# Radar-Based Sensing vs. Camera-Based Vision in Healthcare

## Why Physics Determines Privacy, Anonymity, and Data Protection

**QUMEA AG**

December 2025 | Version 1.0

## Abstract

This white paper examines the differences between radar-based sensing systems and camera-based computer vision systems in healthcare monitoring contexts, with a particular focus on data protection, anonymity, and the distinction between anonymous and anonymized data. It analyzes the technical nature of raw sensor data, the role of artificial intelligence in interpretation, and the regulatory and ethical implications under modern data protection frameworks. It further explains why radar and cameras cannot be treated as equivalent modalities: the wavelength regime determines what identity-bearing visual information can exist in the captured signal. The paper adopts a neutral, technology-agnostic perspective and draws on publicly available scientific, regulatory, and industry sources.

## 1. Introduction

Digital sensing technologies are increasingly used in healthcare environments to support patient safety, including fall prevention and mobility monitoring. While such systems promise efficiency and improved outcomes, they also raise significant questions regarding privacy, data protection, and cybersecurity. Two dominant technological approaches are radar-based motion sensing and camera-based computer vision. Although both can rely on artificial intelligence for interpretation, they differ fundamentally in the type of data they generate and the associated privacy implications.

## 2. Conceptual Distinction Between Personal, Anonymized, and Anonymous Data

Under the European Union General Data Protection Regulation (GDPR), personal data is defined as any information relating to an identified or identifiable natural person (Regulation (EU) 2016/679, Article 4(1)). Visual data, including images and video streams, are generally considered personal data because individuals may be identified directly or indirectly from visual features. Any operation performed on such data—such as collection, storage, analysis, transmission, or deletion—constitutes *processing* under the GDPR, regardless of whether it is automated (Regulation (EU) 2016/679, Article 4(2), https://gdpr.eu/article-4-definitions/).

Anonymized data refers to personal data that has undergone processing intended to irreversibly prevent identification. However, scientific literature indicates that anonymization techniques applied to visual data, such as pixelation or blurring, may not always eliminate re-identification risks, particularly when combined with auxiliary data or advanced machine learning techniques (https://arxiv.org/abs/2209.12046).

Anonymous data, in contrast, refers to data that does not relate to an identifiable individual from the point of collection onward. Such data does not require subsequent anonymization because no personal identifiers are ever captured.

## 3. Physical Limits of Electromagnetic Sensing and Spatial Resolution

All electromagnetic sensing systems are subject to fundamental physical limits that determine what information can be captured at the point of measurement. While different sensing modalities may use different system architectures, the decisive factor governing achievable spatial resolution and information content is the wavelength of the electromagnetic radiation employed.

Optical cameras operate at wavelengths on the order of hundreds of nanometers (e.g., ~0.4–0.7 µm for visible light, corresponding to 450-750 Terahertz), whereas radar systems used for indoor and healthcare applications typically operate at wavelengths on the order of several millimeters (e.g., ~12.5–3.9 mm, corresponding to 24–77 GHz). This difference of four to five orders of magnitude (i.e. factor 100’000) has direct consequences for achievable spatial resolution.

According to classical wave optics, the minimum resolvable spatial detail of any electromagnetic sensing system is limited by diffraction. A commonly used approximation is given by the Rayleigh criterion:

$$\Delta x \approx \frac{\lambda}{2 \cdot \mathrm{NA}}$$

where λ denotes the wavelength and NA the numerical aperture of the sensing system. In practice, NA is bounded by sensing geometry (aperture size and distance), which means that improved spatial resolution requires proportionally larger apertures and/or shorter wavelengths.

As a consequence, matching the spatial detail available from optical imaging at millimeter-wave wavelengths would require apertures that are orders of magnitude larger than practical for indoor healthcare deployments. This constraint also applies to synthetic-aperture approaches and does not depend on implementation choices.

This limitation applies to the information present in the measured electromagnetic field itself. Machine learning and artificial intelligence can only exploit information that exists in the input signal; they cannot reconstruct spatial detail that is absent due to wavelength-limited resolution. Consequently, claims that software or AI could “recover” camera-like visual detail from radar data contradict fundamental principles of wave physics.

As a result, optical cameras are capable of resolving fine spatial features such as facial characteristics, text, and object details. Radar systems operating below 100 GHz, by contrast, are physically incapable of resolving such features, even under ideal conditions and regardless of computational complexity.

Importantly, this limitation is not a design choice but a consequence of wave physics. No post-processing technique can reconstruct spatial information that was never present in the measured signal.

These resolution limits arise from classical wave optics and are independent of quantum-mechanical effects. Quantum noise considerations do not alter the fundamental wavelength-dependent bound on spatial resolution in radar or optical sensing systems.

These physical constraints determine, at the point of data acquisition, whether a sensing modality can in principle capture identity-bearing visual information. The implications of these physical limits for radar-based sensing and camera-based computer vision systems are discussed in the following sections.

### 3.1. Visual Comparison of Sensing Modalities and Identity-Bearing Information

Visually compared, the key differences between common sensing modalities used in healthcare monitoring become immediately apparent. The underlying physical principles – particularly wavelength and signal representation – determine the type of information that can exist in the captured data at the point of acquisition. This distinction is central to understanding why radar-based sensing and camera-based computer vision differ fundamentally with respect to spatial resolution, identity-bearing visual information, and resulting data protection implications:

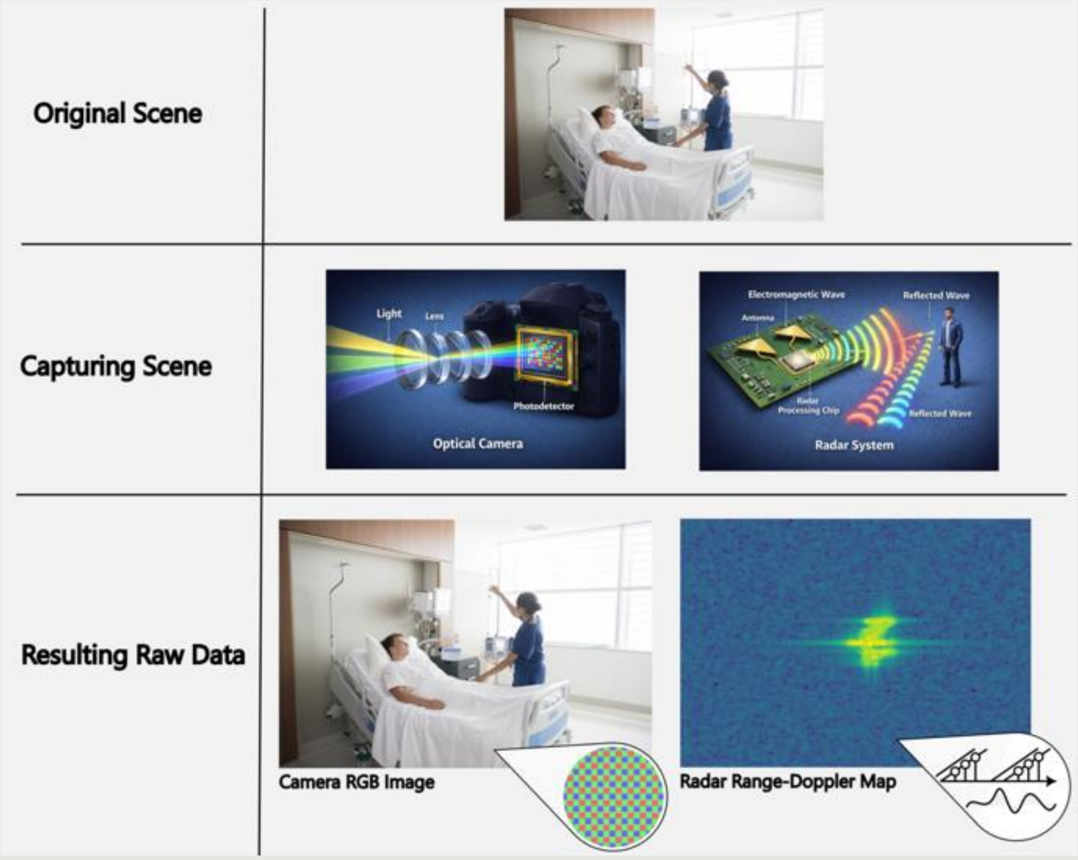

*Source: QUMEA AG, 2025*

## 4. Technical Characteristics of Radar-Based Sensing Data

Radar-based sensing systems emit electromagnetic waves and analyze reflected signals to derive information about motion, position, and velocity. The resulting raw data typically consists of multidimensional time-series measurements such as range, Doppler shift, and angle of arrival. These data representations are non-visual and not directly interpretable by humans without specialized signal processing and machine learning models.

Radar data does not and cannot contain facial features, textures, or other camera-like visual characteristics. As such, it lacks inherent visual identifiers that could reasonably be used to identify individuals. The interpretation of radar data relies on abstract patterns of movement rather than visual appearance (https://qumea.com/en/technology/).

Consistent with the wavelength-limited constraints described in Section 3, radar measurements encode coarse spatial and temporal information about reflecting objects, limited to scales on the order of the radar wavelength. Consequently, radar data lacks the information necessary to reconstruct faces, text, or other identity-bearing visual features, even in principle.

## 5. Technical Characteristics of Camera-Based Computer Vision Data

Camera-based systems rely on optical sensors that capture images or video frames. By definition, any optical camera produces a spatially resolved image at the sensor plane; whether this image is stored, transmitted, or immediately processed does not change the fact that it is captured. These frames are composed of pixel values that directly represent the visual appearance of people and environments. Even when such systems apply anonymization techniques, the original image data is captured before transformation, creating potential exposure to personal data.

CMOS and CCD sensors convert incident photons into pixel-level electrical signals prior to any software-based processing. As a result, visual representations of individuals and environments exist at the hardware level before anonymization, filtering, or inference logic is applied.

Unless physically constrained at the silicon or optical level, the raw image data produced by an optical sensor can in principle be accessed, reconstructed, or repurposed. Anonymization therefore does not prevent the capture of personal data but constitutes a subsequent transformation applied after personal data has already been processed.

Training computer vision systems capable of robust and clinically relevant performance typically relies on large-scale visual training data, often through pre-training on very large image corpora followed by task-specific fine-tuning.

Large-scale empirical studies in computer vision demonstrate that model accuracy and generalization improve systematically as the amount of training data increases. For example, Sun et al. analyze representation learning using datasets ranging from millions to hundreds of millions of images and report consistent performance gains as dataset size grows. Kolesnikov et al. ("Big Transfer" , BiT) similarly show that pre-training on increasingly large image datasets yields improved transfer performance across a wide range of downstream vision tasks.

This reliance implies that the development of camera-based systems involves the collection and processing of substantial volumes of visual data during model development and validation, independent of whether anonymization is applied during deployment. The capturing, storage, transmission, and processing of these images increase the attack surface for cybersecurity incidents, particularly when edge computing devices are used.

## 6. Artificial Intelligence and Interpretation Layers

Artificial intelligence systems can only operate on information present in the input data; they cannot generate spatial or visual detail that was not captured at the sensing stage. Both radar-based and camera-based systems rely on artificial intelligence to interpret sensor data. However, the nature of the input data significantly influences privacy risk. In radar-based systems, machine learning models are trained on abstract motion features and signal patterns. In camera-based systems, models learn from pixel-level visual information, which may encode sensitive attributes even when anonymization is applied.

Scientific studies have shown that machine learning models can inadvertently retain or infer personal characteristics from visual data representations, raising concerns about latent privacy leakage (https://arxiv.org/abs/2209.12046).

### 6.1 Training Data Requirements and Latent Privacy Exposure

Independent of deployment architecture, any artificial intelligence system designed to interpret camera images must be trained on image data. This training phase necessarily involves the processing of visual representations that may contain identifiable individuals, environments, and contextual information.

Even when inference is performed on-device and anonymization is applied in real time, underlying models are typically trained on datasets that include non-anonymized images. This creates an unavoidable exposure to personal data during model development and validation.

Radar-based systems do not face this constraint in the same way, as their training data consists of abstract signal representations that do not encode camera-like visual appearance or identity-bearing facial features.

## 7. Security and Infrastructure Considerations

Recent developments in Denmark illustrate the concrete consequences of camera-based surveillance risks and explain how the issue entered the parliamentary agenda. A parliamentary question raised in September 2025 regarding surveillance technology and security triggered a review of deployed camera systems, which in December 2025 led to the disclosure that surveillance cameras from Chinese manufacturers Hikvision and Dahua were in use in a regional hospital despite prior political assurances to the contrary. The case underscores that camera systems can introduce regulatory, cybersecurity, and geopolitical risks, including loss of oversight, exposure to foreign-manufactured hardware outside approved security infrastructure, and subsequent political and public accountability.

The cybersecurity risk associated with a sensing system is closely linked to the sensitivity of the raw data accessible at the hardware and infrastructure level. Camera-based systems often depend on network-connected edge devices for real-time processing. Public vulnerability disclosures and vendor security advisories for embedded AI hardware platforms demonstrate that such devices can be targets for remote exploitation and data exfiltration (e.g., NVIDIA security advisories: https://nvidia.custhelp.com/). When visual data is involved, successful attacks may result in exposure of highly sensitive information.

In addition to edge computing infrastructure, the camera hardware itself constitutes a potential attack surface. Network-connected cameras have repeatedly been shown to contain exploitable vulnerabilities, enabling unauthorized access to live video streams or stored image data. In such cases, compromise of the camera device may allow direct access to unprocessed visual data prior to any anonymization or on-device filtering.

Radar-based systems, while not immune to cybersecurity risks, generally process data that is insensitive in nature due to the absence of camera-like visual identifiers.

## 8. Ethical and Regulatory Implications

From a regulatory perspective, systems that do not process personal data fall outside many data protection obligations, reducing compliance complexity. Ethical analyses in healthcare technology emphasize privacy-by-design principles, advocating for systems that minimize data sensitivity at the point of collection rather than relying on post-processing safeguards.

In addition, both the General Data Protection Regulation (GDPR) and national data protection laws such as the Swiss Federal Act on Data Protection (DSG) explicitly incorporate the principle of proportionality. This principle requires that, for any given purpose, data processing must be limited to what is necessary and appropriate, and that less intrusive alternatives must be preferred where they can reasonably achieve the same objective. Consequently, when addressing safety-related problems such as fall prevention, solutions that achieve the intended outcome while collecting the least amount of sensitive data are considered ethically and legally preferable. Systems that rely on inherently data-minimizing sensing modalities align more closely with proportionality requirements than approaches that collect highly detailed personal data and subsequently attempt to reduce sensitivity through anonymization.

Within a clinical context, proportionality also has implications for clinical appropriateness. Ethical and clinical literature on sensor-based surveillance indicates that highly intrusive monitoring technologies may impose psychological, dignity-related, or trust-related burdens on certain patient populations, particularly individuals with cognitive impairment or reduced capacity to consent. In such cases, the use of more intrusive sensing modalities may be considered clinically inappropriate, or contextually contraindicated, if comparable safety benefits can be achieved through less intrusive means. These considerations suggest that the intrusiveness of a sensing modality constitutes a clinically relevant parameter and that proportionality functions not only as a legal requirement, but also as a constraint on clinically acceptable technology selection.

The trade-off between data utility and privacy is particularly pronounced in camera-based systems, where stronger anonymization often reduces analytical performance. Radar-based systems avoid this trade-off by operating on non-identifying signals.

## 9. Conclusion

Radar-based sensing and camera-based computer vision represent fundamentally different approaches to healthcare monitoring because they operate in different physical regimes that determine what information can be captured at source.

While both radar and camera systems use electromagnetic radiation, equating the two based on this fact alone is scientifically incorrect. Optical cameras operate in a wavelength regime that inherently supports the capture of fine spatial detail and visual appearance, including identity-bearing features. Radar systems operating below 100 GHz do not. This distinction is enforced by physical law and cannot be altered by system design, software configuration, or artificial intelligence. Consequently, radar-based systems can be anonymous at source with respect to identity-bearing visual information, whereas camera-based systems can only be anonymized after personal data has already been captured.

Radar systems generate non-visual, abstract data that does not contain identity-bearing visual information at source and is therefore fundamentally different from image-based personal data, whereas camera systems inherently capture visual representations that are closely linked to individual identity. While both technologies can support safety-related applications, their privacy, security, and regulatory profiles differ substantially.

Future research and policy discussions should continue to evaluate sensing technologies not only on performance metrics but also on the nature of the data they generate and the associated societal implications.

**References**

- European Parliament and Council of the European Union. General Data Protection Regulation (EU) 2016/679. https://gdpr.eu/
- Carmichael, Z., et al. "On the Privacy Risks of Passive Human Sensing." arXiv, 2022. https://arxiv.org/abs/2209.12046
- Swiss Medtech. "AI in Healthtech – Radar-Based Patient Monitoring." 2025. https://swiss-medtech.ch/
- QUMEA. Technology overview. https://qumea.com/en/technology/
- Sun, C., Shrivastava, A., Singh, S., & Gupta, A. "Revisiting the Unreasonable Effectiveness of Data in Deep Learning Era." ICCV 2017. https://arxiv.org/abs/1707.02968
- Kolesnikov, A., Beyer, L., Zhai, X., Puigcerver, J., Yung, J., Gelly, S., & Houlsby, N. "Big Transfer (BiT): General Visual Representation Learning." ECCV 2020. https://arxiv.org/abs/1912.11370
- Hestness, J., et al. "Deep Learning Scaling is Predictable, Empirically." arXiv:1712.00409. https://arxiv.org/abs/1712.00409
- NVIDIA Corporation. "Jetson Embedded Computing Platforms – Security Advisories." https://nvidia.custhelp.com/
- Lahr, J., Schulze, N., Wüst, L., Beisbart, C., Bruhin, L. C., Ienca, M., Nef, T., Trachsel, M., & Klöppel, S. (2025). Ethics of Sensor-Based Surveillance of People with Dementia in Clinical Practice. Sensors, 25(7), 2252. https://doi.org/10.3390/s25072252
- Limits of Resolution: The Rayleigh Criterion: https://openbooks.lib.msu.edu/collegephysics2/chapter/limits-of-resolution-the-rayleigh-criterion-2/
- Goodman, J. W. Introduction to Fourier Optics. Roberts & Company, ISBN 978-0974707723
- Skolnik, M. I. Radar Handbook. McGraw-Hill, ISBN: 978-0071485470
- Richards, M. A. Fundamentals of Radar Signal Processing. McGraw-Hill, ISBN: 978-1260468717
- Holst, G. C. CMOS/CCD Sensors and Camera Systems. SPIE Press, ISBN: 9781510627116