-204. [https://doi.org/10.1016/0167-2789\(86\)90240-x](https://doi.org/10.1016/0167-2789(86)90240-x)
- Goodnow, C. C. (2007). Multistep Pathogenesis of Autoimmune Disease. *Cell*, *130*, 25-35. <https://doi.org/10.1016/j.cell.2007.06.033>
- Greensmith, J., Whitbrook, A., & Aickelin, U. (2010). *Artificial Immune Systems*. arXiv:1006.4949.
- Hamida, S. U., Chowdhury, M. J. M., Chakraborty, N. R., Biswas, K., & Sami, S. K. (2024). Exploring the Landscape of Explainable Artificial Intelligence (XAI): A Systematic Review of Techniques and Applications. *Big Data and Cognitive Computing*, *8*, Article 149. <https://doi.org/10.3390/bdcc8110149>
- Harman, F. (2026). Incremental Learning in Medical Imaging: A Comprehensive Survey of Deep Neural Network Advances and Challenges. *IEEE Access*, *14*, 15812-15840. <https://doi.org/10.1109/access.2026.3659025>
- Janeway, C. A. (1989). Approaching the Asymptote? Evolution and Revolution in Immunology. *Cold Spring Harbor Symposia on Quantitative Biology*, *54*, 1-13. <https://doi.org/10.1101/sqb.1989.054.01.003>
- Kawai, T., & Akira, S. (2010). The Role of Pattern-Recognition Receptors in Innate Immunity: Update on Toll-Like Receptors. *Nature Immunology*, *11*, 373-384. <https://doi.org/10.1038/ni.1863>
- Kitano, H. (2004). Biological robustness. *Nature Reviews Genetics*, *5*, 826-837. <https://doi.org/10.1038/nrg1471>
- Jerne, N. K. (1974). Towards a Network Theory of the Immune System. *Annales D'Immunologie*, *125C*, 373-389.
- Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S. et al. (2017). Artificial Intelligence in Healthcare: Past, Present and Future. *Stroke and Vascular Neurology*, *2*, 230-243. <https://doi.org/10.1136/svn-2017-000101>
- Medzhitov, R., & Janeway, C. (2000). Innate Immunity. *New England Journal of Medicine*, *343*, 338-344. <https://doi.org/10.1056/nejm200008033430506>
- Mukherjee, D., Paul, K., Banerjee, R. et al. (2024). From Reactive to Proactive: The Rise of Sensor-Based Chronic Disease Management. *Journal of Nonlinear Analysis and Optimization*, *15*, 70-77.
- Murphy, K., Weaver, C., & Berg, L. (2016). *Janeway's Immunobiology* (9th ed.). Garland Science.
- Murray, P. J. (2007). The JAK-STAT Signaling Pathway: Input and Output Integration. *The Journal of Immunology*, *178*, 2623-2629. <https://doi.org/10.4049/jimmunol.178.5.2623>
- Neefjes, J., Jongsma, M. L. M., Paul, P., & Bakke, O. (2011). Towards a Systems Understanding of MHC Class I and MHC Class II Antigen Presentation. *Nature Reviews Immunology*, *11*, 823-836. <https://doi.org/10.1038/nri3084>
- Nunes de Castro, L., & Von Zuben, F. (1999). *Artificial Immune Systems: Part I—Basic Theory and Applications*. https://www.researchgate.net/publication/228712042_Artificial_Immune_Systems_Part_I-Basic_Theory_and_Applications
- O'Shea, J. J., & Murray, P. J. (2008). Cytokine Signaling Modules in Inflammatory Responses. *Immunity*, *28*, 477-487. <https://doi.org/10.1016/j.immuni.2008.03.002>

- Park, H., Choi, J. E., Kim, D., & Hong, S. J. (2021). Artificial Immune System for Fault Detection and Classification of Semiconductor Equipment. *Electronics, 10*, Article 944. <https://doi.org/10.3390/electronics10080944>
- Perelson, A. S. (1989). Immune Network Theory. *Immunological Reviews, 110*, 5-36. <https://doi.org/10.1111/j.1600-065x.1989.tb00025.x>
- Plotkin, S. A. (2005). Vaccines: Past, Present and Future. *Nature Medicine, 11*, S5-S11. <https://doi.org/10.1038/nm1209>
- Pulendran, B., & Ahmed, R. (2011). Immunological Mechanisms of Vaccination. *Nature Immunology, 12*, 509-517. <https://doi.org/10.1038/ni.2039>
- Rappuoli, R. (2018). Vaccines: New Opportunities. *Proceedings of the National Academy of Sciences, 115*, 2329-2336.
- Rock, K. L., Reits, E., & Neefjes, J. (2016). Present Yourself! By MHC Class I and MHC Class II Molecules. *Trends in Immunology, 37*, 724-737. <https://doi.org/10.1016/j.it.2016.08.010>
- Rodríguez, C. et al. (2020). *Future—Adaptive Intelligence and Lifelong Systems*. Fupress.org. <https://fupress.org/journal/AILS/index.php/journal>
- Rose, N. R., & Mackay, I. R. (1998). The Autoimmune Diseases. *New England Journal of Medicine, 338*, 1832-1841.
- Russell, S., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.
- Salwa, H. et al. (2025). *Continual Learning: Overcoming Catastrophic Forgetting for Adaptive AI Systems*. Institute of Electrical and Electronics Engineers (IEEE).
- Sithungu, S. P., & Ehlers, E. M. (2023). SupervisedimmuneNet: Training Artificial Immune Networks Using a Supervised Learning Approach for Improved Multi-Class Classification. In *Proceedings of the 2023 6th International Conference on Computational Intelligence and Intelligent Systems* (pp. 118-123). ACM. <https://doi.org/10.1145/3638209.3638227>
- Souza, S. F., Lima, F. P. D. A., & Chavarette, F. R. (2020). A New Artificial Immune System Based on Continuous Learning for Pattern Recognition. *Revista de Informática Teórica e Aplicada, 27*, 34-44. <https://doi.org/10.22456/2175-2745.102061>
- Swann, J. B., & Smyth, M. J. (2007). Immune Surveillance of Tumors. *Journal of Clinical Investigation, 117*, 1137-1146. <https://doi.org/10.1172/jci31405>
- Timmis, J. (2010). Artificial Immune Systems. In C. Sammut, & G. Webb (Eds.), *Encyclopedia of Machine Learning* (pp. 40-44). Springer. https://doi.org/10.1007/978-0-387-30164-8_33
- Timmis, J., Knight, T., de Castro, L. N., & Hart, E. (2004). An Overview of Artificial Immune Systems. In: R. Paton, H. Bolouri, M. Holcombe, J. H. Parish, & R. Tateson (Eds.), *Computation in Cells and Tissues* (pp. 51-59). Springer.
- Tonegawa, S. (1983). Somatic Generation of Antibody Diversity. *Nature, 302*, 575-581. <https://doi.org/10.1038/302575a0>
- Topol, E. (2019). *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books.
- Tufail, M., Jiang, C., & Li, N. (2025). Immune Evasion in Cancer: Mechanisms and Cutting-Edge Therapeutic Approaches. *Signal Transduction and Targeted Therapy, 10*, Article No. 227. <https://doi.org/10.1038/s41392-025-02280-1>
- Wang, L., Zhang, X., Su, H., & Zhu, J. (2015). A Comprehensive Survey of Continual Learning: Theory, Method and Application. *Journal of Latex Class Files, 14*, Article 1.

- Xie, L., Pelissier, A., Shao, Y. et al. (2026). *AI Developments for T and B Cell Receptor Modeling and Therapeutic Design*. arXiv:2601.17138.
- Yu, K. H., Beam, A. L., & Kohane, I. S. (2018). Artificial Intelligence in Healthcare. *Nature Biomedical Engineering*, 2, 719-731. <https://doi.org/10.1038/s41551-018-0305-z>