



A systematic review of intermediate fusion in multimodal deep learning for biomedical applications

Valerio Guarrasi ^{a,1,*}, Fatih Aksu ^{b,a,1}, Camillo Maria Caruso ^{a,1}, Francesco Di Feola ^{c,1},
Aurora Rofena ^{a,1}, Filippo Ruffini ^{a,1}, Paolo Soda ^{a,c,d,1,**}

^a Research Unit of Computer Systems and Bioinformatics, Department of Engineering, Università Campus Bio-Medico di Roma, Rome, Italy

^b Department of Biomedical Sciences, Humanitas University, Milan, Italy

^c Department of Diagnostics and Intervention, Radiation Physics, Biomedical Engineering, Umeå University, Umeå, Sweden

ARTICLE INFO

Keywords:

Joint fusion
Fusion techniques
Data integration
Biomedical data
Healthcare
Data fusion

ABSTRACT

Deep learning has revolutionized biomedical research by providing sophisticated methods to handle complex, high-dimensional data. Multimodal deep learning (MDL) further enhances this capability by integrating diverse data types such as imaging, textual data, and genetic information, leading to more robust and accurate predictive models. In MDL, differently from early and late fusion methods, intermediate fusion stands out for its ability to effectively combine modality-specific features during the learning process. This systematic review comprehensively analyzes and formalizes current intermediate fusion methods in biomedical applications, highlighting their effectiveness in improving predictive performance and capturing complex inter-modal relationships. We investigate the techniques employed, the challenges faced, and potential future directions for advancing intermediate fusion methods. Additionally, we introduce a novel structured notation that standardizes intermediate fusion architectures, enhancing understanding and facilitating implementation across various domains. Our findings provide actionable insights and practical guidelines intended to support researchers, healthcare professionals, and the broader deep learning community in developing more sophisticated and insightful multimodal models. Through this review, we aim to provide a foundational framework for future research and practical applications in the dynamic field of MDL.

1. Introduction

Deep learning has transformed the landscape of computational approaches in biomedical research, offering unique capabilities in handling complex, high-dimensional data [1]. Multimodal deep learning (MDL), which integrates multiple data types, has emerged as an innovative approach, leveraging the power of deep learning algorithms to interpret and integrate diverse data types, enhancing the robustness and accuracy of predictive models [2]. Unlike traditional unimodal methods, MDL utilizes a variety of data sources such as imaging, textual data, and genetic information, to name a few, providing a more comprehensive understanding of complex biomedical phenomena. This approach is particularly relevant in a field where data diversity and volume are rapidly expanding, offering unprecedented opportunities for advancements in diagnosis, treatment, and patient care.

In the domain of multimodal learning, the integration of data from multiple sources can be achieved through various fusion techniques, *i.e.*, early, late and intermediate fusion. Early fusion combines features at the data level, potentially losing unique modality-specific characteristics. Formally, let x_1, x_2, \dots, x_n represent feature vectors from n different modalities. The fusion function \mathcal{F} combines these features into a single vector:

$$x = \mathcal{F}(x_1, x_2, \dots, x_n) \quad (1)$$

which is then fed into a learning model f , resulting in the output y :

$$y = f(x) \quad (2)$$

Late fusion, on the other hand, occurs at the decision level, often missing opportunities for deeper interaction between modalities. Formally, the feature vectors x_1, x_2, \dots, x_n are fed into independent models

* Corresponding author.

** Corresponding author at: Department of Diagnostics and Intervention, Radiation Physics, Biomedical Engineering, Umeå University, Umeå, Sweden.
E-mail addresses: valerio.guarrasi@unicampus.it (V. Guarrasi), fatih.aksu@st.hunimed.eu (F. Aksu), camillomaria.caruso@unicampus.it (C.M. Caruso), francesco.feola@umu.se (F. Di Feola), aurora.rofena@unicampus.it (A. Rofena), filippo.ruffini@unicampus.it (F. Ruffini), p.soda@unicampus.it, paolo.soda@umu.se (P. Soda).

¹ These authors contributed equally to this work.

f_1, f_2, \dots, f_n , producing separate predictions y_1, y_2, \dots, y_n . The fusion function \mathcal{F} then combines these predictions:

$$y = \mathcal{F}(y_1, y_2, \dots, y_n) \quad (3)$$

Intermediate fusion, striking a balance, integrates data at the feature extraction stage, allowing for a more effective combination of modality-specific features. Formally, the feature vectors x_1, x_2, \dots, x_n are processed by separate components of the model f_1, f_2, \dots, f_n , yielding intermediate representations h_1, h_2, \dots, h_n . These intermediate representations are fused using \mathcal{F}

$$h = \mathcal{F}(h_1, h_2, \dots, h_n) \quad (4)$$

and then a final component f processes h to produce the output:

$$y = f(h) \quad (5)$$

Further details on each component can be found in Section 3.1. Among the various fusion techniques employed in MDL, intermediate fusion stands out for its ability to effectively integrate information at essential stages of the learning process, potentially leading to more accurate and robust models, as deepened in Section 3.7.5. This method is especially beneficial in biomedical applications where different data types, like imaging and genomic information, need to interact closely to produce meaningful insights. The main feature of intermediate fusion lies in its ability to preserve and utilize the distinct qualities of each data type, enhancing the model's capability to handle the complexities inherent in biomedical data.

The utility of deep learning in the context of intermediate fusion in multimodal settings, particularly within biomedical research, is marked by its ability to process and fuse distinct data modalities at a more abstract level [3]. In intermediate fusion, deep learning models are particularly capable at entangling the complex, nonlinear relationships that often exist between different modalities in biomedical data. This understanding is critical for accurately interpreting the multimodal nature of such data, where each modality carries its own information. By leveraging deep learning's ability in feature extraction and representation learning, intermediate fusion methods can effectively bridge the gap between these diverse data sources, leading the way for more insightful biomedical analysis. This approach not only may enhance the accuracy of predictive models but also may unveil patterns and interactions that might be missed by more conventional methods.

The application of MDL in the biomedical domain faces unique challenges, notably due to the heterogeneity and high dimensionality of the data [4]. Biomedical datasets often comprise varied data modalities, each with distinct characteristics and scales, making it difficult to achieve effective integration and analysis. Intermediate fusion, applied within the deep learning framework, addresses these challenges efficiently. By fusing data at the feature level, it allows to effectively manage the heterogeneity by learning a shared representation space, where the complex interaction between unimodal and multimodal features is explored. Moreover, intermediate fusion helps in mitigating the curse of dimensionality, a common issue in biomedical datasets [5], by facilitating dimensionality reduction and feature selection in a way that preserves the most informative aspects of each modality.

While deep learning and intermediate fusion offer significant advantages in handling multimodal biomedical data, they are not without limitations [6]. One of the primary concerns is the black-box nature of deep learning models, which often leads to a lack of transparency and interpretability [7]. This is particularly problematic in biomedical settings where understanding the reasoning behind a model's decision can be as critical as the decision itself. Additionally, these models require large amounts of data to perform optimally, which can be a challenge in situations where data is scarce or privacy concerns limit data availability [8]. Intermediate fusion, despite its benefits, also faces challenges in optimally balancing the contributions of different modalities, especially when the data from one modality is sparse or

of lower quality [2]. Furthermore, the computational complexity of deep learning models, tangled with the complexity of intermediate fusion, can lead to significant resource demands, posing challenges in terms of computational cost and time [9]. Acknowledging these limitations is crucial for advancing the field, as it guides researchers in seeking improvements and developing more robust, transparent, and efficient methods. But in the end, the application of intermediate fusion techniques is still evolving, particularly in high-stakes areas such as disease diagnosis, treatment planning, and outcome prediction, where the precise modeling of complex biological interactions can significantly impact the decisions made by clinicians and healthcare providers [10].

1.1. Objective of the review

The primary objective of this systematic review is to comprehensively analyze and formalize the current intermediate fusion methods of MDL utilized in biomedical applications. The field of intermediate fusion in MDL, particularly within biomedical contexts, is still in its early stages. As such, there is a significant need for a structured analysis and for introducing a formal notation that not only categorizes and evaluates these methods, but also provides a framework that can be extended beyond the biomedical domain.

This systematic review is guided by a set of research questions, designed to shed light on the field of intermediate fusion in MDL, particularly within biomedical applications:

1. *Which are the intermediate fusion techniques currently employed in biomedical deep learning models?* This question aims to identify and categorize the various methods of intermediate fusion, highlighting their unique features and applications in the biomedical context.
2. *What are the existing challenges and limitations associated with these intermediate fusion methods?* This seeks to understand the drawbacks and areas where current methods fall short, providing insights into potential directions for future improvement.
3. *What are the future directions that should be pursued for intermediate fusion methods in biomedical applications?* This examination focuses on exploring potential directions for advancement in intermediate fusion, highlighting emerging trends, potential improvements, and unexplored areas that could significantly enhance the application and efficacy of these techniques in biomedical contexts.
4. *How can the structured notation proposed in this review enhance the understanding and application of intermediate fusion techniques beyond just biomedical applications?* This question explores the broader implications of the review's findings, considering how the proposed notation and categorization can be applied to other domains within deep learning.

Answering these questions will not only provide a thorough understanding of the current state of intermediate fusion in MDL for biomedical applications but will also offer a foundation for future research, extending the insights gained to a wider range of applications. A quantitative meta-analysis of performance metrics is beyond the scope of this review.

The findings and discussions presented in this review are particularly relevant to a diverse group of professionals and researchers:

- **Researchers:** This review holds significant value for researchers at the intersection of deep learning and biomedical applications. It offers comprehensive insights into the current state of intermediate fusion methods, their performance outcomes, and potential research gaps. Researchers can use this information to guide future investigations, develop more sophisticated models, and explore new paths in MDL.

Table 1

Comparison of this survey with previous multimodal learning surveys. The Scope column specifies the survey's application area, where Non-Specific refers to a broad, field-agnostic focus, while Biomedical Focus highlights surveys targeting healthcare and life sciences. The Modality column describes the range of data types covered, with Multiple indicating a wide variety of data and Limited denoting restricted types of modalities, often just one or two. The Fusion column distinguishes between general fusion methods (Early/Intermediate/Late Fusion) and those focused only on Intermediate Fusion. The Notation column indicates whether the survey provides a systematic notation to describe multimodal methods. A checkmark indicates that a specific focus or method is included in the survey, while a cross signifies that it is not.

Survey	Scope		Modality		Fusion		Notation
	Non-Specific	Biomedical	Multiple	Limited	All	Intermediate	
[2,11–22]	✓	✗	✓	✗	✓	✗	✗
[23–25]	✗	✓	✓	✗	✓	✗	✗
[26–32]	✗	✓	✗	✓	✗	✗	✗
This Survey	✗	✓	✓	✗	✗	✓	✓

- **Healthcare Professionals and Practitioners:** Healthcare professionals, including radiologists, pathologists, and clinicians who frequently engage with multimodal imaging or diagnostic techniques, will find this review informative. It provides an overview of the latest advancements in MDL and its potential impact on clinical practice. Understanding these fusion methods can assist healthcare practitioners in evaluating the reliability and applicability of deep learning models, ultimately enhancing patient care and diagnostic accuracy.
- **Deep Learning Community:** Beyond the biomedical domain, this review is also helpful for professionals and researchers working with MDL across various fields. The introduction of a structured notation and analysis of intermediate fusion methods is extendable to other applications, offering a versatile framework that can be adapted to different modalities and datasets. This broader applicability makes the review a crucial resource for anyone interested in the practical and theoretical aspects of MDL.

This survey is accompanied by two supplementary documents that contain all the key information discussed in Section 3 for each individual paper, so that in the main manuscript we can avoid the use of excessive citations and notations. Specifically:

- **Supplementary Material A:** includes a table listing all the fields discussed in Sections 3.2, 3.3, 3.4, 3.5, 3.6, 3.7. The table features each analyzed paper as a row and the discussed fields as columns. Readers can refer to this table for specific information about each paper in any of the analyzed fields.
- **Supplementary Material B:** provides a detailed analysis of the fusion mechanisms used in each article. For every article, it presents the corresponding mathematical notation, graphical representation, and categorization, all of which are defined in Section 3.4. While the main text explains this methodology using selected examples, readers can refer to this material for the complete formalization developed for each individual paper.

Both supplementary documents have been uploaded to our GitHub repository <https://github.com/cosbidev/Intermediate-Multimodal-Fusion-Bio>, ensuring easy accessibility for further exploration.

1.2. Comparison with previous surveys

This review distinguishes itself within the landscape of multimodal learning literature by focusing specifically on intermediate fusion methods in the biomedical domain. To clearly illustrate the novel contributions of our work, we present Table 1, which compares our survey with previous ones based on key features such as focus area, fusion methods discussed, modality coverage, and the introduction of structured notation. While existing surveys, out of the biomedical field, such as [2,11–22], have broadly addressed multimodal machine learning across various fields, they have not deeply explored intermediate fusion techniques. Similarly, reviews in the biomedical sphere, as [23–25], have touched upon the integration of multimodal machine learning but have not specifically honed in on intermediate fusion methods. In the context of biomedical applications, some previous surveys have

been limited in their categorization, focusing primarily on application domains rather than detailed methodological distinctions [26–29]. Additionally, certain reviews have concentrated solely on a limited set of modalities, often considering only two modalities [30–32]. Our survey, therefore, fills a crucial gap by not only concentrating on intermediate fusion within biomedical applications but also by introducing a structured notation system with potential applicability across diverse domains in deep learning.

2. Methodology

2.1. Inclusion/exclusion criteria

Our systematic review adheres to specific inclusion and exclusion criteria to ensure a focused and relevant analysis of intermediate fusion in MDL for biomedical applications:

- **Date:** We include studies published up to the present year (2024) to capture the most recent advancements in the field.
- **Language:** Only articles published in English are considered to ensure accessibility and standardization in analysis.
- **Study Design:** We exclude review articles, non-peer-reviewed sources such as opinion pieces, editorials, letters, and conference abstracts, as they do not provide original research or detailed methodology.
- **Target Application:** The focus is on articles that explore at least one biomedical application. Articles centered on non-biomedical applications or unrelated fields are excluded.
- **Modalities:** We include studies utilizing multiple modalities, considering both homogeneous (e.g., imaging) and heterogeneous types (e.g., imaging combined with tabular data). Studies focusing solely on a single modality or not involving the fusion of multiple modalities are excluded.
- **Fusion:** Articles employing deep learning architectures for fusion, particularly those using intermediate fusion in an end-to-end manner, are included. Articles primarily using early or late fusion methods or non-deep learning approaches are excluded.

2.2. Search strategy

Our systematic review implemented a comprehensive search strategy to identify relevant literature focusing on intermediate fusion methods in MDL within biomedical applications. The search strategy was meticulously crafted to encompass a wide range of pertinent studies while maintaining specificity to the review's objectives.

The literature search was conducted across multiple electronic databases, ensuring a thorough and diverse collection of articles. The databases included are: PubMed, IEEE Xplore, Scopus and Google Scholar.

To capture the breadth of literature relevant to our review, we utilized a set of synonyms and keywords across three main categories:

- Multimodal Deep Learning:

$$A = \left\{ \begin{array}{l} \text{“Multimodal Deep Learning”}, \\ \text{“Multimodal Fusion”}, \\ \text{“Multimodal Representation Learning”} \end{array} \right\}$$

- Biomedical Context:

$$B = \left\{ \begin{array}{l} \text{“Biomedical”}, \\ \text{“Medical”}, \\ \text{“Healthcare”}, \\ \text{“Clinical”}, \\ \text{“Biomedicine”}, \\ \text{“Medicine”} \end{array} \right\}$$

- Intermediate Fusion:

$$C = \left\{ \begin{array}{l} \text{“Joint fusion”}, \\ \text{“Intermediate fusion”}, \\ \text{“Mid-level fusion”}, \\ \text{“Feature-level fusion”} \end{array} \right\}$$

The search query was formulated using a combination of these synonyms, structured as follows:

$$\text{Query} = (A_1 \text{ OR } A_2 \text{ OR } \dots) \text{ AND } (B_1 \text{ OR } B_2 \text{ OR } \dots) \text{ AND } (C_1 \text{ OR } C_2 \text{ OR } \dots)$$

This structure allowed for a comprehensive yet focused search, ensuring that the results were pertinent to the scope of our review.

To understand the current state and trajectory of research in this area, we conducted a comprehensive search analysis by examining the publication numbers across various subdomains of multimodal learning, as shown in Fig. 1. We can observe that the significant drop of publications from “Multimodal Learning” to “Multimodal Deep Learning” highlights that, while multimodal learning is a broad field, the incorporation of deep learning techniques is a more focused area. This reflects the growing interest in leveraging deep learning’s feature extraction and representation learning capabilities within multimodal settings. The slight narrowing for “Multimodal Learning in Biomedical Applications” and for “Multimodal Deep Learning in Biomedical Applications” illustrates the increasing application of these methodologies in this field. Furthermore, we notice from the progression of publications in “Multimodal Deep Learning” to “Multimodal Deep Learning via Intermediate Fusion” suggests that intermediate fusion is an emerging and specialized area within the broader field of MDL. The fact that only 2% publications are categorized under “Multimodal Deep Learning in Biomedical Applications via Intermediate Fusion” indicates that this area, despite its promising applications, remains relatively underdeveloped. This gap underscores the need for further research and exploration. These results justify the necessity of your systematic review, which aims to lay the foundational groundwork by defining further research and development in this promising field.

2.3. Study selection

The study selection process for this systematic review was meticulously designed to ensure the inclusion of the most relevant and high-quality studies focusing on intermediate fusion in MDL within biomedical applications. This multi-stage process involved initial screening and full-text assessment.

1. **Initial Screening:** Initially, we reviewed the titles and abstracts of articles identified from the search. Articles that clearly did not meet the inclusion criteria, described in 2.1, were excluded at this stage.
2. **Full-text Assessment:** We retrieved the full texts of articles that seemed potentially relevant based on the initial screening. Each full-text article was assessed for eligibility, focusing on whether

they adhered to the inclusion and exclusion criteria. Articles that did not meet the criteria, described in 2.1, after a full-text review, were excluded. To ensure reliability and minimize bias, this process involved two independent reviewers, among the authors.

In cases of discrepancies regarding the eligibility of an article, these were resolved through discussion. We maintained a record of the reasons for excluding articles during the full-text assessment phase. To systematically track the number of articles at each stage of the selection process, a PRISMA flow diagram, shown in Fig. 2, was utilized. Following the described steps, we resulted in 54 manuscripts related to MDL in biomedical applications via intermediate fusion.

3. Intermediate fusion in biomedical applications

In this section, we present the various components that constitute a model employing intermediate fusion, as represented in Fig. 3. These components are presented in the order they appear in the model: beginning with the *modalities* as inputs (Section 3.2), followed by the *unimodal module* that extracts unimodal features (Section 3.3), the *fusion module* that combines these features (Section 3.4), the *multimodal module* that processes the fused features further (Section 3.5), and finally the *target* as the output of the model (Section 3.6). Then we also analyze the *learning* strategies used to train the model (Section 3.7). Each of these components is detailed in the corresponding subsections, where we introduce their taxonomies and review their state-of-the-art applications in biomedicine.

3.1. Definitions

Intermediate fusion, as already introduced in Section 1, within the domain of MDL, is an approach that involves extracting features from different modalities using specialized unimodal neural networks, and then merging these features into a fused multimodal representation. This fused representation is subsequently fed into another neural network to yield the final outcome. A distinctive feature of intermediate fusion is that, via the back-propagation of the loss, both the unimodal and multimodal modules are trained. This allows for continuous improvement in capturing the intra- and inter-relationships inherent in multimodal data. This methodology is particularly useful in complex fields like biomedicine.

This method, as shown in Fig. 3, can be broken down into several key components: modalities, unimodal modules, fusion module, multimodal module, and the target. Each of these plays a critical role in the overall functionality and efficacy of the fusion process:

- **Modalities:** Modalities, denoted as x_1, x_2, \dots, x_n refer to the different types of data or sources of information used in the model. In the context of biomedical applications, these could include imaging data, textual or tabular clinical notes, genetic information, or other relevant medical measurements.
- **Unimodal Module:** The unimodal modules, denoted as f_1, f_2, \dots, f_n refer to the parts of the network that processes each modality independently. Here, deep learning models are used to extract features, denoted as h_1, h_2, \dots, h_n , specific to each modality.
- **Fusion Module:** The fusion module \mathcal{F} is the core of intermediate fusion, which takes the features extracted by unimodal modules and combines them into a unified multimodal representation, denoted as h . This fusion can be performed using various techniques like concatenation, averaging, or more complex operations that account for interactions between modalities.
- **Multimodal Module:** After fusion, the combined features are fed into the multimodal module f . This part of the network processes the fused representation to extract insights that are only apparent when considering all modalities together.

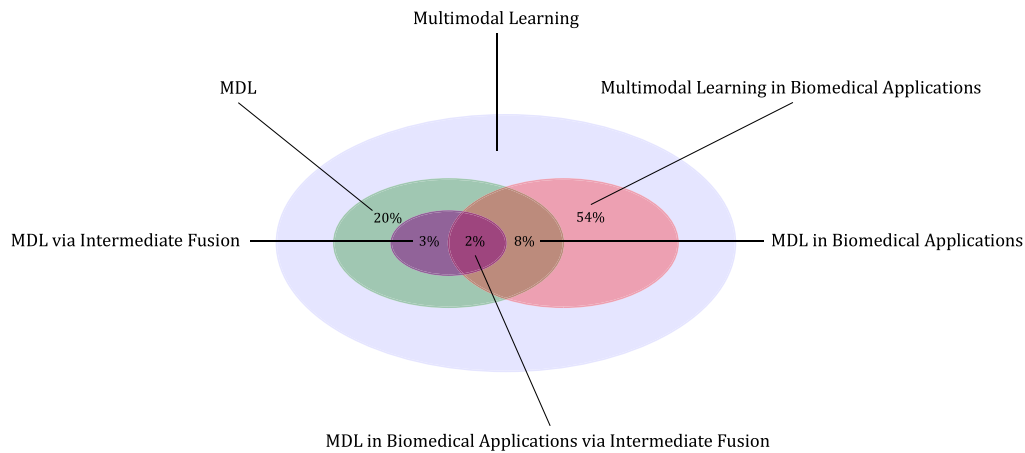


Fig. 1. Venn Diagram of Search Hits by Category. This diagram displays the percentage of search hits for each specified category, ranging from general fields to more specialized areas, i.e., “Multimodal Learning”, “Multimodal Learning in Biomedical Applications”, “Multimodal Deep Learning”, “Multimodal Deep Learning in Biomedical Applications”, “Multimodal Deep Learning via Intermediate Fusion”, “Multimodal Deep Learning in Biomedical Applications via Intermediate Fusion”.

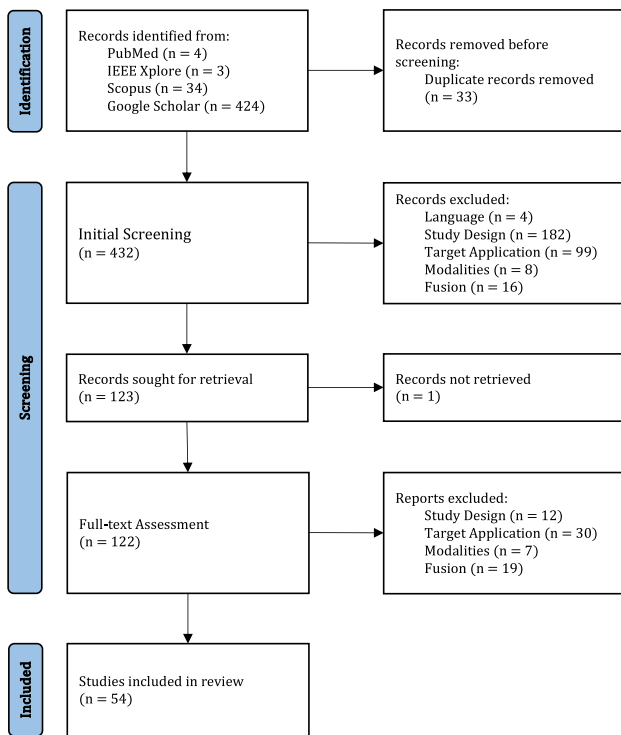


Fig. 2. PRISMA Flow Chart of Literature Selection Process. This flow chart outlines the systematic process of screening and selecting studies for inclusion in the review. It details the number of records identified, included, and excluded at each stage of the search and selection process, from initial database search through to the final included studies.

- **Target:** The target y refers to the output of the network. In biomedical applications, this could be a diagnosis, prognosis, or any other relevant medical outcome. The target guides the choice of the loss function, which in turn influences the back-propagation to iteratively improve the layers part of the unimodal, fusion and multimodal modules.

The subsequent sections will delve into the specifics of each component, analyzing their functions and significance in the context of intermediate fusion for biomedical applications.

3.2. Modalities

MDL with intermediate fusion represents a rapidly growing field in Artificial Intelligence (AI), where information from multiple modalities is integrated at intermediate stages of a deep neural network. Each modality provides unique information, enriching the model’s comprehension of context and thereby enhancing its performance. The different types of data offer complementary perspectives on complex medical scenarios, and their fusion provides a deeper understanding. For instance, medical image analysis can be integrated with textual data extracted from medical reports or laboratory data to obtain a more comprehensive view of a medical condition. Consequently, integrating multiple biomedical modalities strengthens predictive accuracy and fortifies models against input data fluctuations, amplifying their reliability and utility in practical applications.

3.2.1. Type of modalities

From our analysis of the 54 articles included in our survey, we identified 6 macro-modalities:

- **Imaging:** this modality encompasses various medical images that serve crucial roles in both diagnosis and ongoing monitoring of medical conditions. X-rays [33], Magnetic Resonance Imaging (MRI) [34], Computed Tomography (CT) [33], Positron Emission Tomography (PET) [35], Ultrasounds (US) [36] are included in this category.
- **Tabular:** under this modality fall tabular clinical data (spanning from demographic information to laboratory measurements), genomics data (which are used to analyze disease risk, genetic predisposition, or treatment response) and data related to drugs chemical features to analyze their functionality.
- **Text:** textual data comprises free-text reports and transcripts documenting question–answer dialogues between healthcare professionals and patients.
- **Time Series:** this category encompasses various types of signals. Specifically, we distinguish between electrical signals, clinical time series, and tactile signals. Electrical signals are derived from electrocardiograms (ECG) [37], electroencephalograms (EEG) [38], electroglottographs (EGG) [39], electromyography (EMG) [40], and electrodermal activity (EDA) [41]. Clinical time series typically consists of measurements recorded during patients’ hospital stays, which may include vital signs, laboratory results, and other metrics. Tactile signals, utilized in a single article [42], are recorded using a specialized tactile glove equipped with sensors.

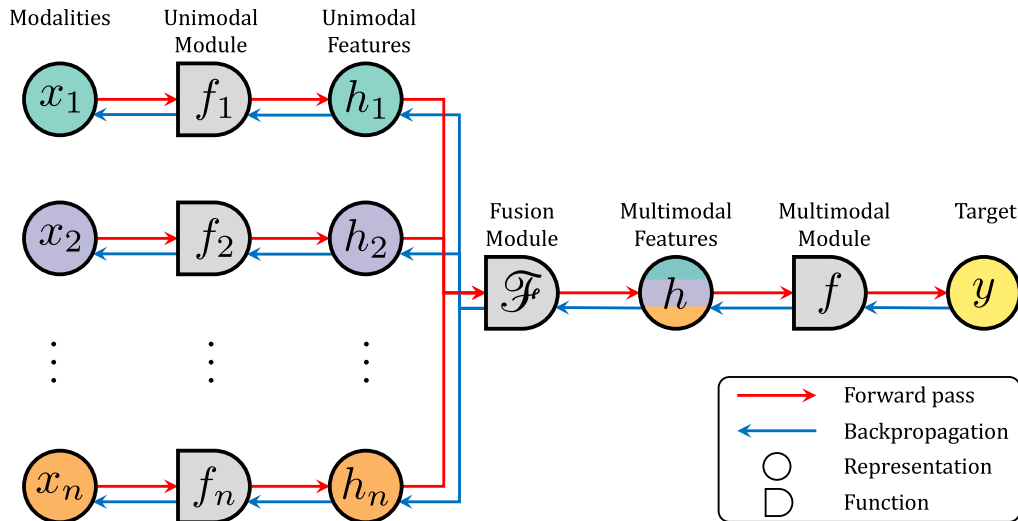


Fig. 3. Schematic Representation of Intermediate Fusion in MDL. This diagram illustrates the intermediate fusion approach within the framework of MDL. It begins with various data modalities x_1, x_2, \dots, x_n each represented in distinct colors (green, purple, orange), progressing through specialized unimodal modules f_1, f_2, \dots, f_n where individual features h_1, h_2, \dots, h_n are extracted, maintaining the color coding of their respective modalities. These features are then integrated by the fusion module \mathcal{F} , resulting in a multimodal feature representation h that is a blend of all the input colors, signifying the fusion of modal features. The subsequent multimodal module f processes these integrated features, aiming to achieve the specified target y shown in yellow. Red arrows indicate the forward pass through the network, *i.e.*, the inference phase, while blue arrows represent the backpropagation, facilitating the training of both unimodal and multimodal components. Different symbols distinguish between data representations and function operations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

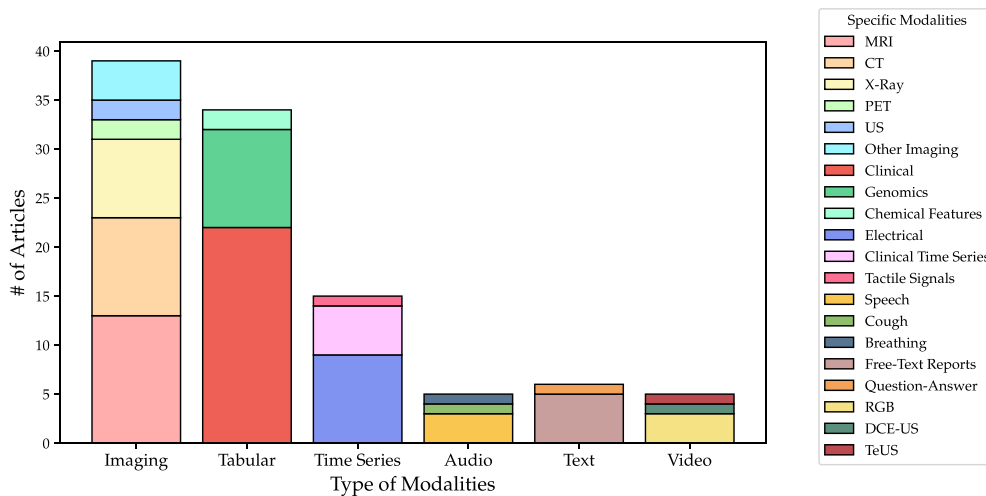


Fig. 4. Stacked bar plot showing the type of modalities used in the analyzed articles. The bar for each macro-modality is segmented, providing information on the specific modalities included.

- **Video:** video data can be related to imaging techniques over time, incorporating techniques like Temporal Enhanced Ultrasounds (TeUS) [43] and Dynamic Contrast-Enhanced Ultrasonography (DCE-US) [44], but it can also include diverse RGB videos capturing movements of hands or the entire body.
- **Audio:** it includes speech but also audio of coughing and breathing.

Fig. 4 illustrates the distribution of the aforementioned modalities used in the analyzed articles. It is observed that imaging and tabular data constitute the most frequently used modalities, representing 37% and 35% of the total modalities, respectively. Among imaging modalities, MRI scans are the most prevalent, accounting for 33% of cases, including variations such as functional MRI (fMRI) and structural MRI (sMRI). Following MRI, CT scans account for 26%, X-rays for 21%, and PET scans and US both for 5%. Other types of imaging data, including histopathology images, microflow imaging, 2D fundus images, 3D optical coherence tomography scans, and skin lesion images, are less

frequently used and collectively account for 10% of the imaging modality. Among the tabular modalities, the vast majority is represented by clinical data at 65%, followed by genomics data at 29% and chemical features at 6%. The prevalence of imaging and tabular modalities reflects the wide availability of diagnostic images and tabular data in the medical context, which we will also see in the publicly available datasets described in Section 3.2.3. Imaging technologies are crucial for diagnosing and monitoring various conditions, as they are well-integrated into medical practice, with established protocols and trained professionals to interpret the results. Tabular data are structured information essential for a comprehensive assessment and management of patients and for personalized medicine. Furthermore, it can be noted that time series represent 15% of modalities, suggesting a growing interest in temporal data analysis in the medical field. In particular, 60% of time series are represented by electrical signals, 33% by clinical time series, and the remaining 7% by tactile signals. On the other hand, text, audio, and video represent smaller percentages, 5%, 4%, and 4%,

respectively, indicating a lower availability of data in these forms or a lower applicability in the considered medical contexts. Indeed, while textual data is prevalent and highly informative in healthcare, it often requires substantial preprocessing to be useful for deep learning due to its unstructured nature and variability in terminology. Audio data is less frequently collected and stored compared to other data types. Although the use of audio in medical diagnostics shows promise for specific applications, such as the automated diagnosis of respiratory conditions from breath sounds [15], it is not yet as universally applicable or validated as other modalities like imaging. Video data requires significant storage and processing capabilities and is often specific to certain fields, such as surgery. In many medical contexts, other data types provide sufficient information for diagnosis and treatment, thereby reducing the reliance on video. Specifically, in text, audio, and video modalities, the most represented categories are free-text reports at 83%, speech at 60%, and RGB video at 60%.

In the articles we have analyzed, various modalities are combined through intermediate fusion techniques, enabling deep learning models to achieve a more comprehensive understanding of medical contexts. Section 3.2.7 will explore the different combinations of modalities utilized in these articles.

3.2.2. Data origin

Given the importance of modalities for MDL in biomedicine, we analyzed the origin of the datasets used in each of the 54 articles included in our survey. We specifically distinguished whether the data was real or synthetic and this analysis reveals a significant trend: a substantial majority of the studies rely on real data, while only a marginal proportion, specifically 2 out of 54 articles [45,46], also incorporate synthetic data generated by generative adversarial networks for models' training. This observation aligns with the expectation that real data are typically reliable and representative of reality, especially in the biomedical domain. At the same time, it underscores that synthesizing data using generative adversarial networks still poses an open technical challenge that researchers may choose to address to bridge the gap in big data in the biomedical domain.

3.2.3. Data source

The aforementioned real data sources mostly come from publicly available datasets. Indeed, only about 30% of the articles make use of private datasets. The presence of a significant number of articles utilizing public datasets suggests that accessible and usable data resources exist for research in this field, enabling researchers to conduct studies and develop models without necessarily having to create or acquire private data. Acquiring, annotating, and managing private datasets can be costly and time-consuming, and their use may sometimes be limited by privacy concerns. In contrast, utilizing existing public datasets can be a convenient and practical solution for many researchers, further promoting comparability and reproducibility in research.

A total of 52 distinct public datasets are used in the analyzed articles, with Table 2 presenting the ones used in at least two articles. Detailed information on public datasets used in only one article can be found in Supplementary Material A.

The datasets listed in Table 2 primarily include modalities within the macro-categories of imaging, *e.g.*, CT scans, X-rays, and MRI, and tabular data, *e.g.*, clinical and genomics data. As we explored in Section 3.2.1, these are the most commonly utilized modalities in the analyzed articles, given their prevalence in the medical domain. The table shows that the most commonly used datasets, *e.g.*, COBRE and MIMIC-CXR, appear in only 4 articles. This clearly indicates the absence of a benchmark dataset utilized in the majority of studies, making it difficult to establish the overall effectiveness of the models.

3.2.4. Sample size

A crucial factor to consider when training MDL models with intermediate fusion techniques is the sample size of the dataset being used,

i.e., the number of instances in a dataset. Intermediate fusion involves combining different input modalities into a shared layer within a neural network to enhance the overall performance of the model. This process is particularly complex and preferably requires large datasets so that the model has enough examples from each modality to learn the relationships between them effectively. Additionally, since intermediate fusion utilizes large networks with numerous weights, it is unsuitable for datasets with limited samples. The sample sizes of the datasets involved in the articles vary as follows: less than 100 samples for 7 datasets, 101 to 500 samples for 24 datasets, 501 to 1000 samples for 14 datasets, 1001 to 1500 samples for 6 datasets, and more than 1501 samples for 27 datasets.

These findings underscore the prevalence of moderate-sized datasets (with 101–1000 samples) in MDL research. However, a notable proportion (27 datasets) surpass the threshold of 1501 samples. These larger datasets hold promise for training robust models capable of capturing diverse variations, potentially yielding more generalizable outcomes. Conversely, 7 datasets comprise fewer than 100 samples, indicating that dealing with small datasets is still a challenge in MDL. Small datasets may lead to overfitting, where models perform well on the training data but fail to generalize to unseen data. Researchers working with small datasets may need to employ techniques such as data augmentation or transfer learning to mitigate this challenge.

3.2.5. Class distribution

In addition to the necessity of employing large datasets, if one intends to tackle a classification problem with intermediate fusion techniques, it is equally important to have balanced datasets. Balancing the diverse classes within the dataset is essential to prevent the model from being biased towards a particular class at the expense of others [81]. However, real-world data, especially in the medical field, often exhibit imbalanced distributions. Since intermediate fusion methods are particularly sensitive to class imbalance, we further analyze the articles that tackle classification tasks to determine whether they use balanced datasets. To this end, we employed the Gini-index [82], denoted as G , computed as:

$$G = 1 - \sum_{i=1}^{|C|} P(C_i)^2 \quad (6)$$

where C is the set of classes, $|C|$ is the number of classes, and $P(C_i)$ is the class prior probability. As a consequence, datasets can be defined as:

- **Balanced:** if $0.48 < G \leq 0.5$;
- **Imbalanced:** if $0.42 < G \leq 0.48$;
- **Highly Imbalanced:** if $0 \leq G \leq 0.42$.

As an example, in a binary classification problem, these ranges of Gini-index correspond to having a balanced dataset when the class priors range from 50–50 to 41–59, an imbalanced dataset when the class priors range from 40–60 to 31–69, and a highly imbalanced dataset when the class priors range from 30–70 to 0–100.

Fig. 5 graphically shows the distribution of balanced, imbalanced, and highly imbalanced datasets used in the reviewed articles. It is worth noting that the combined number of imbalanced and highly imbalanced datasets, 24 and 13, respectively, exceeds the total number of balanced datasets, which stands at 20. Only 37% of articles that use imbalanced or highly imbalanced datasets address the issue of class imbalance through the use of specific techniques. The preferred ones include the use of the Synthetic Minority Over-sampling Technique [83] as employed in [61–63,84], or the use of specific loss functions such as the Label-Distribution-Smooth Margin loss [85] used by [86], the weighted version of the Binary Cross-Entropy Loss used by [87], and the Focal Loss [88] used by [56,89]. However, the presence of studies based on unbalanced datasets highlights the need to pay more attention to this phenomenon. Without such caution, there is a risk of training models

Table 2
Publicly available datasets used at least in two articles on intermediate learning in biomedical research.

Dataset	# of Articles	Modalities	Used Modalities	Availability
COBRE [47]	4	fMRI, sMRI, clinical data, genomics data	[48–50]: fMRI, sMRI, genomics data [51]: fMRI, sMRI, clinical data	link
MIMIC-CXR [52]	4	chest X-ray, free-text reports	[53,54]: chest X-ray, free-text reports [55]: chest X-ray [56]: chest X-ray, free-text reports	link
ADNI [57]	3	MRI, clinical data, genomics data	[46,58,59]: MRI, clinical data	link
ASCERTAIN [60]	3	electrical signals (EEG, ECG, EDA), facial activity data	[61–63]: electrical signals (ECG, EDA)	link
CLAS [64]	3	electrical signals (ECG, EDA), photoplethysmography signals, accelerometer signals	[61–63]: electrical signals (ECG, EDA)	link
fBIRN [65]	3	fMRI, sMRI, genomics data	[48–50]: fMRI, sMRI, genomics data	link
MRPC [66]	3	fMRI, sMRI, genomics data	[48–50]: fMRI, sMRI, genomics data	on request
NSCLC-Radiogenomics [67]	3	CT, PET/CT, genomics data, clinical data	[68–70]: CT, genomics data	link
CPTAC-PDA [71]	2	US, CT, MRI, PET, histopathology images, clinical data	[72,73]: CT, clinical data	link
MIMIC III [74]	2	clinical data, clinical time series, free-text reports	[75]: free-text reports, clinical time series [76]: clinical time series, clinical data	link
MIMIC-IV [77]	2	clinical data, clinical time series	[55]: clinical time series; [56]: clinical data	link
TCGA-KIRP [78]	2	CT, MRI, PET, clinical data	[72,73]: CT, clinical data	link
TCGA-LIHC [79]	2	CT, MRI, PET, clinical data	[72,73]: CT, clinical data	link
TCGA-STAD [80]	2	CT, clinical data	[72,73]: CT, clinical data	link




Fig. 5. Distribution of Balanced, Imbalanced, Highly Imbalanced datasets.

that are less robust, characterized by biases in their predictions and overfitting on dominant classes.

3.2.6. Missing modalities

A common challenge in biomedical multimodal data is the presence of one or more missing modalities, an issue emerging when certain expected data modalities for a given sample are unavailable. This can happen due to various factors, such as registration errors, technical limitations in capturing specific data types, or simply because certain

modalities were not collected for all the patients in the dataset. In the context of MDL, it is essential to handle the presence of missing modalities appropriately, as this can pose difficulties in the data processing for the model. Despite the frequency of datasets with missing modalities in the biomedical field [24], only 4 articles [55,56,59,76] deal with datasets where a modality is missing to a certain extent, developing models specifically designed to handle this situation. More details are available in Section 3.7.3. Conversely, 6 articles addressed the issue through imputation techniques for missing modalities in




Fig. 6. Macro-modalities combinations in bimodal articles before (a) and after (b) feature extraction. The numbers on the edges represent the frequency of articles in which two modalities, indicated as nodes, are employed together. The width of the edges is drawn proportional to the corresponding number of articles. The sizes of the nodes are also arranged according to the number of articles that the modalities are used for.

the datasets. The imputation techniques employed include the use of neural networks and machine learning algorithms, as in [72,73], or the replacement of missing values with mean or predefined values, as in [75,84,90,91]. The remaining articles, however, do not address the issue of missing modalities: they use datasets without missing modalities or remove samples with missing modalities.

3.2.7. Modalities combination

We define the multimodal dimension characterizing the studies exploiting the number of modalities provided as input to deep learning models. We note that a significant proportion of studies (74% of articles) utilize bimodal datasets, indicating a prevalent trend in research. Additionally, 20% of the articles employ trimodal datasets, while a smaller fraction, only 6%, engage with high-modal datasets, characterized by the incorporation of more than three modalities. Fig. 6(a) illustrates the macro-modalities combinations within the datasets used in the bimodal articles. The numbers on the edges represent the frequency of articles in which two modalities, indicated as nodes, are employed together. The width of the edges is drawn proportional to the corresponding number of articles. The sizes of the nodes are also arranged according to the number of articles that the modalities are used for. We observe that the most prevalent combination is the imaging modality with the tabular modality, appearing in 16 articles, which accounts for 40% of the bimodal articles. Among these, the most common pairings involve CT scans and clinical data [72,73,87,89,92], followed by MRI scans with clinical data [46,58,59,90], and then CT scans with genomics data [68–70]. These findings align with expectations because, as outlined in Section 3.2.1, the imaging and the tabular are the most common macro-modalities and, within them, CT, MRI, clinical data, and genomics data are prevalent. The integration of imaging and tabular data offers a more comprehensive view of patients, with images providing detailed visual information on anatomy and tissue structures, while tabular data may contain clinically relevant information not found in images.

Fig. 6(a) shows that around 13% of bimodal articles, *i.e.*, 5 studies, utilize datasets combining two distinct types of imaging modalities. They include the combination of different MRI variants [93,94], CT and PET scans [95]. It is reasonable to speculate that different imaging modalities can provide complementary information on anatomical and functional tissue characteristics; their combined use can improve diagnostic accuracy by more clearly identifying anomalies or areas of interest and, in a deep learning model, can enhance predictive capability.

Additionally, 10% of the bimodal articles, *i.e.*, 4 studies, utilize datasets combining the imaging and text modalities. Notably, X-rays paired with free-text reports is the most prevalent combination [45,53,54], providing direct visual information from X-rays alongside corresponding medical interpretations.

Moving on to the trimodal domain, the most recurring macro-modalities combination is the use of two imaging modalities and one tabular modality, encompassing 55% of the trimodal articles. Within this subset, studies predominantly combine different MRI variants with genomics [48–50] or clinical data [51,96]. This trend is consistent with what happens in the bimodal domain, and the same observations apply.

Conversely, within the high-modal domain, no discernible trend emerges due to the limited number of articles falling into this category [97–99].

For more comprehensive insights into macro-modalities combinations and specific modalities combinations for each article, please refer to Supplementary Material A.

3.2.8. Feature extraction

As discussed in Section 3.2.7, understanding how modalities are combined is fundamental. Equally crucial, however, is considering the application of feature extraction techniques to data before they are fed into the models. Feature extraction is not part of the training process and should not be confused with the unimodal module, which will be described in Section 3.3. Feature extraction can play a pivotal role in preparing data for training deep learning models, involving the transformation of raw data into a set of representative features. These features serve as the actual input for the unimodal models and may belong to a different category of modalities than the original data. Fig. 6(b) illustrates the combination of macro-modalities in bimodal articles following feature extraction. Comparing it with Fig. 6(a), with which Fig. 6(b) shares the structure, it allows us to visualize the differences between before and after feature extraction. The emergence of the graph modality following feature extraction is immediately evident. In particular, there are 2 articles that, through the feature extraction process, transform the raw data, video and time series [100], and audio and time series [101], in the graph modality. An interesting trend is also the transformation of modalities into the tabular type after the feature extraction process [102]. In fact, as noted from Fig. 6(a), only 2 articles originally use datasets where both combined modalities are tabular. However, as shown in Fig. 6(b), the number of articles combining tabular data with other tabular data increases to 8 after feature extraction. In particular, this type of transformation is adopted




Fig. 7. Methods for extracting input features from various modalities. *Raw*: No feature extraction techniques applied; the raw data is used directly. *Handcrafted*: Features are manually engineered by experts to suit the problem domain, resulting in processed data used as input for the models. *Learned*: Features are extracted by training a neural network or machine learning model, which processes the raw data through forward and backward passes, providing the processed input data for the models.

by the 3 articles that used datasets with the combination of the time series modalities [61–63], by 2 articles that used an imaging and a tabular modality [69,70], and by one article that used a video and a time series combination [103]. Similarly, the text modality is converted into the tabular modality in 3 out of 5 articles [45,53,75], and this trend is also repeated in trimodal and high-modal articles. In particular, as declared in Section 3.2.7, in the trimodal field, the most frequently present combination of modalities in datasets is with two imaging modalities and one tabular modality but in 33% of the cases [48,49], after feature extraction, the data actually provided as input to the unimodal models all belong to the tabular modality. In the high-modal field instead, having only 3 articles, it is not possible to define a trend of the most used combination of modalities but, in one case [98], the original imaging modality used in combination with different tabular modalities is in turn converted to tabular before being provided as input to the unimodal model. This trend shows increased confidence among researchers in working with tabular data.

However, feature extraction is not always employed, and in these cases, data is used in its raw form as input for the unimodal models. Conversely, when feature extraction is applied, the input features for the unimodal models can be defined as handcrafted or learned, depending on the method of feature extraction utilized. Fig. 7 shows the 3 ways in which input features can be provided to models. In handcrafted feature extraction, features are typically engineered by human experts based on their understanding of the problem domain, which is crucial. Conversely, in learned feature extraction, features are automatically learned from raw data using a neural network or other machine learning models. In this case, instead of manually designing features, a model is trained with a forward and a backward pass to automatically discover relevant features from the data during the learning process.

The pie charts in Fig. 8 provide an overview of how feature extraction is performed for macro-modalities and for specific modalities belonging to the most represented macro-modalities, *i.e.*, imaging, tabular, and time series. Fig. 8(a) illustrates the 6 macro-modalities – imaging, tabular, time series, text, audio, and video – within the datasets before feature extraction. Each modality is represented proportionally within the innermost circumference, while the outer circumference is segmented to reflect the portion of data used in its raw form (denoted with R), the portion subjected to handcrafted feature extraction (H), and the portion undergoing learned feature extraction (L). Furthermore, for the 3 most represented macro-modalities – imaging, time series, tabular – Fig. 8(b), (c), (d), respectively, show how feature extraction is applied to the specific modalities included in the macro ones.

Looking at Fig. 8, the following emerges:

- In the context of imaging macro-modality, raw data is utilized in 80% of cases. In 16% of cases, concerning CT [69] and MRI [48–51], handcrafted feature extraction is employed. Only in 4% of cases, solely involving CT [70,98], the features are learned.

- For the tabular macro-modality, in 87% of cases, the used data are raw. In 13% of cases, all related to chemical features [99], features extraction is handcrafted.
- For the time series macro-modality, raw data are used in 50% of cases. Conversely, feature extraction is handcrafted in 45% of cases [61–63,100,103,104] and features are learned in 5% of cases [101]. The cases where data is not used in raw form exclusively involve electrical signals.
- For the text macro-modality, in 57% of cases features are raw, while in 43% of cases [45,75,90] features are learned.
- For the video macro-modality, in 40% of cases features are raw, in another 40% of cases features are learned [103,105] and in 20% of cases [100] features are handcrafted.
- For the audio macro-modality, in 60% of cases features are raw, in 20% of cases features are learned [101] and in another 20% of cases features are handcrafted [104].

Based on this evidence, we can state that the use of raw data is prevalent compared to the extraction of handcrafted or learned features. This is particularly evident in imaging and tabular modalities, where 80% and 87% of cases, respectively, use raw data. This suggests that deep learning models taking a single modality as input, described in Section 3.3, can directly process this type of modality without the need for feature extraction. For the remaining macro-modalities – time series, text, audio, and video – while the prevalence of raw data usage persists, the percentages are less skewed compared to imaging and tabular scenarios. This underscores the fact that the distinct nature of the data may necessitate different processing approaches. While imaging and tabular data lend themselves to more direct interpretation and benefit from raw data utilization, time series, text, audio, and video may require more extensive processing to extract relevant information.

After analyzing modalities, the following Section will provide a detailed examination of the deep learning models employed for processing unimodal data, prior to the intermediate fusion of the different modalities.

3.3. Unimodal module

The power of intermediate fusion lies in its ability to capture both intra- and inter-modal relationships inherent in multimodal data. In contrast to early fusion, each modality is processed through a dedicated branch denoted as f_i in Fig. 3, extracting modality-specific features, *i.e.*, the marginal representations denoted as h_i in the same figure. Extracting these marginal representations, facilitated by a global loss derived from fused features, allows for adaptability in relation to other modalities. This process empowers neural networks to derive representations that optimize the overarching goal. The rest of this section reports on the analysis of the unimodal networks adopted in the literature, shedding light on the preferences and trends observed in the field.

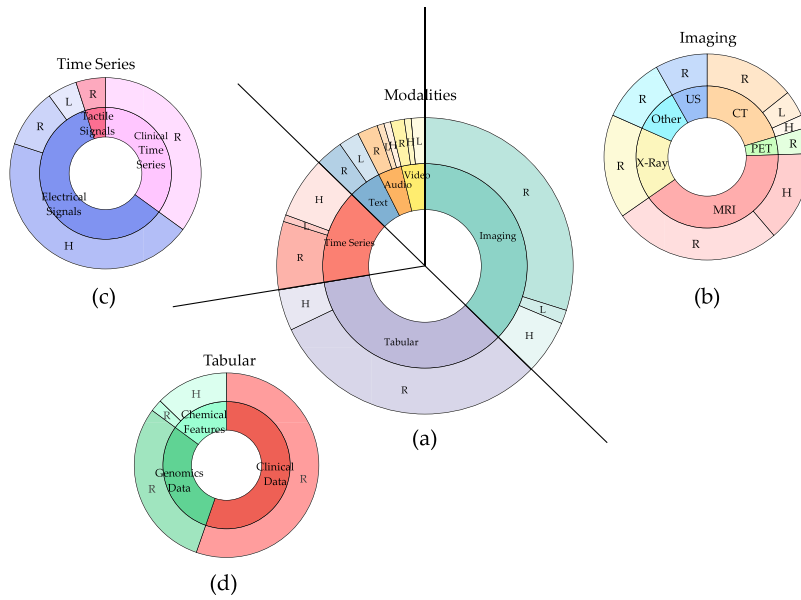


Fig. 8. Pie charts providing an overview of how feature extraction is performed for macro-modalities (a) and specific modalities belonging to imaging (b), time series (c), and tabular (d). R indicates that data are used in their raw form, H indicates that handcrafted feature extraction is performed, and L indicates that learned feature extraction is performed.

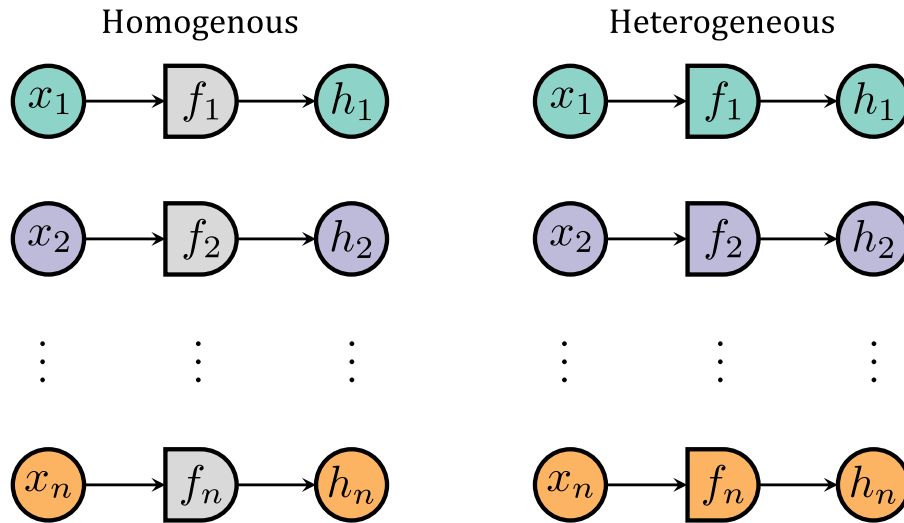


Fig. 9. Schematic representation of homogeneity: f_1, f_2, \dots, f_n represent unimodal modules and the colors represent the type of architecture/modality. Unimodal modules are homogeneous when all the modules have the same type of architecture (as in the figure left), and heterogeneous if there are at least two different types of architectures, modality-specific (as in the figure right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.3.1. Architecture types

The analysis of the literature suggests categorizing the reviewed papers into two classes: homogeneous and heterogeneous. We say that a multimodal model is homogeneous when all constituent unimodal networks share a uniform type of architecture; conversely, we say it is heterogeneous when at least one network diverges from the others, as schematically depicted in Fig. 9. In our inspection of the 54 articles within the scope of this survey, we discerned that 24 articles exhibit homogeneity, signifying uniformity across their unimodal networks, whereas 30 articles manifest heterogeneity, indicating the presence of diverse network types. Therefore, it is challenging to identify a trend in uniformity, as this classification appears intricately linked to the specific characteristics of input modalities. With the same type of modality, unimodal modules tend to have the same type of architecture and vice versa.

We identified six distinct types of architectures employed as unimodal modules in the literature:

- **Convolutional Neural Network (CNN):** Specialized for processing data with a grid-like topology, such as images, using convolutional layers to capture spatial hierarchies [106].
- **Fully Connected Neural Network (FCNN):** Composed of layers in which each neuron is connected to every neuron in the previous and following layers, suitable for learning non-spatial hierarchies such as those in tabular data [107].
- **Recurrent Neural Network (RNN):** Designed for sequential data, with the ability to retain information from previous inputs by looping the output back into the network [107].
- **Graph Neural Network (GNN):** Operating on graph structures, they leverage node and edge information to capture complex relationships and interdependencies [108].
- **Autoencoder (AE):** They aim to compress input into a lower-dimensional representation and then reconstruct it, used for tasks like denoising or dimensionality reduction [107].

- **Transformer (TF):** They utilize self-attention mechanisms to weigh the significance of different parts of the input data differently, excelling in parallel processing and handling sequential data without the need for recurrence [109].

Among the 54 articles investigated, 39 incorporated CNNs for at least one modality, 24 utilized FCNNs, and 10 employed RNNs. TFs, GNNs, and AEs were selected in 4, 3, and 2 articles, respectively. Furthermore, 5 articles opted to employ raw data without employing any unimodal network for at least one modality in their methodologies. These findings match with the fact that a large fraction of works use images and clinical tabular data, as shown in Section 3.2.1

3.3.2. Modality distributions

We now turn our attention to the specific modalities employed by these unimodal architectures. Fig. 10 illustrates the distribution of modalities across each architecture type. These modalities represent the original data type which, in some instances, undergoes feature extraction through either handcrafted methods or deep learning algorithms as described in Section 3.2.8. Regarding the utilization of architecture types across different modalities, the distribution is as follows: CNNs are utilized in 63 modules, with the imaging modality accounting for 38 of these. Time series data follows with 13 modules, tabular data is used in 5 modules, and audio, video, and text data are employed in 3, 3, and 1 modules respectively. Among these, 8 modules use handcrafted features for time series data, while 1 module each for text and video modalities employs learned features, with the remaining modules processing raw data. FCNNs are utilized across 41 modules, with 32 dedicated to tabular and 7 to the imaging modality. Time series and video modalities each account for one module. Given the inefficiency of processing raw images with FCNNs, a feature extraction step is typically employed, with handcrafted methods used in 5 instances and learned features in 2. For time series data, handcrafted methods are employed, whereas video modality features are extracted using deep learning methods. RNNs serve 11 modules, 5 of which process time series data, and the remaining 6 are equally divided between tabular and text modalities. Ten of these modules use raw data, with only one, pertaining to text data, employing learned features. GNNs are favored for 6 modules, including 2 each for imaging and time series data, and 1 each for audio and video modalities. While imaging modalities are processed in their raw form, the others use extracted features, with one time series modality using handcrafted methods and the other using learned methods. Audio features are extracted via deep learning, whereas video features employ handcrafted methods. TFs are utilized in 4 modules: 2 for tabular data and 2 for text data, with learned features employed for one of the text modalities and the rest using raw data. AEs are employed in 2 modules, both for the imaging modality; one uses raw data and the other handcrafted features. Lastly, in all 5 instances where raw data is processed without any unimodal architecture, the data belongs to the tabular modality. In summary, the exploration of unimodal architectures reveals a broad utilization across various data modalities, with a clear preference for certain architectures in specific contexts. CNNs dominate in versatility, especially in imaging, while FCNNs are preferred for tabular data, and RNNs excel in time series and textual data processing. GNNs and TFs show specificity in handling complex and sequential data structures, respectively. The deployment of these architectures is significantly influenced by the nature of the data, raw or processed through handcrafted or learned features, underscoring the importance of feature extraction and the interaction between data types and architectural choices in designing effective deep learning systems.

3.3.3. Architectural combinations

Another analysis that we have conducted involved determining the types of architectures used in conjunction as unimodal modules. Fig. 11 presents a graph illustrating the combinations of networks utilized. The

numbers on the edges represent the frequency of articles in which two architectures, indicated as nodes, are employed together. The width of the edges is proportional to the corresponding number of articles, and the sizes of the nodes are also arranged according to the number of unimodal modules that the specific architecture is used for. Finally, the pie charts inside the nodes represent the different types of modalities. As observed in the graph, the most prevalent connections are CNNs with themselves, followed by CNNs with FCNNs; 19 articles utilized at least two CNNs together as unimodal modules, while 13 articles employed one or more CNNs with one or more FCNNs. The next most frequent connection is within FCNN itself, employed by 7 articles. 5 articles preferred to use an RNN accompanied by an FCNN, and 4 by a CNN. In all the cases where raw data is used for at least one modality, CNNs are chosen for the other modalities. There are two occurrences of AE and FCNN used together, and two with AE and RNN. In fact, these two instances correspond to the same two articles [48,50] that employ three architectures together: FCNN, AE, and RNN. When examining the use of TFs, three articles chose to use them with CNN, while among the other two, one preferred to use them with an FCNN and the other with an RNN. One interesting observation is that GNNs are exclusively used in conjunction with each other and are not employed with any other type of architecture. This trend can be attributed to the specialized nature of graph-structured data and the need to maintain consistent graph representations throughout the model, preserving crucial structural information that other architectures might struggle to handle effectively. In conclusion, our analysis of the types of architectures used in conjunction as unimodal modules provides valuable insights into the trends and preferences within the field: CNNs emerge as the most commonly utilized architecture, often paired with themselves or with FCNNs, again owing to the prevalence of medical imaging and clinical tabular data which are mostly employed together as seen in Fig. 6. The exclusive pairing of GNNs with each other highlights their distinct role within the landscape of neural network architectures.

3.3.4. Dimensionality transformation

Unimodal modules are responsible for extracting features, often referred to as marginal representations, from individual modalities which are subsequently utilized by the fusion block. A primary advantage of this approach is the ability to represent high-dimensional inputs through more compact unimodal features, thereby condensing the information into a vector of reduced dimensionality. Conversely, it can also expand the dimension of the information by creating new information from the interaction of low-dimensional input features. The information is considered condensed if the marginal representation size is smaller than the input size, expanded if the marginal representation size is greater, and equal if the sizes are the same. Fig. 12 depicts a schematic representation of this concept.

Fig. 13 illustrates the degree of condensation/expansion, with the data categorized based on the architecture type and modality involved. It is calculated by dividing the marginal representation size by the input size. Among the articles surveyed for this review, only 39% provide information about the input size and marginal representation size for at least one modality. Fig. 13 depicts the total number of modules where the corresponding ratio is applied in these articles. The majority of modules condense information by a factor of one to two. CNNs notably condense more information than other architectures, particularly in imaging modalities, which inherently entail larger sizes. Conversely, smaller condensation ratios, or even expansions, are predominantly observed in modules utilizing FCNNs with tabular modalities. This observation aligns with the nature of tabular data, primarily composed of clinical and demographic information with substantially smaller sizes compared to other modalities. There are only two papers [58,59], both implemented by the same author, that prefer a marginal representation dimension equal to the input dimension, hence a ratio of 1, both utilizing an FCNN for tabular data. Several modules exhibit extreme expansion ratios, a less common occurrence in the literature. One

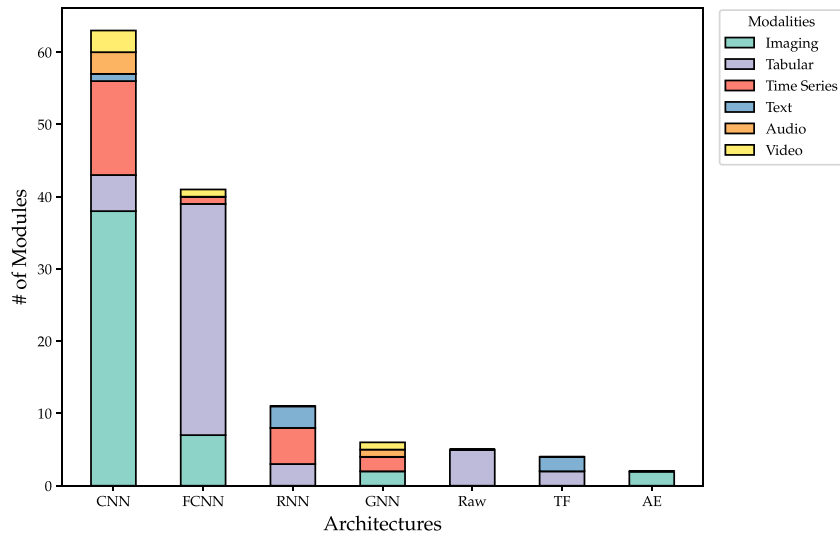


Fig. 10. Distribution of modalities according to type of architectures.

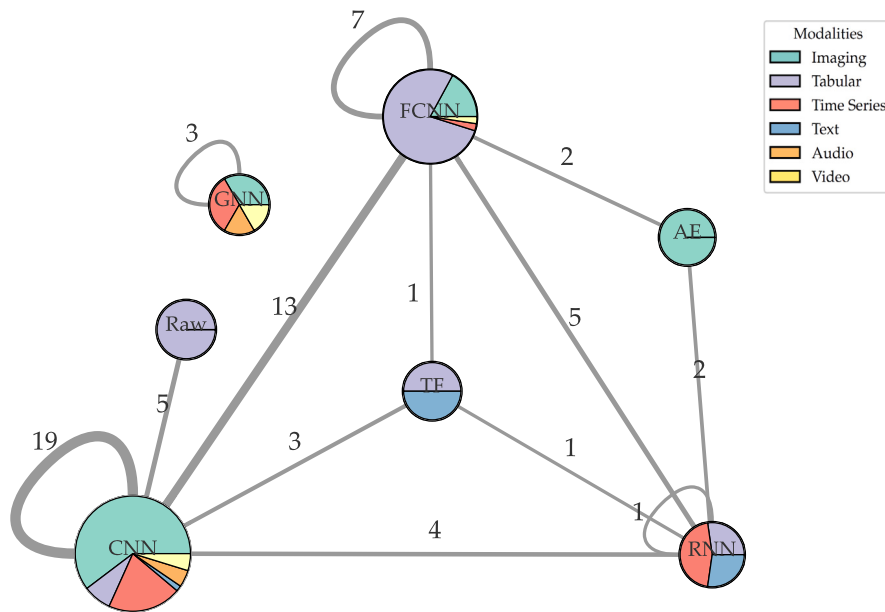


Fig. 11. Network graph illustrating the interconnectivity between different types of architectures. The numbers on the edges represent the frequency of articles in which two architectures, indicated as nodes, are employed together. The width of the edges is drawn proportional to the corresponding number of articles and the size of the nodes is arranged according to the number of modules that the architecture is used for. The pie charts inside the nodes indicate the proportions of the different modalities that are used with the corresponding architecture.




Fig. 12. Schematic representation of condensation/expansion: x_1 represents the input and h_1 represents the marginal representation. The size of the circles represents the size of the vectors. The information is considered condensed if the size of the marginal representation is smaller than the input size, and expanded if it is larger. The dimension of the information remains the same if the size of the marginal representation is equal to the input size.

contributing factor to this phenomenon is the endeavor to maintain equal marginal representation sizes for each modality. For example, in the case of [90], where one modality is imaging and the other is tabular data, the imaging input size is 50,000 while the input size of the tabular data is only 18. The authors opt for an equal marginal representation size of 512. Consequently, while the condensation ratio for the imaging

modality is 0.01, the expansion ratio for the tabular modality is 28.44. Similarly, in [56], the tabular modality has a considerably smaller size (15) compared to the other two modalities, imaging (50,176) and text (393,216), with authors again selecting an equal marginal representation size of 512. Consequently, the tabular modality exhibits an expansion ratio of 34 alongside condensation ratios of 0.01 for




Fig. 13. Histogram of the condensation/expansion ratios divided according to architecture (bars on the left) and modality (bars on the right).

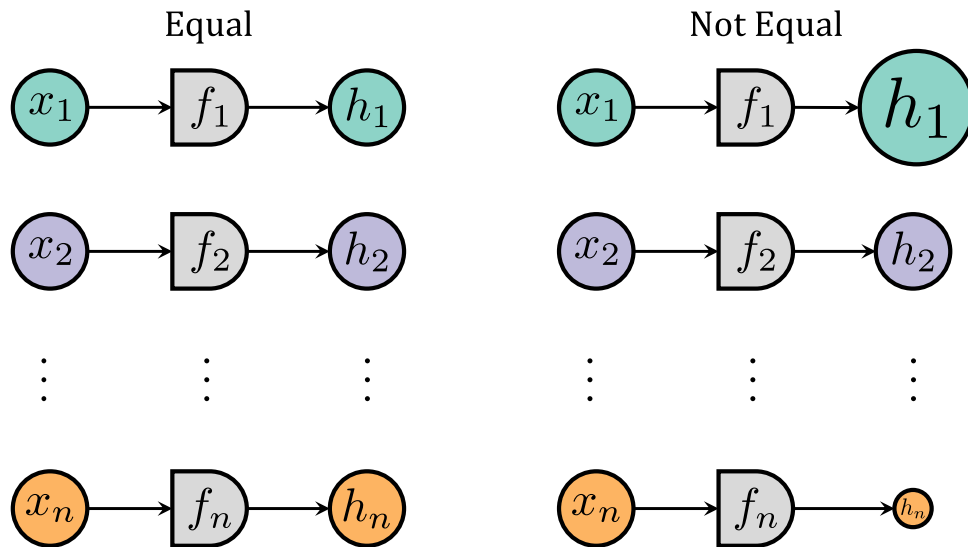


Fig. 14. Distribution of the marginal representations: h_1, h_2, \dots, h_n represent the marginal representations and the size of the circles represents the size of the feature vectors. Marginal representations are distributed equally if all of them have the same size and not equally if at least two of them have different sizes.

imaging and 0.001 for text. Fig. 14 represents a schematic of the distribution of the marginal representations. Marginal representations are distributed equally if all of them have the same size and not equally if at least two have different sizes.

In fact, many studies prefer equally distributed marginal representations. Among the 18 articles providing marginal representation sizes for all modalities, only five [55,58,59,91,110] feature non-equally distributed marginal representations. Cahan et al. [87], who prefer to have equal marginal representation sizes for each modality, state that the rationale behind this approach is to prevent any unintended bias towards one modality over others. While this rationale is sensible, it is uncertain whether all the different modalities can be represented with an equal number of features. All in all, the degree of information condensation/expansion depends on the input modality and the architecture that processes the corresponding input. Our analysis indicates that the imaging modality, when processed with CNNs, exhibits the highest condensation ratios, while tabular data processed with FCNNs demonstrates the highest expansion ratios.

Unimodal representations are extracted using separate unimodal modules. These representations are then forwarded to the next step,

which we refer to as the fusion module, where they are fused to create a common multimodal representation. The following section presents the newly proposed notation and taxonomy alongside the analysis of the fusion modules adopted in the literature.

3.4. Fusion module

Intermediate multimodal fusion involves integrating data from multiple modalities (e.g., text, image, audio, video) at an intermediate data representation of the modalities rather than at the data or the output level, as in early and late fusion approaches respectively. This approach enables an interaction between the modalities that are inherently unattainable in early fusion methods, where the early integration of modalities hinders the disentanglement of contributions for each data type. Likewise, it overcomes the limitations of late fusion techniques, which often rely on independently learned disjoint spaces for each modality.

At the core of the intermediate fusion approach there is the fusion module, denoted as \mathcal{F} in Fig. 3: it fuses the hidden representations extracted from the n unimodal modules h_1, \dots, h_n into the unified

multimodal representation h . In other words, the fusion module aims to map the features of different modalities into a common embedding space, which can be obtained through different strategies exploiting the hidden representation spaces of the modalities. These approaches, characterized by different degrees of abstraction and characterization of the feature spaces, intend to maximize the complementarity of the modalities during the learning phase. Indeed, there is no unified approach to maximize the informativeness of the extracted representations, but each approach uses more or less complex fusion structures to this end. These may involve different fusion operations, achieved by a single fusion or multiple fusions, which may combine all available information, eventually also using h_1, \dots, h_n multiple times.

Recognizing the central role of the fusion module we hereby introduce a formal notation specifically designed to highlight the fusion process detailed in the next Section 3.4.1. This notation framework aims to facilitate a thorough understanding of the critical aspects of multimodal fusion, namely *what*, *how many*, *when* and *how* to fuse: these three issues will be further analyzed in Sections Section 3.4.3, 3.4.4, 3.4.5 and 3.4.6. Hence, we will categorize the existing intermediate multimodal fusion approaches and offer a framework to distinguish the various architectural design choices available when designing an intermediate fusion multimodal model. It is worth noting that our notation framework applies beyond the medical domain, *i.e.*, the focus of this survey. Moreover, Section 3.4.7 explores the correlation between the unimodal modules and the fusion module architecture adopted.

3.4.1. Systematic notation for multimodal fusion

We now introduce a comprehensive mathematical notation system to facilitate a rigorous analysis of the multimodal intermediate fusion strategies employed in the revised studies. It enables the detailed description and precise representation of the nature, localization and timing of the fusion mechanisms designed in the architectures. A fusion is defined as the operation that integrates input data, as modalities and/or previous fusions, producing a unified representation.

We denote a generic single fusion as:

$$\mathcal{F}_i = \bullet(\alpha_j^l, \alpha_k^m, \dots)_{\rightarrow} \quad (7)$$

where the components of the notation are:

- **Fusion function** (\mathcal{F}): identifies the fusion function that takes as input multiple feature representations, *e.g.*, modalities or outputs of other fusions, and returns a single fused representation.
- **Fusion ordering** (i): Denotes the position of the fusion within the architecture, following a sequential progression from input to output. Therefore, the highest fusion ordering index denotes the number of fusions performed in the fusion module.
- **Inputs** (α_j, α_k , etc.): Represent the input modalities that can be the sample x represented in the original feature space, or the output of previous fusions \mathcal{F} . For example, two different inputs j and k could be represented as x_j and x_k , \mathcal{F}_j and \mathcal{F}_k , or a combination such as x_j and \mathcal{F}_k .
- **Layer depth** (l and m): Indicate the number of trainable layers through which inputs α_j and α_k have been processed before their integration. This indicates when the fusion of information occurs. An example could be a fusion function that takes as input α_j^2, α_k^3 , implying that the two input representations passed through two and three layers, respectively, before the fusion.
- **Fusion operation** (\bullet): Symbolizes the type of operation performed between inputs α_j, α_k , etc. To standardize the analysis, we identified 5 fusion categories: *Concatenation*, *Attention*, *Tensor-operation*, *Calibration*, and *Knowledge-sharing*, each represented by specific symbols listed in Table 3 and further detailed in Section 3.4.6. The analysis of these operations provides insight into how each fusion is performed.
- **Final Fusion** ($_{\rightarrow}$): if present, it specifies that the considered fusion function is the last in the architecture and its output coincides

Table 3
Symbols representing the fusion categories.

Fusion Category	Symbol
Concatenation	\oplus
Attention	Δ
Tensor-operation	\boxplus
Calibration	$\$$
Knowledge-sharing	\equiv




Fig. 15. Graph representation of the fusion strategy applied in [110], used as example to present our formal notation. Here, the modalities and the fusion nodes are represented in red and gray, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with the latent fused representation h , given as input to the multimodal module f . Conversely, if not present, it denotes that the fusion under consideration is at the beginning or in the middle of the fusion structure.

After identifying all the components used in the model, it is possible to draw a directional bipartite graph $G(N, E)$ representing how the modalities and fusions are combined in each fusion approach, where N is composed of 2 types of nodes, modalities and fusions, respectively, whilst E is composed of directed edges which connect the inputs to a fusion.

To better illustrate the use of the proposed notation in the analysis of a multimodal fusion approach, we examine the intermediate fusion strategy from [110], which integrates 3 modalities, specifically MRI, X-ray and tabular data, represented with x_1, x_2 , and x_3 , respectively, through the following fusion operations:

$$\begin{aligned} \mathcal{F}_1 &= \oplus(x_1^{21}, x_2^{25}) \\ \mathcal{F}_2 &= \oplus(\mathcal{F}_1^0, x_3^0)_{\rightarrow} \end{aligned} \quad (8)$$

This structure features 2 fusion operations: \mathcal{F}_1 and \mathcal{F}_2 . The first fusion \mathcal{F}_1 concatenates the modalities x_1 and x_2 after processing them through 21 and 25 convolutional layers, respectively. The second fusion \mathcal{F}_2 combines the output of \mathcal{F}_1 , directly without any additional trainable layer, with the x_3 modality. The latter modality, since it does not go through any unimodal module, is used in its raw form, x_3^0 . Fusion \mathcal{F}_2 yields the final multimodal representation of the module, as indicated by the \rightarrow symbol, which is then fed to the multimodal module. This complete set of operations inside the fusion module is also represented by its graph in Fig. 15.

We adopt this notation and graph representation for the comprehensive analysis of all 54 articles reviewed, which is reported in Supplementary Material B. This document examines each method offering four key pieces of information: first, we report the model image, taken from its respective article; second we formalize the fusion strategy with our approach thus reporting a set of equations; third, we offer the graph detailing the fusion structure; fourth, it also reports further taxonomies we have defined, which will be detailed in the next section.

3.4.2. Fusion module taxonomy

The systematical classification of the reviewed fusion methodologies by using the notation introduced so far not only offers to the readers a systematic analysis and categorization of the literature, but it also




Fig. 16. Structures types: (a) Single Sudden; (b) Multiple Sudden; (c) Multiple Gradual; (d) Multiple Multiflow. First row: all Sudden architectures; Second row: gradual architecture. The dashed gray arrows for the Multiple Sudden, Multiflow and Gradual indicate that each node might be used as additional inputs for subsequent fusion nodes.

permits us to highlight *what* is fused, *how many* times the fusion is performed, *when* the fusion occurs, and *how* the fusion is performed.

These categorizations are intended to group the fusion methods according to their fundamental characteristics and principles, thus seeking an answer to some fundamental questions about intermediate fusion strategies for multimodal inputs:

- What elements are involved in the fusion?
- How many times does the fusion occur between the modalities?
- When is the fusion performed between the modalities?
- How is the fusion among the modalities performed?

We aim to provide a coherent and generalized understanding of the main fusion approaches applied in the field, facilitating a deeper comprehension of their implications and applications in multimodal research and proposing suitable and coherent taxonomies for the classification and categorization of any work in this field. However, the proposed taxonomies may not cover all the possible multimodal intermediate fusion methodologies, but it is applicable to the works considered in this review. Nevertheless, the introduction of our notation system aims at simplifying the analysis of any multimodal fusion method, thus making more straightforward the extension of the taxonomies presented in this review.

In the subsequent sections, we present our proposed taxonomies for multimodal fusion strategies. To illustrate these taxonomies, we draw upon select articles from the literature as examples. However, it is important to note that our discussion of these examples is intended to be illustrative rather than exhaustive. We aim to provide sufficient context to elucidate our classification system, while an exhaustive in-depth analysis of each specific approach can be found in Supplementary Material B.

3.4.3. What?

Regarding what is being fused, in Section 3.2.1, we identified the main categories of modalities, including imaging, tabular, text, time-series, video, and audio. Furthermore, our proposed notation allows for

clear visualization of the fusions performed between different modalities and other multimodal representations before getting to the final output of the proposed model. For a more comprehensive understanding, readers can refer to Supplementary Material B, where we have specified the modalities used by each fusion function in the respective methods.

3.4.4. How many?

Going forward, we quantitatively assess the fusion strategies adopted within the corpus of the reviewed literature, by categorizing the architectures based on whether they opt for a *single* fusion mechanism or *multiple* fusions throughout the model's structure. The predominance of *single* studies (35 out of 54) demonstrates a preference for this simpler approach over those employing *multiple* fusions (19 out of 54). This observation underscores a preference towards intermediate fusion strategies that favor simpler and straightforward modalities integration, relying on robust unimodal representations. This trend is related to the vast typologies of data adopted in multimodal fusion in the biomedical field, where data curation and processing play a central role, and the fusion techniques are not fully investigated. Furthermore, these single fusion approaches often investigate more sophisticated multimodal modules able to analyze complex multimodal representations; for an in-depth discussion on the characteristics and implications of these multimodal modules, refer to Section 3.5.

3.4.5. When?

Upon evaluating when the fusion occurs among different modalities within multimodal architectures, two primary categorizations emerged from our analysis.


The first categorization focuses on the timing of the fusions that integrate the different modalities: whether all modalities are fused simultaneously in a single operation, gradually across multiple ones, or if there are multiple components of the network elaborating the same inputs. Firstly, we distinguish *sudden* architectures, characterized by the simultaneous fusion of all modalities in a single fusion, from *gradual* architectures, where modalities are fused progressively and

incorporated in subsequent and distinct fusions. Secondly, we identify the *multi-flow* categorization as a specific group of methods that employ multiple sequences of fusion operations that process information from a set of considered modalities, defined as *fusion flow*. Therefore, the *multi-flow* structure is characterized by the generation of multiple distinct multimodal representations, each reflecting a unique integration and elaboration of all the input modalities, which are then fused into a unique fused representation. The graphical representation of single sudden, multiple sudden, multiple gradual and multi-flow strategies are shown in Fig. 16 in panel a, b, c and d, respectively.

To elucidate the application of the above categorization, we present an example drawn from the list of reviewed articles for each category of the taxonomy.

The set of fusions in the architecture proposed in [95] can be represented as follows:


$$\begin{aligned} \mathcal{F}_1 &= \mathbb{A}(x_1^4, x_2^5) \\ \mathcal{F}_2 &= \mathbb{A}(x_1^3, \mathcal{F}_1^1) \\ \mathcal{F}_3 &= \mathbb{A}(x_1^2, \mathcal{F}_2^1) \end{aligned} \quad (9)$$



Here, \mathcal{F}_1 qualifies the approach as sudden due to the integration of both modalities x_1 and x_2 in a single fusion. Subsequent fusions \mathcal{F}_2 and \mathcal{F}_3 , which integrate modality x_1 with previous fusion outputs, do not alter this classification, as the initial fusion already incorporated all the considered modalities, even if \mathcal{F}_2 and \mathcal{F}_3 are adopted to fuse different unimodal representations (x_1^3 and x_1^2) and other fusion functions (\mathcal{F}_1^1 and \mathcal{F}_2^1).

Conversely, the architecture presented in [97] exemplifies the gradual approach and can be described with the following fusion functions:


$$\begin{aligned} \mathcal{F}_1 &= \mathbb{M}(x_1^1, x_2^2) & \mathcal{F}_2 &= \mathbb{M}(x_2^3, x_3^1) \\ \mathcal{F}_3 &= \mathbb{M}(x_2^3, \mathcal{F}_1^2, \mathcal{F}_2^2) & \mathcal{F}_4 &= \mathbb{M}(x_4^x, \mathcal{F}_3^2) \\ \mathcal{F}_5 &= \mathbb{M}(x_5^0, \mathcal{F}_4^7) \end{aligned} \quad (10)$$



\mathcal{F}_1 fuses the modalities x_1 , x_2 and x_3 , while the fusions \mathcal{F}_4 and \mathcal{F}_5 gradually integrate x_4 and x_5 , respectively, thus depicting the gradual approach for the integration of the modalities. Through this fusion approach the integration between the modalities is performed at different abstraction levels of the architecture, resulting in a hierarchical processing ability of the model.

As an example of multi-flow method, in [104] the authors proposed a multi-flow architecture, whose fusions can be represented as follows:

$$\begin{aligned} \mathcal{F}_1 &= \mathbb{M}(x_1^1, x_2^1) & \mathcal{F}_2 &= \mathbb{M}(x_1^2, x_2^2) \\ \mathcal{F}_3 &= \mathbb{M}(x_1^3, x_2^3) & \mathcal{F}_4 &= \mathbb{M}(x_1^4, x_2^4) \\ \mathcal{F}_5 &= \mathbb{M}(\mathcal{F}_1^3, \mathcal{F}_2^3, \mathcal{F}_3^3, \mathcal{F}_4^3) \end{aligned} \quad (11)$$




This model, integrating 2 modalities, x_1 and x_2 , through 4 parallel and independent fusions \mathcal{F}_1 , \mathcal{F}_2 , \mathcal{F}_3 , and \mathcal{F}_4 , exemplifies a multi-flow methodology. Indeed, it presents 4 distinct fusion flows, each elaborating all the modalities on different abstraction levels of the unimodal representations modules. Finally, \mathcal{F}_5 integrates each different fusion flow in a single representation, denoting a multi-level fusion features space.

The second categorization delves into the abstraction level of the modalities at which the fusion occurs, as denoted by the superscript index indicating the number of trainable layers each modality, or fusion, has passed through before being fed to the configured fusion module. We differentiate between *synchronous* and *asynchronous* fusion methods. The former is characterized by the processing of each modality through an identical number of trainable layers prior to fusion, as indicated by matching superscripts. Conversely, the latter integrates modalities that have passed through different numbers of trainable layers, thus fusing at varying levels of abstraction. It is worth noting that the gradual fusion approach intrinsically aligns with asynchronous methodologies: the sequential integration of modalities at distinct stages inevitably leads to varying degrees of abstraction among the modalities being fused, thereby precluding a synchronous fusion process for gradual architectures.

Examples of synchronous and asynchronous methods are found in [103] (Eq. (12)) and [111] (Eq. (13)), respectively. Despite their use of the same modalities and fusion function, the distinction lies in the superscript numbers for each modality before fusion:

$$\mathcal{F}_1 = \mathbb{M}(x_1^1, x_2^1) \quad (12)$$

$$\mathcal{F}_1 = \mathbb{M}(x_1^3, x_2^2) \quad (13)$$



The synchronous method [103], integrates modalities x_1 and x_2 after they have both been processed through a singular trainable layer, ensuring uniform abstraction levels at the fusion point. Conversely, the asynchronous method [111], processes the modalities x_1 and x_2 through different unimodal modules, composed of 3 layers for x_1 and 2 for x_2 before their fusion.

To summarize the considerations reported so far, Table 4 reports the list of the reviewed articles categorized in sudden, gradual and multi-flow, and distinguished by the adoption of single and multiple fusions approaches. On the one hand, the single fusion approaches are all categorized in sudden strategy, given the adoption of single fusion operation; on the other hand, we observe that most of the multiple methods apply the sudden strategy (11 articles) over the more complex gradual (5 articles) and multi-flow (3 articles) approaches. This trend highlights the preference for merging the modalities all at once with sudden approaches in the biomedical field, underlying the will to focus intermediate fusion learning on fine-grained and exhaustive unimodal features, simplifying the subsequent fusion operations. The prevalence of this approach may be attributed to the absence of a priori knowledge regarding modality hierarchies in the biomedical domain, precluding the direct establishment of a definitive fusion strategy. Consequently, the proposed methodologies predominantly leverage well-established unimodal architectures capable of extracting salient information, rather than exploring more intricate fusion strategies. This conservative approach mitigates the risk of introducing unintended biases during the training process, which could potentially compromise the efficacy of multimodal feature learning.

Within single sudden fusion architectures, asynchronous methods predominate (24 articles out of 35), compared to synchronous instances. This trend reflects the widespread adoption of a single fusion

Table 4

Classification of articles by the number of fusion (single/multiple), and fusion timing categorization (sudden/gradual/multi-flow). The articles with the symbol ‡, incorporate a synchronous fusion strategy, whereas those without the symbol use an asynchronous approach.

When?	How Many?	# of Articles	References
Sudden	Single	35	[42,48–50,53–56,69,72,73,76,87,89–92,111–117], [45,61,70,75,84,98,99,103,105,118,119]‡
	Multiple	11	[68,95,96], [62,63,93,94,100,101,120,121]‡
Gradual	Multiple	5	[51,58,59,97,110]
Multi-flow	Multiple	3	[46,86,104]

to fuse different data types (23 articles out of 35), such as imaging and tabular data, which typically require different preprocessing and feature extraction processes, resulting in heterogeneous unimodal architectures. In contrast, architectures that implement multiple fusions exhibit a more balanced distribution, with 10 asynchronous and 9 synchronous implementations.

3.4.6. How?

We identify five types of approaches that can be used to fuse the data, that we already introduced in Table 3. They are: *Concatenation*, *Attention*, *Tensor-operation*, *Calibration*, and *Knowledge-Sharing*. In the following, we deepen each one as we remark on the need to carefully consider the method's selection and implementation. Indeed, understanding the strengths and limitations of each fusion class is a crucial step for researchers willing to make informed choices that optimize the integration of unimodal data, potentially leading to more sophisticated and effective multimodal systems.

Concatenation. It is by far the most commonly used approach in the reviewed methods, where multiple input representations or features are combined by joining them along a specific axis or dimension [122]. This method creates a single, larger representation that preserves all the information of its original components. Notably, the majority of single fusion works (29 out of 35) utilize concatenation to merge features processed by unimodal modules. An example is [48], where the authors used concatenation to merge 3 unimodal representations derived from structural and functional neuroimaging and genomic data. However, this approach introduces several limitations for multimodal methods, particularly concerning the high dimensionality of the fused feature space. Indeed, in studies that employ multiple fusions, concatenation is generally adopted as the primary technique to elaborate the intermediate fused representations, thereby inheriting the same limitations identified in single fusion methods. An example is [104], which introduces a multi-flow network where each flow is represented by a concatenation function that merges unimodal features at different levels of abstraction. However, most of the studies (13 over 19) employ concatenation for a structural necessity, because they use it to produce the final version of the fused features after all the other fusion operations have taken place. This further processing of the representations usually results in a reduced and more compact form with respect to those methods that employ the concatenation immediately after the unimodal modules, limiting also possible dimensionality issues and poor correlation between the fused representations.

Tensor-operation. Tensor-operation represents a versatile category within our categorization, and groups together a wide range of mathematical strategies for synthesizing multimodal representations from interactions between unimodal representations [122]. These operations, e.g., Hadamard product, outer product, pooling operations, and aggregation functions, as minimum, maximum and average, enable the generation of representations with varying dimensionalities, potentially lower, equal, or greater than those of the input representations.

In single fusion approaches, only a few studies [45,55,84] delve into this fusion category. For instance, in [45] the authors present an approach that employs an element-wise interaction between tensors of imaging features from chest X-rays and textual features from radiology

reports, that ensures an element-wise interaction between the different modalities involved. This technique generates an interactive map integrating modalities that also retains the input's dimensionality.

In multiple fusion approaches, tensor-operation is mostly used in combination with other fusion operations (13 works out of 19). For example, in [101] the authors introduce a multiple fusion architecture that incorporates both modality-specific networks and modality-shared networks. Combining the knowledge-sharing and the average operation between the modality-specific representations, the model yields a fused representation that exploits the relationships between the features across all the patients and, through the average, maintains the information from the 2 inputs in a final space with the same inputs' dimensionality. However, the effectiveness of approaches employing tensor-operation fusion strategies may be constrained by the limitations inherent to the lack of optimized and dynamic mechanisms that can improve how the modalities interact with each other. Indeed, these fusion operations are not able to dynamically obtain a rich multimodal representation, thus they rely on different parts of the architecture for the elaboration of the multimodal representation.

Attention. The Attention category contains all the methods that use an inter-modality attention mechanism, either in its classic form or a customized version [123]. The characteristic of these approaches lies in the generation of fused features that emerge from the interactions between different modalities, guided by an attention mechanism. By prioritizing features based on their significance in inter-modality relationships, the attention mechanism ensures that the fusion process promotes a multimodal synthesis where the contribution of each modality is weighted according to its relevance and complementarity to the others. This mechanism directs the network's learning focusing on the most pertinent features across the modalities, with a particular emphasis on information reciprocity. Among the reviewed articles, only [112] proposes a single fusion method based on a variation of the attention mechanism. Specifically, the authors propose a multimodal fusion module based on a modality-based attention mechanism. This approach allows for the weighted combination of the input unimodal representations, resulting in a latent representation that captures the combined insights from all modalities while adjusting their relative importance in a context-specific manner. On the other end, the majority of the works that adopt the *Attention* operation to fuse the unimodal representation adopt a multiple fusion approach. This trend is related to the nature of the attention mechanism, which maximizes its effectiveness if integrated multiple times. As an example, in [46] the authors proposed a multiple fusion network designed with cross-attention mechanisms for the inter-modality encoders, able to exploit the inter-modality dependencies among couples of unimodal representations. Another example is [95], where the authors proposed a multiple fusion network, which extracts anatomical CT features, and then intermediately fuses them with PET features through a cross-attention mechanism at multiple depths of the models.

Calibration. The calibration category includes those approaches that merge the unimodal representations into a shared representation of reduced dimensions (Squeeze), generating an intermediate multimodal representation [124]. Then, this representation is used as a calibration signal (Excitation) on all the fusion inputs that generate the signal itself. The approaches that employ calibration fusion are classified as

multiple fusion approaches, given the necessity to employ other fusions to merge the calibration signal with the original input data. Moreover, the generation of the intermediate fused representation is performed generally on features processed with the same number of layers, as evidenced by the majority of synchronous fusion approaches among those that adopt the calibration fusion function [62,63,93,94]. As an example, in [93] a multiple intermediate fusion learning strategy is applied for glioma-grading studies. The authors introduced a module to obtain multimodal information interactions, where the output feature maps from the two unimodal representations are first merged, and then a high-order attention mechanism is applied. Subsequently, the calibration signal obtained is fused with each unimodal input representation. On the contrary, in two other articles [51,68], the calibration is performed in an asynchronous manner, generating intermediate fused representation on features processed with different numbers of layers. As an example in [51] the authors adopt the fusion of two modalities through a calibration process, and at the final layers of the architecture, tabular data are concatenated, resulting in a gradual and asynchronous approach to fuse the modalities.

Knowledge-sharing. This knowledge-sharing category includes all methodologies that learn to extract a useful representation independently of the input representation [123]. In particular, these approaches do not implement a direct interaction between the various unimodal components of the model, instead, they learn the fused multimodal representation through an optimization step, whether it uses a specific loss, a normalization rule (e.g., orthogonalization) or, more simply, back-propagating on a portion of the network that processes different modalities. Creating an intermediate embedding space involves mapping features from different modalities into a common multidimensional space where they can be compared or combined. This approach often involves training with objective functions that encourage similarities or dissimilarities between the content of different modalities, guaranteeing semantic consistency among the feature embeddings. An example is the model proposed in [101], where 2 modalities representations influence each other using a similarity loss, that forces both representations to maintain a fused information consistency between the features spaces. Conversely, in [97] the authors proposed a self-supervised module using an orthogonalization loss to guide the network to learn modality-specific features and maximize the differentiation capability for the final task. A further example is the model reported in [105], where the authors used a dual-stream U-Net structure to implement an intermediate fusion between TeUS and MRI modalities. Both streams share the encoding layer weights in the middle of the network, forcing the backpropagation step to optimize the same bottleneck weights of the network for both unimodal-specific decoder updates.

Fig. 17 presents the distributions of the different categories of fusion operations, distinguished between single and multiple approaches. The left panel in the figure displays the percentages of use of specific types of fusion in relation to the total number of single approaches.

Among the articles employing a single fusion operation, concatenation was the most preferred method, accounting for 82% of cases. This dominance can be attributed to its simplicity and intuitive nature, allowing for straightforward integration of features from different modalities. Other fusion methods were used less frequently: knowledge-sharing (6%), tensor operations (9%), calibration (0%), and attention (3%). This significant imbalance likely stems from concatenation's simplicity and effectiveness in many scenarios, suggesting a data-oriented approach that prioritizes optimizing individual modality representations over complex inter-modal interactions. In fact, concatenation allows the subsequent layers of the network to learn the most relevant cross-modal features, which can be advantageous when the optimal fusion strategy is not known a priori.

In contrast, articles using multiple fusion methods still favored concatenation (46%), but saw a notable increase in tensor operations, rising from 9% in single fusion methods to 32%. This shift reflects the

need for multiple fusion methods for diverse fusion operations throughout the whole fusion module, which often involve tensor or vector manipulations, that lead to a higher level of interactions between the modalities. Regarding other types of fusion, those based on attention mechanisms are increasingly attracting researchers' interest, accounting for 3% and 7% of the fusions in single and multiple architectures, respectively. The attention is preferred in multiple fusion approaches given its effectiveness when applied at different stages of the fusion process, and possesses the ability to achieve a rich representation conditioned by multiple modalities, maintaining a certain degree of explainability thanks to attention maps.

The calibration category, which involves using the calibration signal generated in other subsequent fusions, represents an exclusive approach of multiple architectures. Used in 13% of cases, it is the third most used method in multiple architectures, as it allows a degree of interactivity between unimodal flows within the model that can be managed according to needs.

Indeed, the same module can be used at a single point to generate a single calibration signal for the modalities, or it can be employed multiple times in succession to gradually learn multimodal representations with increasing degrees of abstraction.

Finally, the knowledge-sharing category, which accounts for 6% and 2% of fusions in single and multiple approaches, respectively, emerges as an interesting yet still under-explored approach. The higher percentage of knowledge-sharing operations in single fusion methods compared to multiple fusion approaches can be attributed to the unique characteristics of these methodologies. The preference for knowledge-sharing in single fusion approaches stems from the desire to develop intermediate multimodal methods that centralize this type of fusion, allowing for a more precise evaluation of the effectiveness of this type of fusion. Conversely, in multiple fusion methods, knowledge-sharing operations often play a supplementary role, enriching intermediate representations rather than serving as the primary fusion mechanism, contributing to a richer and multi-faceted approach to multimodal integration.

These patterns underscore the evolving nature of multimodal fusion strategies, where researchers are increasingly exploring sophisticated combinations of techniques to optimize the integration of diverse modalities. As the field progresses, we may see further refinement in the application of the fusion methods presented, potentially leading to more sophisticated and effective fusion architectures. Based on the literature, it appears that most studies have focused on simpler fusion methodologies, *i.e.*, Concatenation and Tensor operations, rather than more complex techniques, *i.e.*, Attention, Knowledge Sharing, and Calibration. While these advanced methods generally offer higher performance, they are also more data-intensive, which poses a challenge in the biomedical field where large datasets are often difficult to obtain.

3.4.7. Unimodal module influence on the fusion module

One of the main aspects that emerges from the analysis is the importance of the dimensionality of the fused space, a crucial factor for the effective integration of unimodal information, because it directly impacts the model's ability to capture and leverage multimodal interactions. Considering the proportions of the unimodal representations in the fused feature space, only 18 out of the 54 articles reported unimodal feature sizes for all the modalities involved in the fusion. Among these 18 articles, 13 of them employ an equal dimension for each unimodal representation, whereas the rest do not adopt this approach and have different dimensions for each unimodal feature vector, as reported in Section 3.3. The choice of homogeneity of the unimodal modules is not a reason for the equal distribution, given that, among these 13 articles, 7 employ homogeneous unimodal modules, whereas the remaining 6 have heterogeneous ones. This almost exact separation between heterogeneous and homogeneous modules shows the trend from the authors to seek a uniform feature space no matter the type of unimodal module involved in processing the unimodal data. As




Fig. 17. Percentage of fusion operation for all the categories, divided by single (a) and multiple (b) fusion approaches.

described in [87], who opts for using equal marginal representation sizes for each modality, this preference could be to avoid bias toward any modality. Although it is known that it can help prevent bias, it remains unclear whether using an equal number of features for each modality is a good choice, given the different nature of the modalities involved.

Furthermore, focusing on the 13 articles that prefer an equal distribution of the unimodal representations, we observe that 12 of them utilize a single sudden approach, whereas 1 employs a multiple sudden approach [121]. This marked prevalence of single sudden approaches using unimodal representations with equal sizes easily correlates with the goal of uniformly analyzing all the modalities involved, giving them equal importance in the final task. Moreover, 19 out of 24 homogeneous approaches employ a synchronous approach. The choice to fuse unimodal representations with homogeneous and synchronous approaches, further confirms a common interest in exploiting uniform and balanced representation spaces. Conversely, the majority of works that employ a heterogeneous approach, 28 out of 30 employ an asynchronous approach, depicting various scenarios, where the fusion occurs between different architectures, data types and fusion operations.

In the next Section, we will delve into the last module of an intermediate fusion approach, called multimodal module and represented as f in Fig. 3. This module, which is provided with the combined features obtained by the fusion module, processes it to extract insights that are only apparent when considering all modalities together.

3.5. Multimodal module

The next component of the intermediate fusion framework, after the fusion operations that generated a multimodal representation, is the multimodal module, denoted with the f block in Fig. 3. This module leverages the representation h derived from the fusion module \mathcal{F} , to obtain the desired target y . The role of the multimodal module strongly depends on the type of architecture employed and the task for which it is used. Following our definition, no fusion operation takes place in this module, but only the final processing of the multimodal representation h to perform the task at hand. In scenarios where the fusion module employs multiple fusion operations, enabling the interaction of various intermediate representations to generate the final fused representation, the multimodal module's role is limited to processing a shallow representation that aligns with the task's representational requirements. Conversely, when the fusion architecture is simpler, such as in single

sudden approaches, the multimodal module assumes a more significant role in processing and refining the multimodal representation.

In analyzing the multimodal module, a crucial aspect is the proportion of trainable parameters with respect to those used in the unimodal module's architectures, to characterize its impact on feature learning and task optimization. In this context, we distinguish 3 cases (Fig. 18) with respect to the unimodal module, which are the cases where unimodal modules have a larger ($M < U$) number of parameters, comparable ($M = U$), or smaller ($M > U$) than the multimodal module. This analysis reflects the role of the module, whether it serves only for final task purposes or acts as a foundational component to process the multimodal features for the final task.

We also characterize the multimodal module on the base of the architecture that it implements, which influences the ability of the model to analyze the multimodal representation h . Firstly, we categorized the multimodal modules' architecture into the following groups:

- **FCNN:** Predominantly employed in both single (29 out of 35) and multiple (17 out of 19) fusion methods, the fully connected architecture plays a crucial role in mapping the feature space to the target space on which the final loss function is computed. This category of multimodal module typically comprises single or multiple fully connected layers. When stacked and combined with non-linear activation functions, FCNNs possess the theoretical capability to approximate any continuous function, rendering them versatile for mapping complex feature spaces to task-specific output spaces. Moreover, FCNNs excel at integrating information from the entirety of the fused feature space simultaneously, making them particularly well-suited for tasks requiring global feature integration.
- **RNN:** This architecture comprises methods that use a multimodal module based on recurrent layers for the prediction task, incorporating single or multiple LSTM or GRU layers to capture dependencies between modalities [55,75,100,104]. RNNs are characterized by their cyclic connections, allowing them to maintain an internal state or memory that can capture temporal dependencies in the input data. This architecture can process inputs sequentially, making them particularly adept at handling variable-length input sequences and for temporal dependency modeling.
- **CNN:** This approach, which features a multimodal module composed of convolutional layers, exploits the ability of the convolutional layers to extract representations from high dimensional




Fig. 18. Three cases illustrating size relationships between unimodal and multimodal modules: $M > U$ indicates the multimodal module is larger than the unimodal module; $M = U$ denotes equal sizes for both modules; $M < U$ signifies the multimodal module is smaller than the unimodal module.

data with a lower computational cost compared to the FCNN approach. Only 3 works use this type of architecture, 2 employing convolutional layers and a final fully connected layer [61,95], whereas the latter uses only convolutional layers [105]. CNNs offer unique strengths in processing structured grid-like data, particularly in extracting hierarchical spatial features. Originally designed for image processing, CNNs have shown remarkable abilities in extracting relevant spatial features from input data, making them particularly effective for modalities with inherent spatial structure. In the multimodal module, CNNs could be considered an alternative way to extrapolate positional-like patterns within the modalities that can match with the final task output.

- **Attention:** It includes methods employing attention-based models. This approach, exploiting the potential of the attention mechanism, can learn how to leverage information from the interaction of the different modalities. Among these works, 2 approaches adopt channel or spatial attention blocks inside their multimodal module [56,94], whereas 3 different works adopt state-of-the-art or customized TF architectures [87,89,117]. Attention allows the model to dynamically weigh the importance of different parts of the input or intermediate representations, enabling more efficient and effective information processing. This approach can be implemented in various forms, including self-attention, cross-attention, and multi-head attention.
- **None:** This is a peculiar category where no trainable layers are employed for the multimodal module. Among the reviewed works, we found a single study that employs this strategy [69]. Here, the authors propose a new intermediate fusion method that optimizes a loss function, considering model uncertainty while estimating correlations among predictions produced by different unimodal modules.

Fig. 19 illustrates the distribution of all these architectures categorized by $M < U$, $M = U$, $M > U$ approaches. We notice that most of the reviewed studies employ unimodal models with larger dimensions than those of the multimodal component (45 over 54). In contrast, only a small portion of the approaches adopt a comparable dimension (7 over 54) or a multimodal module larger than the unimodal modules' dimension (2 over 54).

The vast majority of $M < U$ approaches adopt a FCNN (37 articles) architecture for the multimodal module, whereas 4 methods adopt Attention mechanism to improve the ability to elaborate the fused features representation relatively to the final objective tasks. In these works, the responsibility for learning and elaborating the features representation spaces is mainly shifted to the fusion module. Indeed, in this context, the multimodal module concentrates on fine-tuning the learned representations for the specific objectives of the final task, acting as the last step in a series of transformations to process the fused features for the downstream task.

Conversely, a small portion of the works adopts an unimodal module of equal or smaller dimensions with respect to the multimodal module (red and green slices of the inner circle in Fig. 19). In these

few cases, the multimodal module's role is to transform the initially fused data into a more sophisticated representation useful for the final task, which is capable of exploiting the underlying similarities between the combined modalities in the most effective way. The general trend among all the works underlines the central role of the fusion module rather than the multimodal module. Indeed, this component serves as an integration point but also as a dynamic learning environment where the interplay of different modalities can be explored and optimized. In contrast, most of the approaches use the multimodal module mainly to make the architecture compatible with the downstream tasks.

In the next Section, we will further discuss the downstream biomedical tasks, deepening some other meaningful characteristics of the methods, comprehending transfer learning, explainability, missing modality robustness and the experimental configurations.

3.6. Target

In MDL models, the target represents the output y , as outlined in Fig. 3, which is functional for the task's execution. In the context of biomedical applications, this could be a diagnostic classification, a prognostic regression, or any other relevant task that requires the definition of an outcome. All selected articles perform supervised tasks, involving the use of annotated datasets in which all data instances are labeled. We categorized each article based on the generic task and the medical task. For the first, we identified 4 categories: classification, regression, segmentation, and object detection. The second instead divides into 5 categories: diagnosis, prognosis, detection, drug discovery, and question answering. Fig. 20 shows that the majority perform classification tasks (47 out of 54) with a prevalence of diagnostic tasks (26 out of 54) and prognostics tasks (12 out of 54). Some perform classification tasks for detection purposes (7 out of 54), as in [62] where the authors design a model for stress detection or in [100] where they develop a system for detecting protective behaviors in patients with chronic pain. Only one article addresses classification for drug discovery [84], while another focuses on classification for question answering [86]. In total, 3 articles perform regression, addressing prognostic tasks [70,111,116]. All proposed applications are single-task, except for 2 works: in [105], the authors propose an approach for prostate cancer detection by integrating segmentation and prediction of the cancer likelihood for each pixel; in [94], the authors develop a system for liver tumor detection including classification and bounding box regression for the detection stage.

Focusing on the disease areas involved in intermediate fusion studies, we identified three main clusters, as shown in Fig. 21(a): oncology (22 out of 54), mental health (14 out of 54), and pneumology (9 out of 54). Due to the diverse range of application areas, the remaining publications were grouped into a single cluster labeled as "Others" reporting the specific application for each article. Moreover, our analysis highlights how the nature of the disease influences the choice and use of different data modalities. As shown in Fig. 21(b), oncology relies mostly on genomics data, followed by clinical data, CT, and MRI imaging. Conversely, mental health prioritizes MRI and electrical data,




Fig. 19. Multimodal modules architectures distributions separated by unimodal modules larger (M<U), comparable (M=U), or smaller (M>U) than the multimodal module. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 20. Tasks involved in intermediate fusion studies. Each bar of the bar plot refers to the generic task type (classification, regression, segmentation, object detection) whilst the color legend refers to the medical task involved (diagnosis, prognosis, detection, drug discovery, or question answering). Each task addressed within the same article is considered an individual work. As a consequence, the total count depicted in the bar plot may surpass the number of selected articles. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

e.g., EEG, ECG, whilst pneumology relies mainly on clinical data and X-ray imaging.

3.7. Learning

As we described the key components of an MDL model in the previous sections, we now turn our attention to pivotal techniques for training and deploying MDL models in real-world scenarios.

3.7.1. Transfer learning

Although deep learning has achieved remarkable success, the hunger for data hinders its full potential in the medical field where data scarcity is fueled by high costs, privacy, and ethical issues [125]. To this end, transfer learning has been proposed to mitigate the problem. Inspired by human learning, transfer learning allows the exploitation of knowledge acquired in a similar or different domain, reducing the need to train the entire MDL model from scratch and thus limiting the need for large amounts of data in the domain of interest. Transfer learning consists of two macro-steps: first, the model is trained on large datasets,

a task usually referred to as pre-training [126], to make the model learn general unimodal and/or multimodal feature representations that could be valuable across different related tasks. Second, the pre-trained model is tailored to a specific task, usually leveraging smaller datasets. We refer to this phase as finetuning [127]. Following the tree diagram in Fig. 22, the majority (46 out of 54) train their models from scratch, whilst the remaining (11 out of 54) utilize pre-training, dividing into those pre-training only the unimodal modules (9 out of 54) and those pre-training the entire model (2). With respect to the task at hand, there are different approaches for pre-training a MDL model each with its variations. Some opt to pre-train the unimodal modules in a different supervised task and a different domain (4 out of 54). Following this step, they proceed with either finetuning all modules (3 out of 54) or freezing the unimodal modules while finetuning the multimodal module (1 out of 54). Conversely, some others choose to pre-train the unimodal modules in the same supervised task and related domain (5 out of 54). Subsequently, they either finetune all modules (3 out of 54) or freeze the unimodal modules and finetune the multimodal module in isolation (2 out of 54). Among those pre-training the entire model, we identify those pretraining in the same supervised task and related domain (1 out of 54), and those pretraining in a different unsupervised task and related domain (1 out of 54), followed by finetuning the entire MDL model. Though the lack of large datasets in the medical domain might make transfer learning the optimal solution to address data scarcity, our analysis shows that transfer learning is not widely adopted. Indeed, although it has been successfully employed in several tasks, transfer learning has raised some skepticism especially when employed in the medical field [125]. In recent work, Matsoukas et al. [128] highlight that the core principle of transfer learning is the increasing reuse of learned representations. Therefore, optimal outcomes are achieved when pre-training occurs in domains similar to the specific domain under consideration. This discourages the use of pre-trained models on generic domains like the well-known ImageNet [129]. Although this approach might contribute to speeding up model convergence, it does not guarantee superior performance or more meaningful learned representations. Therefore, the lack of pre-trained backbone architectures on medical domains makes transfer learning a sub-optimal approach to address data scarcity, which paves the way for alternative solutions such as generative data

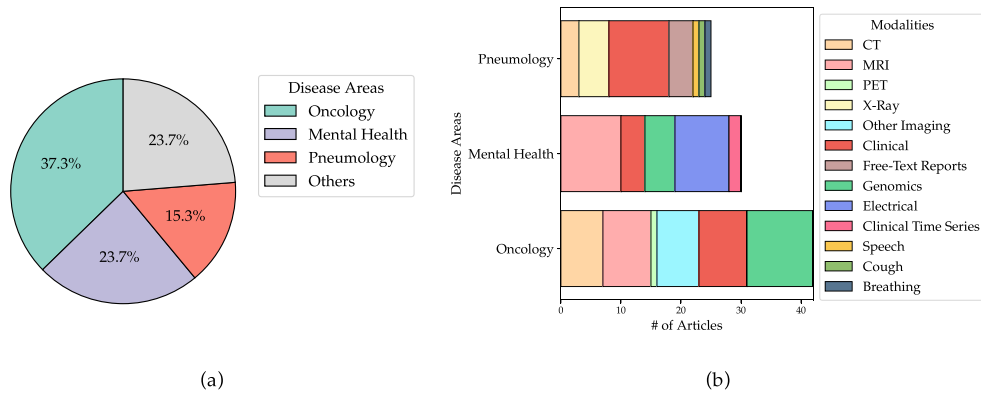


Fig. 21. (a) Disease areas involved in intermediate fusion studies. Others: visual question answering [86], visual impairment [42], overall survival form discharge [75], kidney injury [112], chronic pain [100], placental insufficiency [120], osteoarthritis [110], voice pathology [104], heart failure [91], glaucoma [121], diabetic rethynopathy [121], heart failure mortality prediction [76], phenotype classification [55], all-cause mortality prediction [55]; (b) Exploring data modalities across top three disease areas.

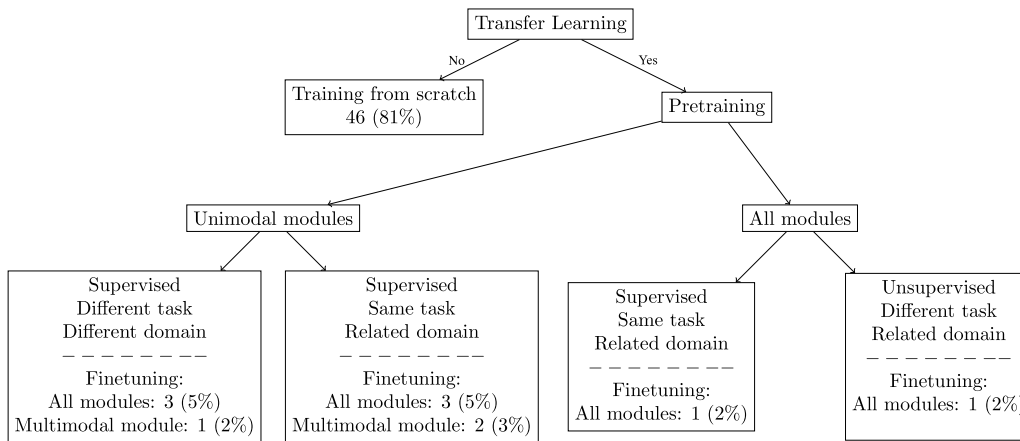


Fig. 22. Transfer Learning in intermediate fusion methods. The tree diagram distinguishes between articles training from scratch the entire MDL model and articles pretraining the unimodal modules in isolation or pretraining the entire network.

augmentation [2]. While our work does not cover generative data augmentation, interested readers can refer to [130] for further information on this topic. Focusing on the articles that perform transfer learning, all utilize one or more pre-trained CNNs [53–56,58,73,110,114,116,117]. Two articles use a pre-trained RNN [53,55], one uses a pre-trained FCNN [58], and another employs a TF [56]. We observed that among those pre-training their networks, 7 utilize well-known CNN architectures, while 3 opted to pre-train their custom models [58,110,117]. The well-known CNN architectures include ResNet-based models [53–55,114,116], InceptionV3 [73] and EfficientNet [56]. While the pre-trained RNNs and FCNN are custom-designed, for the TF, a BERT-based model is used.

3.7.2. Multimodal optimization

In a MDL model, the optimization strategy can significantly influence the multimodal representation, h , returned by the fusion module \mathcal{F} , either through indirect or direct effects. Indirect effects occur when the loss function operates on the output y , thus indirectly influencing the update of the multimodal representation h . Conversely, direct effects involve loss functions that act directly on the feature representations, promoting their interaction and complementarity. Based on our analysis, 51 articles fall into indirect effects, with the majority of them using standard loss functions (e.g., cross-entropy, mean-squared

error, focal loss). Among the indirect effects, it is notable that some approaches consider the multimodal nature of the model when formulating the loss function. This is done with the explicit intent of guiding and balancing the influence of each modality in defining the multimodal representation, h . In [120], the authors use a contrastive learning strategy, using the triplet loss to decouple modality shared and modality-specific feature spaces. To balance the informativeness of each modality before fusion, [118] quantifies the multimodal confidence by computing the true class probability which is then used to weight each unimodal representation before feeding to \mathcal{F} . Similarly, [69] models the uncertainty of individual modalities produced by different learners, thus capturing the pairwise correlation between predictions. To ensure all modalities can equally contribute to the model's output, preventing the model from being biased towards a subset of modalities, [48,49] use a multimodal regularization [131] based on functional entropy whilst [116] uses different learning rates for the unimodal modules to balance the learning of different modalities. Turning our attention to direct effects, we classify them into intra-modal and inter-modal effects. The former involves architectures that use multiple unimodal modules for each modality, aiming to extract distinct representations for a single modality. For instance, in [101], the authors impose an orthogonalization constraint within the individual unimodal embeddings to help the separation between modal




Fig. 23. Robustness to missing modalities. The pie chart highlights the articles whose architectures do not handle missing modalities, those that are designed to be robust against missing modalities, and those that although not specified by the authors could potentially handle them.

shared and modal specific features. As for the inter-modal effects, we find methodologies that fall into the category of knowledge-sharing fusion [93,97,101] as we describe in Section 3.4.

3.7.3. Missing modalities

The term missing modalities denotes the partial or complete absence of one or more modalities within a multimodal instance. In the context of multimodal learning, the presence of missing modalities is a major concern that can undermine the applicability of a model, especially in data-in-the-wild scenarios where incomplete instances are prevalent. In the medical field, this problem is even more pronounced as the combination of missing modalities and data scarcity severely limits the availability of large amounts of data. Looking at Fig. 23, the majority of the articles (45 out of 54) showcase models that lack robustness to missing modalities. The main criticality lies in the way the modalities are fused with a prevalence of simple concatenation-based fusion that requires complete data instances. Only a few (5 out of 54) claim their models can handle missing modalities. In their work, Kohankhaki et al. [56] propose an attention-based mechanism for the fusion layer followed by a pooling operation, which is permutation invariant to its inputs, making the absence of some modalities unproblematic. Alternatively, Hayat et al. [55] treat the multimodal representation as a sequence of unimodal tokens, such that the fusion module performs the aggregation through a recurrence mechanism, e.g., LSTM. A further approach is introduced by Zheng et al. [112], where an exponential decreasing or increasing mechanism is used to gradually plug out or plug in a specific data modality. To simulate a missing modality scenario, Ostertag et al. [59] propose a training strategy that randomly imputes values from the training set with zeros. They indeed demonstrate that training the multimodal architecture with randomly missing data increases its performance. Leveraging a GNN, Jiao et al. [120] first build a graph-based manifold knowledge on complete data samples which is then transferred to incomplete data scenarios, so that when a modality is missing, its knowledge is derived from the constructed manifold information. The remaining 4 articles, while not explicitly addressing the problem, present approaches that potentially could handle missing modalities. This is achieved either through fusion using attention-based techniques [87,89,117] or through element-wise tensor operations [93].

3.7.4. Explainability

To date, we are witnessing a surge in the performance of DL models due to larger databases and increased computational power [132]. Just as task complexity is increasing, model architectures are also becoming more complex, featuring nested nonlinear structures that turn them into black boxes hiding the decision-making process. Lack of interpretability

and transparency hinders the use of DL-based systems in many real-world applications, such as in the medical domain. Indeed, in medical applications, the possibility of establishing a connection between explainability and causality is crucial as it allows human experts to understand the decision-making process of an AI system [114] but also acts as a prerequisite for new insights [132]. Indeed, the possibility to unveil new patterns in data as well as novel associations, e.g., between genes and diseases, through model understanding and introspection is often more crucial than the prediction itself.

To date, explainable AI (XAI) offers several tools to improve the trust and transparency of AI models, but it is still in its infancy, especially in the context of MDL. This is confirmed by our analysis with the majority of articles (34 out of 54) not addressing XAI in their works or only mentioning it as a future development.

The interested reader will find several taxonomies of XAI techniques in the literature, each of which groups them according to a certain criterion [133]. For example, an explanation can be model agnostic or model specific based on its dependence on the model it is trying to explain, or it can be local or global depending on whether the scope is to derive predictions for a specific data point or the entire model. Moreover, explainability can be introduced at different stages of model development, producing pre-modeling, during modeling, and post-hoc modeling explanations.

The specific focus of our review on MDL, leads us to a further categorization of XAI techniques based on the module of the MDL model on which the explanations are produced. This classification results in a distinction between unimodal and multimodal explanations. The former aim to highlight salient features within a single data modality, whilst the latter focus on modality importance, quantifying the heterogeneity of information across modalities as well as their interactions. In Table 5, we list unimodal and multimodal XAI approaches used in the selected articles with their respective reference. Due to their extensive adoption and established success in various domains, CNNs [42,46,51,54,56,62,63,87,91,95,97,105,114,116,118] are the most explained architecture type, followed by FCNNs [46,48,49,56,91,114,116]. Only three publications provide explanations for RNNs [48,49,112], two for TFs [56,87], one for GNNs [120], and another for AEs [48].

3.7.5. Experimental configuration

Designing a comprehensive experimental setup is an essential requirement for objectively assessing MDL models. Based on our analysis, we deem that a thorough experimental setting should include: (A) conducting all the experiments using a robust resampling technique (e.g., cross-validation, bootstrap sampling, jackknife sampling, leave-one-out cross-validation, stratified sampling, random sampling); (B) comparing unimodal and multimodal configurations; (C) comparing various fusion strategies, i.e., early, intermediate, and late; (D) conducting structured statistical tests to determine the statistical significance of model performance; (E) provide for the use of external validation, thus assessing the model's generalization.

Table 6 shows the experimental setting adopted by the selected articles, bringing out that the literature is still lacking in using comprehensive experimental settings. First of all, the use of robust resampling technique (A), which is pivotal for ensuring the reliability and robustness of experimental findings, is limited to just 19 articles. When comparing different experiments, it is typical to compute numerical metrics (e.g., accuracy, F1 score, dice score, etc.) relevant to the task at hand and compare their averages. However, merely comparing average scores is not enough to determine the superiority of one approach over another. A higher average does not necessarily signify a better approach since a wider dispersion of individual results might indicate greater instability. This is why, the average performance of an experiment should always be reported along with their standard deviations.

Table 5
XAI approaches for unimodal and multimodal explanations used in the selected articles.

Unimodal XAI	Multimodal XAI	Articles
Saliency Maps	Modality-wise Attention, Correlation of Interaction	[48,49]
GradCAM, Integrated Gradient	–	[114]
–	Modality Ablation	[75]
Ablation, LIME	–	[91]
SHAP, Saliency Maps	–	[116]
GradCAM	–	[42,54]
3D t-SNE	3D t-SNE	[63]
CAM	–	[51]
Guided-backprojection	–	[95]
LayerCAM	–	[97]
Attention Maps	–	[46]
CAM, UMAP	CAM	[120]
Feature Importance	Modality Importance	[118]
GradCAM++, Integrated Gradients	–	[56]
GradCAM, Feature Importance, t-SNE	t-SNE	[87]
Masking	–	[105]
t-SNE	–	[62]
Feature Importance	Modality-Wise Attention	[112]




Fig. 24. (a) Average performance of multimodal approaches compared to unimodal approaches. \bar{U} denotes average unimodal performance, \bar{M} denotes average multimodal performance, σ denotes the standard deviation; (b) Average performance of intermediate fusion approaches compared to other fusion techniques. \bar{EF} denotes average early fusion performance, \bar{IF} denotes average intermediate fusion performance, \bar{LF} denotes average late fusion performance, σ denotes the standard deviation.

Although our analysis has focused on works addressing fusion techniques, the question “Is multimodal better than unimodal?” is still central in the research community and is preliminary to the investigation of how to represent and combine different data modalities. Indeed, 42 articles include **B** in their setting. In Fig. 24(a), we show that 40 works demonstrate the superiority of the multimodal approach over the unimodal approach. However, 29 out of these 40 report only the average performance without standard deviation. Among those reporting the standard deviation, 5 confirm the superiority of the multimodal approach by showing lower standard deviations than the unimodal approach, 5 report equal standard deviations, and 1 shows higher standard deviations for the multimodal approach. Going back to Table 6, **C** is second in position with 31 articles performing a comparison between fusion strategies. In Fig. 24(b), we show that intermediate fusion outperforms early and late fusion in 24 works. Again, the majority do not report standard deviations (16 out of 24). Among those reporting the standard deviation, 4 corroborate the superiority of intermediate fusion by showing lower standard deviations whilst the other 4 report equal standard deviations for all fusion strategies, *i.e.*, early, intermediate, and late. Only 1 article shows that early fusion outperforms intermediate and late fusion with no difference in standard deviation, whilst 3 articles claim late fusion as the best fusion technique of which 2 do not report standard deviations and 1 shows equal standard deviations between early, intermediate, and late fusion. Only 14 articles assess the statistical significance of their results through structured statistical tests (**D**). Furthermore, only 5 use an external validation set in their experiments (**E**). The use of a complete and comprehensive experimental procedure including **A**, **B**, **C**, **D** and **E** is limited to only 4 articles. In addition, only 11 articles make their code public thus not facilitating transparency and reproducibility.

4. Discussion and conclusion

As we draw this systematic review to a close, it is essential to step back and contemplate the current limitations, the trajectory for future research as well as the practical impact of our findings. This section synthesizes the insights garnered from our in-depth analysis of intermediate fusion methods in MDL for biomedical applications. As mentioned above, this survey includes two supplementary documents for detailed information on individual papers, enhancing readability: Supplementary Material A contains a table summarizing all fields discussed in Section 3 for each paper, Supplementary Material B analyzes the fusion mechanisms of each paper, showing the corresponding fusion notation and graphical representation.

4.1. Limitations, research gaps and future directions

4.1.1. Limitations

While conducting the analyses, we identified significant gaps in the reporting of essential information, leading to reduced methodological transparency in the field of MDL, as shown in Fig. 25. These gaps are not trivial; they are central to ensuring reproducibility, understanding model behavior, and validating the robustness of the proposed solutions. Regarding the modalities’ information, the input size (Section 3.2) and the marginal representation size (Section 3.3.4) are insufficiently reported. Without clarity on input dimensions and the size of representations post unimodal processing, it becomes challenging to grasp the data’s distribution and the extent of feature extraction prior to fusion. Regarding XAI (Section 3.7.4), the lack of reporting on unimodal XAI and even more so on multimodal XAI limits the understanding of the multimodal models. In particular, explainability is fundamental in biomedical applications where understanding model decisions can

Table 6

Experimental configurations used in intermediate fusion studies. A: cross-validation scheme, B: comparison between unimodal and multimodal configurations, C: comparison between different fusion strategies, *i.e.*, early, intermediate, and late, D: statistical tests, E: external validation set, \emptyset : none of the previous but only comparisons of intermediate fusion configurations. The symbol \wedge denotes the logic operator AND. The central column shows the number of articles using a specific experimental setting whilst the right column reports the references.

Experimental Configuration	# of Articles	References
\emptyset	2	[62,63]
A	2	[58,113]
B	13	[42,54,61,62,68,70,73,75,92,96,104,105,111,119]
C	4	[56,72,76,103]
D	1	[86]
A \wedge B	1	[45]
B \wedge C	6	[50,53,87,90,94,112]
B \wedge D	3	[59,91,101,120]
C \wedge D	1	[118]
A \wedge B \wedge C	7	[48,49,51,55,69,110,115]
A \wedge B \wedge D	1	[97]
A \wedge C \wedge D	2	[84,99]
B \wedge C \wedge D	4	[89,98,120,121]
A \wedge B \wedge C \wedge E	1	[95]
A \wedge B \wedge C \wedge D	2	[46,100]
A \wedge B \wedge C \wedge D \wedge E	4	[93,114,116,117]

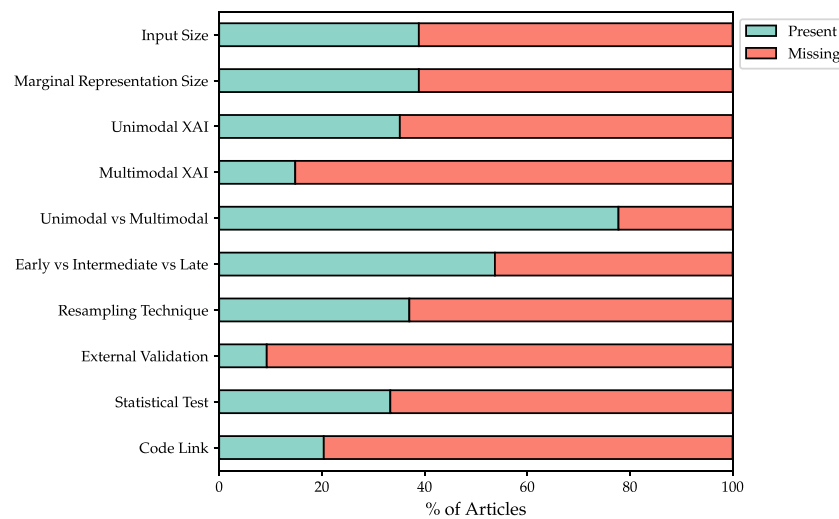


Fig. 25. Distribution of Reported Information in Analyzed Articles. Presents stacked bar charts for various fields, indicating the percentage of information explicitly stated in the articles (green) versus the information missing or not stated (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

impact clinical outcomes, essential for increasing trust and adoption of AI-based systems. Regarding the comparisons (Section 3.7.5), a minority of studies fail to compare unimodal and multimodal approaches, which is a relatively lower omission but still significant for establishing the value of multimodal strategies. A more substantial gap is seen in the comparison of different fusion strategies (Section 3.7.5), with studies not reporting comparisons between early, intermediate, and late fusion, which are pivotal for contextualizing the effectiveness of intermediate fusion. In terms of validation (Section 3.7.5), many studies do not use sampling techniques, and even more overlook external validation. Such omissions call into question the generalizability of the models. Furthermore, statistical tests (Section 3.7.5) to establish the significance of findings are omitted in many of cases, raising concerns about the statistical robustness of the reported outcomes. Finally, regarding the absence of a public code link in publications significantly hinders the field's progression. The availability of code is essential for peer validation, understanding model intricacies, and fostering collaborative improvements. The concerning levels of missing information in these critical areas suggest a need for more rigorous reporting standards in MDL research, and also highlights an opportunity for future work to address these gaps, enhancing the quality and transparency of research within the domain.

4.1.2. Potential research gaps

Moving forward, we hereby outline the primary research gaps identified:

- Limited modalities combination:** As discussed in Section 3.2.7, most studies utilizing intermediate fusion techniques in MDL within the medical field tend to focus on integrating only two types of modalities, typically imaging and tabular data [134]. More complex combinations that involve integrating more than two modalities, including text, video, time series, and audio, are relatively uncommon. This limited bimodal approach, centered on imaging and tabular modalities, may be due to the complexity and challenges associated with merging data from diverse modalities, as well as the need for effective data structuring and preprocessing. However, a more comprehensive multimodal analysis could offer a more detailed view of the patient, enhancing the performance of the models.
- Lack of Multi-scale Data:** In the literature on MDL using intermediate fusion techniques within the biomedical field, there is a notable gap in integrating data of different dimensional scales. Combining macroscopic data (*e.g.*, radiological images) with microscopic data (*e.g.*, histopathology) can offer a more comprehensive understanding of the health status of patients.

Such integration would enable the direct correlation of observable macroscopic changes with underlying molecular alterations, thereby improving the understanding of various pathologies.

- **Lack of benchmark datasets:** As shown in Section 3.2, a notable research gap in the field of intermediate MDL for biomedical applications is the absence of standardized, publicly available benchmark datasets. Similar to the role of ImageNet [135] in evaluating novel models for computer vision classification tasks, the availability of such datasets would allow researchers to compare and validate their deep learning models across different biomedical domains.
- **Suboptimal marginal representation dimensions:** A significant research gap exists regarding the dimensionality of marginal representations. Our analysis reveals no clear trend in how these dimensions are chosen, with many studies arbitrarily defaulting to powers of two, likely influenced by binary representations. More critically, as mentioned in Section 3.3.4 the majority of studies employ equal-dimensional marginal representations across modalities, an approach that often lacks justification and may be suboptimal, considering that different modalities inherently contain varying amounts and types of information. The lack of a systematic approach to determining appropriate dimensions for each modality potentially leads to suboptimal fusion performance, loss of critical modality-specific information, and inefficient use of computational resources.
- **Insufficient consideration of multimodal fusion strategies:** As shown in Section 3.4, our analysis reveals that there is a persistent tendency in the field to prioritize unimodal feature extraction approaches over addressing the complexities of multimodal fusion. This bias has led to the development of oversimplified architectures for intermediate fusion, leaving critical aspects unexplored, as shown in Section 3.4.6 where we analyzed the fusion operations performed in the fusion module. Key challenges remain unresolved, such as determining optimal methods for combining network branches that process different modalities, managing varying input sizes and dimensionalities, and ensuring efficient information flow across the fused network. Furthermore, our investigation indicates that the majority of fusion operations are performed without strong theoretical foundations or assumptions regarding the relationships between modality features. This lack of a consistent approach results in methodologies that do not take full advantage of the complementarity and correlation of the characteristics of the different modalities. Consequently, the potential for integrating information from different sources remains largely unrealized, limiting the effectiveness of current multimodal fusion strategies in biomedical applications. Addressing these limitations requires a paradigm shift in research focus, emphasizing the development of more sophisticated, theoretically grounded fusion techniques that can effectively capture and utilize inter-modal relationships.
- **Limited Multimodal Module exploration:** The multimodal module, responsible for the final elaboration of fused information from the fusion module and its mapping to the task representation space, is a critical component in intermediate MDL. Our analysis, as detailed in Section 3.5, reveals a significant research gap in understanding the most effective architectures for extracting valuable task-specific representations from fused features. The vast majority of studies rely heavily on FCNN architectures for this crucial stage. While FCNNs offer different advantages, this predominant focus overlooks the potential benefits of more advanced approaches that could optimize the final learning elaboration from fused feature representations. There is a notable scarcity of research exploring alternative architectures such as RNNs, CNNs, attention mechanisms, and graph-based methods in this context. This limitation is particularly concerning given the diverse nature of biomedical tasks and datasets. Different architectural approaches may offer unique advantages in capturing complex temporal dependencies (RNNs), spatial patterns (CNNs), feature importance (attention mechanisms), or relational structures (graph-based methods) within the fused representations. The lack of comprehensive comparative studies exploring these varied approaches hinders our ability to determine optimal architectures for specific biomedical applications.
- **Limited interpretability and explainability:** Deep learning models in biomedical applications often lack interpretability, making it challenging to understand and explain the decisions made by the models [136]. As shown in Section 3.7.4, our analysis shows that XAI is still an understudied topic in most of the current state-of-the-art of biomedical multimodal intermediate fusion approaches. This oversight highlights significant research gaps: there is a lack of standardized methods for evaluating model interpretability, inadequate methods to enhance the transparency of complex models, and insufficient integration of XAI methods tailored to multimodal data. Bridging these gaps is essential to ensure AI models are both accurate and interpretable, fulfilling the requirements for deployment in safety-critical applications such as healthcare.
- **Transferability and generalization across biomedical domains:** Many existing studies in intermediate MDL focus on specific biomedical domains or tasks, which limits the transferability and generalization of the models developed. There exists a notable research gap exploring strategies that facilitate transfer learning and domain adaptation to enhance the generalization capabilities of these models across various biomedical fields. Moreover, investigating the transferability of learned representations and developing techniques for detecting and mitigating domain shifts are critical areas that remain underexplored.
- **Robustness to missing modalities:** The ability to handle missing modalities is paramount for the deployment of MDL models in real-world scenarios, especially in such a complex and heterogeneous field as healthcare. As shown in Section 3.7.3, our analysis reveals that most of the current biomedical multimodal intermediate fusion approaches are still inadequate to be deployed in real-world medical applications [137]. This is evidenced by the prevalence of simple concatenation-based fusion strategies, which fail to make models robust to missing modalities. Consequently, most existing models lack inherent robustness to missing modalities during both training and inference, significantly hindering their applicability in real-world healthcare settings.
- **Real-time and resource-constrained implementation:** Despite the promise of intermediate MDL in biomedical applications, the real-time implementation and deployment of these models on resource-constrained devices or environments is an ongoing research challenge. Future research could investigate techniques that enable efficient computation, model compression, and optimization to accommodate real-time requirements and resource limitations. Exploring hardware acceleration techniques or designing specialized architectures suitable for intermediate MDL in resource-constrained settings could help bridge this research gap.
- **Ethical considerations and potential biases:** As intermediate MDL models become more pervasive in biomedical applications, it is crucial to address ethical considerations and potential biases that may arise from these models. Research could aim to investigate the potential biases, fairness, and transparency issues associated with the deployment of these models, with a focus on ensuring patient privacy, data anonymization, and reducing biases related to race, gender, or socioeconomic factors. Designing frameworks that enable the identification and mitigation of

biases within intermediate MDL models would contribute towards ensuring the ethical use of these technologies in biomedical applications.

While there is a clear path forward, marked by research gaps and opportunities for enhancement, the potential for these sophisticated computational techniques to aid patient care and biomedical research remains beyond question. It is with a collective effort toward addressing these gaps, through rigorous research, transparent reporting, and collaborative innovation, that we can harness the full potential of MDL. As we stand on the peak of these advancements, the convergence of technology, medicine, and data science prospects a future where healthcare is more precise, personalized, and preventive.

4.1.3. Future research

The exploration of intermediate fusion methods in MDL has uncovered numerous avenues for future research, especially pertinent to advancing the capabilities and applications within the biomedical sphere. These potential directions are critical for progressing the field and addressing the identified limitations that currently challenge the efficacy and applicability of these methods.

- **Multimodal dimension and modalities combination:** Future research should prioritize investigating trimodal and high-modal combinations, as well as incorporating diverse modalities like time series, text, audio, and video. Our analysis has shown that most current studies are bimodal, predominantly utilizing imaging and tabular data. Exploring a wider range of data types and multimodal dimensions could be essential for gaining a more comprehensive understanding of clinical scenarios, thereby enhancing diagnostic and predictive accuracy.
- **Integration of Multi-scale Data:** One area of future research could be focused on integrating multi-scale data from different modalities for better analysis and understanding in biomedical applications. This could involve developing deep learning architectures that can effectively handle and fuse data at different spatial and temporal scales, such as combining microscopic and macroscopic data in cancer research.
- **Creation of benchmark datasets:** Future research should focus on the creation and curation of large-scale, high-quality benchmark datasets that encompass multiple biomedical domains, ensuring a wide range of data modalities, such as imaging, genomics, clinical data, and others. These datasets should be annotated with standardized labels and metadata to support various deep learning tasks, enabling consistent evaluation and comparison of deep learning models.
- **Optimization of marginal representation dimensions:** Future research should investigate the optimal dimensionality of marginal representations in intermediate fusion. Current research shows no clear trend in how these dimensions are chosen which indicates that the selection might be suboptimal. This presents a significant research opportunity to explore the impact of varying dimensionalities on fusion performance; methods for determining optimal dimensions for each modality, taking into account the inherent complexity and information density of different data types; and strategies for balancing the representation of modalities with significantly different feature spaces without losing critical information. Addressing this gap could lead to more efficient and effective multimodal fusion models that better capture the unique characteristics of each modality.
- **Deepening Multimodal Architectures:** Future research in intermediate MDL for biomedical applications should prioritize the development of innovative and integrated fusion strategies. By exploring novel approaches to combine information from diverse sources, taking also into account the data heterogeneity in the biomedical field, more advanced techniques to harmoniously exploit the most informative features could be obtained. Moreover, the actual research lacks of advanced techniques able to map the multimodal features into the final task output, limiting the effective capability of the model to elaborate the multimodal nature of the input data. This paradigm shift in research focus should emphasize creating more sophisticated, theoretically grounded fusion techniques capable of effectively capturing and leveraging inter-modal relationships. Such advancements will be critical in maximizing the impact of multimodal analysis on diagnostic accuracy, treatment planning, and overall patient care in the biomedical field.
- **Explainability and Interpretability:** Deep learning models often lack interpretability, a crucial aspect in biomedical applications where decision-making is critical. Future research should aim to improve the explainability of these models in biomedical contexts. This involves creating methods that provide clear interpretations or explanations of the decisions made by MDL models, enabling clinicians and researchers to understand and trust their predictions. To achieve this, there is plenty of room for advancement. Future research should prioritize integrating XAI principles starting from the initial design phase, favoring inherently explainable architectures. Moreover, developing standardized evaluation metrics for model interpretability, and enhancing user interfaces to better communicate model decisions can further advance this field. To this end, a joint effort between researchers and clinicians will also be crucial to ensure that the developed solutions are both technically robust and practically applicable in real-world biomedical settings.
- **Transfer Learning and Domain Adaptation:** Biomedical datasets are often limited in size and differ in distribution across different hospitals or research centers. Future research should focus on investigating transfer learning and domain adaptation techniques in MDL to overcome the challenge of limited data and domain shift. By exploring these techniques, models trained on one dataset or domain could be effectively applied to other datasets or domains, thereby overcoming data scarcity and improving generalizability. This research direction has the potential to improve the robustness, applicability, and resilience of multimodal deep-learning models in biomedical applications.
- **Privacy and Ethical Considerations:** Future research on MDL for biomedical applications must address privacy, ethics, and regulatory compliance. Beyond the upcoming AI Act, frameworks such as GDPR, HIPAA, ISO/IEC 27001, and the NIST AI RMF provide essential guidelines for data security, transparency, and risk management. Implementing privacy-preserving techniques like federated learning and differential privacy will safeguard data confidentiality while ensuring compliance with informed consent and data governance principles. Addressing bias and fairness is equally crucial, as demographic disparities in training data can lead to skewed model predictions. Techniques such as fairness-aware learning and adversarial debiasing should be explored to promote equitable AI outcomes.
- **Real-time Applications and Clinical Decision Support:** Investigating real-time applications of MDL in clinical decision support systems is a promising area of future research. This could involve developing models that can process and analyze multimodal data streams in real-time, providing timely insights or recommendations to healthcare professionals for accurate diagnosis or treatment planning.
- **Federated Learning:** As healthcare data grows not only in volume but also in sensitivity, maintaining privacy and data security becomes crucial. Federated learning presents a promising solution by enabling the training of models on decentralized datasets without needing to share the data itself. Future research could

explore how federated learning can be applied to intermediate fusion in MDL, particularly in biomedical applications. Investigating this approach would address privacy concerns, particularly with patient data, and could enhance model training by leveraging diverse datasets from multiple institutions without compromising data security.

- **Multimodal Medical Foundation Models:** Advances in MDL have made great strides in improving diagnosis, prognosis, and treatment decisions. However, current methodologies remain largely task-specific (narrow AIs), resulting in poor generalization and lacking broad medical understanding and contextual information. In contrast, the recent rise of foundation AI models offers a chance to rethink biomedical AI systems, enriching them with more general capabilities. Foundation models have gained popularity in language, and their scope has gradually expanded to other domains, such as computer vision. These models are often trained with self-supervision and contrastive learning approaches, fully exploiting the informativeness of large-scale unlabeled data. This enhances scalability but also compels the models to predict parts of the inputs, making them richer and potentially more useful than models trained on a limited labeled space. Foundation models can be effectively adapted to several downstream tasks and settings using in-context learning or few-shot fine-tuning, with the same set of model weights. Despite their promise, foundation models in healthcare are still in an early stage, facing several limitations such as insufficient domain-specific training data, challenges in handling the diverse and complex nature of medical information, and the need for rigorous validation in clinical settings. Given the potential of foundation models, it is imperative to investigate their multimodal capabilities, leading to the opportunity to define a unified medical AI system that can interpret complex multimodal data. The new paradigm offered by foundation models in conjunction with multimodal learning expertise leads to the discovery of inter/cross-modality information but also potential knowledge intersection between seemingly different tasks and data. Moreover, there will be opportunities to explore the pivotal role of text as the “inter-lingua” of all modalities [138]. This will result in a holistic orchestration between data modalities and tasks that equips the AI system with robust biomedical knowledge, to be deployed in real clinical applications.

4.2. Implications for practice

Our review sheds light on specific practical implications that can significantly advance research and application in biomedical fields through multimodal AI. Here are detailed benefits and impacts:

- **Enhanced Diagnostic Accuracy:** By integrating diverse data modalities through intermediate fusion, diagnostic models become more precise. For example, in oncology, combining imaging data with genomic information enables the precise classification of tumor types, leading to earlier and more accurate cancer diagnoses. This comprehensive approach allows clinicians to detect anomalies that might be missed when using a single data modality, ultimately contributing to more effective treatment strategies.
- **Personalized Patient Care:** The deep insights provided by intermediate fusion into complex patient data are vital for personalizing treatment plans. In chronic disease management, integrating electronic health records with real-time data from wearable devices can tailor interventions to individual patient needs. For instance, in diabetes care, combining blood glucose levels with lifestyle data helps in adjusting insulin doses more accurately, enhancing patient outcomes.

- **Efficient Data Utilization:** With the ever-increasing volumes of data in healthcare, intermediate fusion methods stand out by enabling the synthesis of heterogeneous data streams into actionable insights. Hospitals can employ these methods to integrate lab results, medical imaging, and patient histories into unified dashboards for clinicians, reducing information overload and improving decision-making efficiency. This not only streamlines workflows but also ensures critical information is readily accessible during patient consultations.
- **Advancement in Research:** Applying intermediate fusion techniques can lead to significant breakthroughs in biomedical research. For instance, in pharmacogenomics, fusing genetic profiles with drug response data accelerates the discovery of how genetic variations affect individual responses to medications. This can lead to the development of more effective drugs with personalized dosing regimens, minimizing adverse effects and improving therapeutic outcomes.
- **Education and Training:** The integration of intermediate fusion methods into educational curricula for healthcare professionals highlights the shift towards data-centric medical practices. Medical schools and training programs can incorporate courses on AI and data fusion technologies, equipping future healthcare providers with the necessary skills to utilize advanced diagnostic tools. This ensures that upcoming professionals are proficient in interpreting complex multimodal data, which is increasingly essential in modern healthcare settings.
- **Resource Optimization:** Improvements in diagnostic and treatment accuracy reduce unnecessary procedures and optimize the allocation of medical resources. For example, more accurate diagnostics can decrease the need for invasive surgeries or redundant testing, thereby cutting costs and reducing patient risk. Healthcare systems can reallocate these saved resources to areas with greater need, enhancing overall patient care quality.
- **Ethical and Privacy Considerations:** As multimodal fusion becomes more embedded in clinical practice, it is imperative to navigate the ethical and privacy aspects diligently. Implementing robust data governance policies, such as encryption and anonymization techniques, ensures that patient data is securely handled. Compliance with regulations not only builds patient trust but also protects healthcare institutions from legal repercussions, making ethical data management a cornerstone of practice.

By highlighting these aspects, this review emphasizes the critical role of intermediate fusion in transforming the biomedical landscape. The practical applications outlined demonstrate how intermediate fusion can be directly implemented to improve diagnostic processes, personalize patient care, and optimize resources in real-world healthcare settings. The adoption of these advanced computational techniques paves the way for more effective, personalized, and data-driven approaches in healthcare, positioning the field to leverage the full potential of AI in improving patient outcomes.

CRediT authorship contribution statement

Valerio Guarrasi: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Fatih Aksu:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Camillo Maria Caruso:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Francesco Di Feola:** Writing – review & editing,

Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Aurora Rofena:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Filippo Ruffini:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Paolo Soda:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Fatih Aksu, Camillo Maria Caruso, Aurora Rofena and Filippo Ruffini are Ph.D. students enrolled in the National Ph.D. in Artificial Intelligence, course on Health and Life Sciences, organized by Università Campus Bio-Medico di Roma.

This work was partially founded by: (i) Università Campus Bio-Medico di Roma under the program “University Strategic Projects” within the project “AI-powered Digital Twin for next-generation lung cancer care (IDEA)”; (ii) from PRIN 2022 MUR 20228MZFAA-AIDA (CUP C53D23003620008); (iii) from PNRR MUR project PE0000013-FAIR; (iv) from Project ECS 0000024 Rome Technopole, - CUP C83C22000510001, NRP Mission 4 Component 2 Investment 1.5, Funded by the European Union - NextGenerationEU. (v) from Kempestifelserna grant n. JCSMK24-0094

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.imavis.2025.105509>.

Data availability

Data will be made available on request.

References

- [1] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444.
- [2] D. Ramachandram, G.W. Taylor, Deep multimodal learning: A survey on recent advances and trends, *IEEE Signal Process. Mag.* 34 (6) (2017) 96–108.
- [3] D. Lahat, T. Adali, C. Jutten, Multimodal data fusion: an overview of methods, challenges, and prospects, *Proc. IEEE* 103 (9) (2015) 1449–1477.
- [4] J. Sui, T. Adali, Q. Yu, J. Chen, V.D. Calhoun, A review of multivariate methods for multimodal fusion of brain imaging data, *J. Neurosci. Methods* 204 (1) (2012) 68–81.
- [5] S.E. Viswanath, P. Tiwari, G. Lee, A. Madabhushi, A.D.N. Initiative, Dimensionality reduction-based fusion approaches for imaging and non-imaging biomedical data: concepts, workflow, and use-cases, *BMC Med. Imaging* 17 (2017) 1–17.
- [6] T. Ching, D.S. Himmelstein, B.K. Beaulieu-Jones, A.A. Kalinin, B.T. Do, G.P. Way, E. Ferrero, P.-M. Agapow, M. Zietz, M.M. Hoffman, et al., Opportunities and obstacles for deep learning in biology and medicine, *J. R. Soc. Interface* 15 (141) (2018) 20170387.
- [7] M.A. Ahmad, C. Eckert, A. Teredesai, Interpretable machine learning in healthcare, in: *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, 2018, pp. 559–560.
- [8] E. Vayena, A. Blasimme, I.G. Cohen, Machine learning in medicine: addressing ethical challenges, *PLoS Med.* 15 (11) (2018) e1002689.
- [9] E. Strubell, A. Ganesh, A. McCallum, Energy and policy considerations for modern deep learning research, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, (09) 2020, pp. 13693–13696.
- [10] S. Min, B. Lee, S. Yoon, Deep learning in bioinformatics, *Brief. Bioinform.* 18 (5) (2017) 851–869.
- [11] W. Guo, J. Wang, S. Wang, Deep multimodal representation learning: A survey, *IEEE Access* 7 (2019) 63373–63394.
- [12] S. Jabeen, X. Li, M.S. Amin, O. Bourahla, S. Li, A. Jabbar, A review on methods and applications in multimodal deep learning, *ACM Trans. Multimed. Comput. Commun. Appl.* 19 (2s) (2023) 1–41.
- [13] J. Gao, P. Li, Z. Chen, J. Zhang, A survey on deep learning for multimodal data fusion, *Neural Comput.* 32 (5) (2020) 829–864.
- [14] K. Bayouhd, R. Knani, F. Hamdaoui, A. Mtibaa, A survey on deep multimodal learning for computer vision: advances, trends, applications, and datasets, *Vis. Comput.* 38 (8) (2022) 2939–2970.
- [15] P.P. Liang, A. Zadeh, L.-P. Morency, Foundations & trends in multimodal machine learning: Principles, challenges, and open questions, *ACM Comput. Surv.* 56 (10) (2024) 1–42.
- [16] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, A.Y. Ng, Multimodal deep learning, in: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 689–696.
- [17] C. Zhang, Z. Yang, X. He, L. Deng, Multimodal intelligence: Representation learning, information fusion, and applications, *IEEE J. Sel. Top. Signal Process.* 14 (3) (2020) 478–493.
- [18] P. Xu, X. Zhu, D.A. Clifton, Multimodal learning with transformers: A survey, *IEEE Trans. Pattern Anal. Mach. Intell.* (2023).
- [19] T. Baltrušaitis, C. Ahuja, L.-P. Morency, Multimodal machine learning: A survey and taxonomy, *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (2) (2018) 423–443.
- [20] J. Summaira, X. Li, A.M. Shoib, S. Li, J. Abdul, Recent advances and trends in multimodal deep learning: A review, 2021, arXiv preprint arXiv:2105.11087.
- [21] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (8) (2013) 1798–1828.
- [22] Y. Zhu, Y. Wu, N. Sebe, Y. Yan, Vision+ x: A survey on multimodal learning in the light of data, *IEEE Trans. Pattern Anal. Mach. Intell.* (2024).
- [23] J.N. Acosta, G.J. Falcone, P. Rajpurkar, E.J. Topol, Multimodal biomedical AI, *Nature Med.* 28 (9) (2022) 1773–1784.
- [24] S.R. Stahlschmidt, B. Ulfenborg, J. Synnørgren, Multimodal deep learning for biomedical data fusion: a review, *Brief. Bioinform.* 23 (2) (2022) bbab569.
- [25] Y.-D. Zhang, Z. Dong, S.-H. Wang, X. Yu, X. Yao, Q. Zhou, H. Hu, M. Li, C. Jiménez-Mesa, J. Ramirez, et al., Advances in multimodal data fusion in neuroimaging: Overview, challenges, and novel orientation, *Inf. Fusion* 64 (2020) 149–187.
- [26] L. Heiliger, A. Sekuboyina, B. Menze, J. Egger, J. Kleesiek, Beyond medical imaging—a review of multimodal deep learning in radiology, *Authorea Prepr.* (2023).
- [27] J. Lipkova, R.J. Chen, B. Chen, M.Y. Lu, M. Barbieri, D. Shao, A.J. Vaidya, C. Chen, L. Zhuang, D.F. Williamson, et al., Artificial intelligence for multimodal data integration in oncology, *Cancer Cell* 40 (10) (2022) 1095–1110.
- [28] F. Behrad, M.S. Abadeh, An overview of deep learning methods for multimodal medical data mining, *Expert Syst. Appl.* 200 (2022) 117006.
- [29] T. Zhou, S. Ruan, S. Canu, A review: Deep learning for medical image segmentation using multi-modality fusion, *Array* 3 (2019) 100004.
- [30] Y. Xu, Deep learning in multimodal medical image analysis, in: *Health Information Science: 8th International Conference, HIS 2019, Xi’an, China, October 18–20, 2019, Proceedings 8*, Springer, 2019, pp. 193–200.
- [31] S.-C. Huang, A. Pareek, S. Seyyedi, I. Banerjee, M.P. Lungren, Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines, *NPJ Digit. Med.* 3 (1) (2020) 136.
- [32] L. Antonelli, M.R. Guarracino, L. Maddalena, M. Sangiovanni, Integrating imaging and omics data: a review, *Biomed. Signal Process. Control.* 52 (2019) 264–280.
- [33] M. Leil, M. Kachelrieß, Computed tomography 2.0: new detector technology, AI, and other developments, *Invest. Radiol.* 58 (8) (2023) 587–601.
- [34] J. Qian, H. Li, J. Wang, L. He, Recent advances in explainable artificial intelligence for magnetic resonance imaging, *Diagnostics* 13 (9) (2023) 1571.
- [35] C. Beegle, N. Hasani, R. Maass-Moreno, B. Saboury, E. Siegel, Artificial intelligence and positron emission tomography imaging workflow: technologists’ perspective, *PET Clin.* 17 (1) (2022) 31–39.
- [36] Y.-T. Shen, L. Chen, W.-W. Yue, H.-X. Xu, Artificial intelligence in ultrasound, *Eur. J. Radiol.* 139 (2021) 109717.
- [37] Z.I. Attia, D.M. Harmon, E.R. Behr, P.A. Friedman, Application of artificial intelligence to the electrocardiogram, *Eur. Heart J.* 42 (46) (2021) 4717–4730.
- [38] X. Zhou, C. Liu, Z. Wang, L. Zhai, Z. Jia, C. Guan, Y. Liu, Interpretable and robust ai in eeg systems: A survey, 2023, arXiv preprint arXiv:2304.10755.
- [39] M.-C. Nguyễn, Exploring Machine Learning Perspectives for Electrolottographic Signals (Ph.D. thesis), LIG (Laboratoire informatique de Grenoble), 2023.
- [40] S. Gao, J. Gong, B. Chen, B. Zhang, F. Luo, M.O. Yerabakan, Y. Pan, B. Hu, Use of advanced materials and artificial intelligence in electromyography signal detection and interpretation, *Adv. Intell. Syst.* 4 (10) (2022) 2200063.
- [41] H.F. Posada-Quintero, K.H. Chon, Innovations in electrodermal activity data collection and signal processing: A systematic review, *Sensors* 20 (2) (2020) 479.

- [42] X. Li, M. Huang, Y. Xu, Y. Cao, Y. Lu, P. Wang, X. Xiang, AviPer: assisting visually impaired people to perceive the world with visual-tactile multimodal attention network, *CCF Trans. Pervasive Comput. Interact.* 4 (3) (2022) 219–239.
- [43] S. Azizi, S. Bayat, P. Yan, A. Tahmasebi, J.T. Kwak, S. Xu, B. Turkbey, P. Choyke, P. Pinto, B. Wood, et al., Deep recurrent neural networks for prostate cancer detection: analysis of temporal enhanced ultrasound, *IEEE Trans. Med. Imaging* 37 (12) (2018) 2695–2703.
- [44] Y. Dong, Y. Qiu, D. Yang, L. Yu, D. Zuo, Q. Zhang, X. Tian, W.-P. Wang, E.M. Jung, Potential application of dynamic contrast enhanced ultrasound in predicting microvascular invasion of hepatocellular carcinoma, *Clin. Hemorheol. Microcirc.* 77 (4) (2021) 461–469.
- [45] S. Shetty, A. Mahale, et al., Multimodal medical tensor fusion network-based DL framework for abnormality prediction from the radiology CXRs and clinical text reports, *Multimedia Tools Appl.* (2023) 1–48.
- [46] X. Sun, W. Guo, J. Shen, Toward attention-based learning to predict the risk of brain degeneration with multimodal medical data, *Front. Neurosci.* 16 (2023) 1043626.
- [47] C. Aine, H.J. Bockholt, J.R. Bustillo, J.M. Cañive, A. Caprihan, C. Gasparovic, F.M. Hanlon, J.M. Houck, R.E. Jung, J. Lauriello, et al., Multimodal neuroimaging in schizophrenia: description and dissemination, *Neuroinformatics* 15 (2017) 343–364.
- [48] M.A. Rahaman, J. Chen, Z. Fu, N. Lewis, A. Iraj, T.G. van Erp, V.D. Calhoun, Deep multimodal predictome for studying mental disorders, *Hum. Brain Mapp.* 44 (2) (2023) 509–522.
- [49] M.A. Rahaman, J. Chen, Z. Fu, N. Lewis, A. Iraj, V.D. Calhoun, Multi-modal deep learning of functional and structural neuroimaging and genomic data to predict mental illness, in: 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society, EMBC, IEEE, 2021, pp. 3267–3272.
- [50] M.A. Rahaman, Y. Garg, A. Iraj, Z. Fu, J. Chen, V. Calhoun, Two-dimensional attentive fusion for multi-modal learning of neuroimaging and genomics data, in: 2022 IEEE 32nd International Workshop on Machine Learning for Signal Processing, MLSP, IEEE, 2022, pp. 1–6.
- [51] R. Liu, Z.-A. Huang, Y. Hu, Z. Zhu, K.-C. Wong, K.C. Tan, Attention-like multimodality fusion with data augmentation for diagnosis of mental disorders using MRI, *IEEE Trans. Neural Networks Learn. Syst.* (2022).
- [52] A.E. Johnson, T.J. Pollard, N.R. Greenbaum, M.P. Lungren, C.-y. Deng, Y. Peng, Z. Lu, R.G. Mark, S.J. Berkowitz, S. Hornig, MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs, 2019, arXiv preprint arXiv:1901.07042.
- [53] A. Lobantsev, N. Gusarova, A.S. Vatian, A.A. Kapitonov, A.A. Shalyto, Comparative assessment of text-image fusion models for medical diagnostics, *Информационно-управляющие системы* (5) (2020) 70–79.
- [54] J. Zeng, K. Niu, Y. Lu, S. Pei, Y. Zhou, Z. Guo, MIFTP: A multimodal multi-level independent fusion framework with improved twin pyramid for multilabel chest X-Ray image classification, in: 2022 IEEE 34th International Conference on Tools with Artificial Intelligence, ICTAI, IEEE, 2022, pp. 1112–1119.
- [55] N. Hayat, K.J. Geras, F.E. Shamout, MedFuse: Multi-modal fusion with clinical time-series data and chest X-ray images, in: Machine Learning for Healthcare Conference, PMLR, 2022, pp. 479–503.
- [56] M. Kohankhaki, A. Ayad, M. Barhoush, B. Leibe, A. Schmeink, Radiopaths: Deep multimodal analysis on chest radiographs, in: 2022 IEEE International Conference on Big Data (Big Data), IEEE, 2022, pp. 3613–3621.
- [57] C.R. Jack Jr., M.A. Bernstein, N.C. Fox, P. Thompson, G. Alexander, D. Harvey, B. Borowski, P.J. Britson, J. L. Whitwell, C. Ward, et al., The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods, *J. Magn. Reson. Imaging: An Off. J. Int. Soc. Magn. Reson. Med.* 27 (4) (2008) 685–691.
- [58] C. Ostertag, M. Visani, T. Urruty, M. Beurton-Aimar, Long-term cognitive decline prediction based on multi-modal data using Multimodal3DSiameseNet: transfer learning from Alzheimer's disease to Parkinson's disease, *Int. J. Comput. Assist. Radiol. Surg.* 18 (5) (2023) 809–818.
- [59] C. Ostertag, M. Beurton-Aimar, M. Visani, T. Urruty, K. Bertet, Predicting brain degeneration with a multimodal siamese neural network, in: 2020 Tenth International Conference on Image Processing Theory, Tools and Applications, IPTA, IEEE, 2020, pp. 1–6.
- [60] R. Subramanian, J. Wache, M.K. Abadi, R.L. Vieriu, S. Winkler, N. Sebe, ASCERTAIN: Emotion and personality recognition using commercial sensors, *IEEE Trans. Affect. Comput.* 9 (2) (2016) 147–160.
- [61] K. Radhika, V.R.M. Oruganti, Deep multimodal fusion for subject-independent stress detection, in: 2021 11th International Conference on Cloud Computing, Data Science & Engineering (Confluence), IEEE, 2021, pp. 105–109.
- [62] R. Kuttala, R. Subramanian, V.R.M. Oruganti, Multimodal hierarchical CNN feature fusion for stress detection, *IEEE Access* 11 (2023) 6867–6878.
- [63] K. Radhika, V.R.M. Oruganti, Stress detection using CNN fusion, in: TENCON 2021-2021 IEEE Region 10 Conference, TENCON, IEEE, 2021, pp. 492–497.
- [64] V. Markova, T. Ganchev, K. Kalinkov, Clas: A database for cognitive load, affect and stress recognition, in: 2019 International Conference on Biomedical Innovations and Applications, BIA, IEEE, 2019, pp. 1–4.
- [65] D.B. Keator, T.G. van Erp, J.A. Turner, G.H. Glover, B.A. Mueller, T.T. Liu, J.T. Voyvodic, J. Rasmussen, V.D. Calhoun, H.J. Lee, et al., The function biomedical informatics research network data repository, *Neuroimage* 124 (2016) 1074–1079.
- [66] B.M. Adhikari, L.E. Hong, H. Sampath, J. Chiappelli, N. Jahanshad, P.M. Thompson, L.M. Rowland, V.D. Calhoun, X. Du, S. Chen, et al., Functional network connectivity impairments and core cognitive deficits in schizophrenia, *Hum. Brain Mapp.* 40 (16) (2019) 4593–4605.
- [67] S. Bakr, O. Gevaert, S. Echeagaray, K. Ayers, M. Zhou, M. Shafiq, H. Zheng, J.A. Benson, W. Zhang, A.N. Leung, et al., A radiogenomic dataset of non-small cell lung cancer, *Sci. Data* 5 (1) (2018) 1–9.
- [68] G. Hou, L. Jia, Y. Zhang, W. Wu, L. Zhao, J. Zhao, L. Wang, Y. Qiang, Deep learning approach for predicting lymph node metastasis in non-small cell lung cancer by fusing image–gene data, *Eng. Appl. Artif. Intell.* 122 (2023) 106140.
- [69] H. Wang, V. Subramanian, T. Syeda-Mahmood, Modeling uncertainty in multimodal fusion for lung cancer survival analysis, in: 2021 IEEE 18th International Symposium on Biomedical Imaging, ISBI, IEEE, 2021, pp. 1169–1172.
- [70] V. Subramanian, M.N. Do, T. Syeda-Mahmood, Multimodal fusion of imaging and genomics for lung cancer recurrence prediction, in: 2020 IEEE 17th International Symposium on Biomedical Imaging, ISBI, IEEE, 2020, pp. 804–808.
- [71] The clinical proteomic tumor analysis consortium pancreatic ductal adenocarcinoma collection (CPTAC-PDA), 2018, <https://www.cancerimagingarchive.net/collection/cptac-pda/#citations>. Online; (Accessed 2024-04-16).
- [72] A.B. Menegotto, C.D.L. Becker, S.C. Cazella, Computer-aided diagnosis of hepatocellular carcinoma fusing imaging and structured health data, *Heal. Inf. Sci. Syst.* 9 (1) (2021) 20.
- [73] A.B. Menegotto, C.D. Lopes Becker, S.C. Cazella, Computer-aided hepatocarcinoma diagnosis using multimodal deep learning, in: Ambient Intelligence—Software and Applications—, 10th International Symposium on Ambient Intelligence, Springer, 2020, pp. 3–10.
- [74] A.E. Johnson, T.J. Pollard, L. Shen, L.-w.H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L.A. Celi, R.G. Mark, Data descriptor: MIMIC-III, a freely accessible critical care database, *Sci. Data* 3 (160035) (2016) 1–9.
- [75] K. Niu, K. Zhang, X. Peng, Y. Pan, N. Xiao, Deep multi-modal intermediate fusion of clinical record and time series data in mortality prediction, *Front. Mol. Biosci.* 10 (2023) 1136071.
- [76] M. Ma, X. Hao, J. Zhao, S. Luo, Y. Liu, D. Li, Predicting heart failure in-hospital mortality by integrating longitudinal and category data in electronic health records, *Med. Biol. Eng. Comput.* (2023) 1–17.
- [77] A. Johnson, A. Johnson, L. Bulgarelli, T. Pollard, S. Hornig, L.A. Celi, R. Mark, MIMIC-IV (version 2.2), 2023, <https://physionet.org/content/mimiciv/2.2>. Online; (Accessed 2024-04-16).
- [78] M. Linehan, R. Gautam, S. Kirk, Y. Lee, C. Roche, E. Bonaccio, J. Filippini, K. Rieger-Christ, J. Lemmerman, R. Jarosz, The cancer genome atlas cervical kidney renal papillary cell carcinoma collection (TCGA-KIRP) (version 4), 2016, <https://www.cancerimagingarchive.net/collection/tcga-kirp/>. Online; (Accessed 2024-04-16).
- [79] B. Erickson, S. Kirk, Y. Lee, O. Bathe, M. Kearns, C. Gerdes, K. Rieger-Christ, J. Lemmerman, The cancer genome atlas liver hepatocellular carcinoma collection (TCGA-LIHC) (version 5), 2016, <https://www.cancerimagingarchive.net/collection/tcga-lihc/>. Online; (Accessed 2024-04-16).
- [80] F. Lucchesi, N. Aredes, The cancer genome atlas stomach adenocarcinoma collection (TCGA-STAD) (version 3), 2016, <https://doi.org/10.7937/K9/TCIA.2016.GDHL9KIM>. Online; (Accessed 2024-04-16).
- [81] A. Auremma Citarella, L. Di Biasi, F. De Marco, G. Tortora, ENTAIL: yet another amyloid fibrils classifier, *BMC Bioinformatics* 23 (1) (2022) 517.
- [82] Y. Yuan, L. Wu, X. Zhang, Gini-impurity index analysis, *IEEE Trans. Inf. Forensics Secur.* 16 (2021) 3154–3169.
- [83] N.V. Chawla, K.W. Bowyer, L.O. Hall, W.P. Kegelmeyer, SMOTE: synthetic minority over-sampling technique, *J. Artificial Intelligence Res.* 16 (2002) 321–357.
- [84] S. Seo, Y. Kim, H.-J. Han, W.C. Son, Z.-Y. Hong, I. Sohn, J. Shim, C. Hwang, Predicting successes and failures of clinical trials with outer product-based convolutional neural network, *Front. Pharmacol.* 12 (2021) 670670.
- [85] J. Ren, C. Yu, X. Ma, H. Zhao, S. Yi, et al., Balanced meta-softmax for long-tailed visual recognition, *Adv. Neural Inf. Process. Syst.* 33 (2020) 4175–4186.
- [86] Y. Li, S. Long, Z. Yang, H. Weng, K. Zeng, Z. Huang, F.L. Wang, T. Hao, A bi-level representation learning model for medical visual question answering, *J. Biomed. Informatics* 134 (2022) 104183.
- [87] N. Cahan, E. Klang, E.M. Marom, S. Soffer, Y. Barash, E. Burshtein, E. Konen, H. Greenspan, Multimodal fusion models for pulmonary embolism mortality prediction, *Sci. Rep.* 13 (1) (2023) 1–15.
- [88] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 2980–2988.
- [89] N. Cahan, E.M. Marom, S. Soffer, Y. Barash, E. Konen, E. Klang, H. Greenspan, Weakly supervised multimodal 30-day all-cause mortality prediction for pulmonary embolism patients, in: 2022 IEEE 19th International Symposium on Biomedical Imaging, ISBI, IEEE, 2022, pp. 1–4.

- [90] G. Holste, S.C. Partridge, H. Rahbar, D. Biswas, C.I. Lee, A.M. Alessio, End-to-end learning of fused image and non-image features for improved breast cancer classification from mri, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 3294–3303.
- [91] A. Bhattacharya, S. Sadasivuni, C.-J. Chao, P. Agasthi, C. Ayoub, D.R. Holmes, R. Arsanjani, A. Sanyal, I. Banerjee, Multi-modal fusion model for predicting adverse cardiovascular outcome post percutaneous coronary intervention, *Physiol. Meas.* 43 (12) (2022) 124004.
- [92] J.V. Sousa, P. Matos, F. Silva, P. Freitas, H.P. Oliveira, T. Pereira, Single modality vs. multimodality: What works best for lung cancer screening? *Sensors* 23 (12) (2023) 5597.
- [93] M. He, K. Han, Y. Zhang, W. Chen, Hierarchical-order multimodal interaction fusion network for grading gliomas, *Phys. Med. Biol.* 66 (21) (2021) 215016.
- [94] C. Pan, P. Zhou, J. Tan, B. Sun, R. Guan, Z. Wang, Y. Luo, J. Lu, Liver tumor detection via a multi-scale intermediate multi-modal fusion network on MRI images, in: 2021 IEEE International Conference on Image Processing, ICIP, IEEE, 2021, pp. 299–303.
- [95] I. Ahmad, Y. Xia, H. Cui, Z.U. Islam, AATS: Anatomy aware tumor segmentation network for PET-CT volumes and images using a lightweight fusion-attention mechanism, *Comput. Biol. Med.* 157 (2023) 106748.
- [96] Y. Li, B. Zou, J. Wu, Y. Dai, H.X. Bai, Z. Jiao, A dynamic multi-modal fusion network for ovarian tumor differentiation, in: 2022 IEEE International Conference on Bioinformatics and Biomedicine, BIBM, IEEE, 2022, pp. 767–772.
- [97] Z. Meng, Y. Zhu, W. Pang, J. Tian, F. Nie, K. Wang, MSMFN: An ultrasound based multi-step modality fusion network for identifying the histologic subtypes of metastatic cervical lymphadenopathy, *IEEE Trans. Med. Imaging* 42 (4) (2022) 996–1008.
- [98] S.-C. Huang, A. Pareek, R. Zamanian, I. Banerjee, M.P. Lungren, Multimodal fusion with deep neural networks for leveraging CT imaging and electronic health record: a case-study in pulmonary embolism detection, *Sci. Rep.* 10 (1) (2020) 22147.
- [99] M.A. Alsherbiny, I. Radwan, N. Moustafa, D.J. Bhuyan, M. El-Waisi, D. Chang, C.G. Li, Trustworthy deep neural network for inferring anticancer synergistic combinations, *IEEE J. Biomed. Heal. Informatics* 27 (4) (2021) 1691–1700.
- [100] G. Cen, C. Wang, T.A. Olugbade, A.C.d.C. Williams, N. Bianchi-Berthouze, Exploring multimodal fusion for continuous protective behavior detection, in: 2022 10th International Conference on Affective Computing and Intelligent Interaction, ACII, IEEE, 2022, pp. 1–8.
- [101] T. Chen, R. Hong, Y. Guo, S. Hao, B. Hu, MS²-GNN: Exploring GNN-based multimodal fusion network for depression detection, *IEEE Trans. Cybern.* (2022).
- [102] G. Di Teodoro, F. Siciliano, V. Guarrasi, A.-M. Vandamme, V. Ghisetti, A. Sönnberg, M. Zazzi, F. Silvestri, L. Palagi, A graph neural network-based model with out-of-distribution robustness for enhancing antiretroviral therapy outcome prediction for HIV-1, *Comput. Med. Imaging Graph.* (2025) 102484.
- [103] J.N. Njoku, A.C. Caliwag, W. Lim, S. Kim, H. Hwang, J. Jung, Deep learning based data fusion methods for multimodal emotion recognition, *J. Korean Inst. Commun. Inf. Sci.* 47 (1) (2022) 79–87.
- [104] H.M. Mohammed, A.N. Omeroglu, E.A. Oral, MMHFNet: Multi-modal and multi-layer hybrid fusion network for voice pathology detection, *Expert Syst. Appl.* 223 (2023) 119790.
- [105] A. Sedghi, A. Mehrtash, A. Jamzad, A. Amalou, W.M. Wells, T. Kapur, J.T. Kwak, B. Turkbey, P. Choyke, P. Pinto, et al., Improving detection of prostate cancer foci via information fusion of MRI and temporal enhanced ultrasound, *Int. J. Comput. Assist. Radiol. Surg.* 15 (2020) 1215–1223.
- [106] Y. LeCun, B. Boser, J.S. Denker, D. Henderson, R.E. Howard, W. Hubbard, L.D. Jackel, Backpropagation applied to handwritten zip code recognition, *Neural Comput.* 1 (4) (1989) 541–551.
- [107] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning internal representations by error propagation, parallel distributed processing, explorations in the microstructure of cognition, ed. de rumelhart and j. mccllland. vol. 1. 1986, *Biometrika* 71 (599–607) (1986) 6.
- [108] M. Gori, G. Monfardini, F. Scarselli, A new model for learning in graph domains, in: Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005, vol. 2, IEEE, 2005, pp. 729–734.
- [109] A. Waswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: NIPS, 2017.
- [110] C. Guida, M. Zhang, J. Shan, Improving knee osteoarthritis classification using multimodal intermediate fusion of X-ray, MRI, and clinical information, *Neural Comput. Appl.* (2023) 1–10.
- [111] S. Oh, S.-R. Kang, I.-J. Oh, M.-S. Kim, Deep learning model integrating positron emission tomography and clinical data for prognosis prediction in non-small cell lung cancer patients, *BMC Bioinformatics* 24 (1) (2023) 1–13.
- [112] K. Zheng, S. Cai, H.R. Chua, M. Herschel, M. Zhang, B.C. Ooi, DyHealth: making neural networks dynamic for effective healthcare analytics, *Proc. the VLDB Endow.* 15 (12) (2022) 3445–3458.
- [113] Y. Zhang, X. Jian, L. Xu, J. Zhao, M. Lu, Y. Lin, L. Xie, iTcep: a deep learning framework for identification of T cell epitopes by harnessing fusion features, *Front. Genet.* 14 (2023) 1141535.
- [114] V. Guarrasi, P. Soda, Multi-objective optimization determines when, which and how to fuse deep networks: An application to predict COVID-19 outcomes, *Comput. Biol. Med.* 154 (2023) 106625.
- [115] E. Perez, S. Ventura, Multi-view deep neural networks for multiclass skin lesion diagnosis, in: 2022 IEEE International Conference on Omni-Layer Intelligent Systems, COINS, IEEE, 2022, pp. 1–6.
- [116] S. Steyaert, Y.L. Qiu, Y. Zheng, P. Mukherjee, H. Vogel, O. Gevaert, Multimodal deep learning to predict prognosis in adult and pediatric brain tumors, *Commun. Med.* 3 (1) (2023) 44.
- [117] Y. Wu, M. Daoudi, A. Amad, Transformer-based self-supervised multimodal representation learning for wearable emotion recognition, *IEEE Trans. Affect. Comput.* (2023).
- [118] Z. Han, F. Yang, J. Huang, C. Zhang, J. Yao, Multimodal dynamics: Dynamical fusion for trustworthy multimodal classification, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 20707–20717.
- [119] H.-A. Rashid, U. Kallakuri, T. Mohsenin, TinyM²Net-V2: A compact low power software hardware architecture for multimodal deep neural networks, *ACM Trans. Embed. Comput. Syst.* (2023).
- [120] J. Jiao, H. Sun, Y. Huang, M. Xia, M. Qiao, Y. Ren, Y. Wang, Y. Guo, GMRL-Net: A graph-based manifold regularization learning framework for placental insufficiency diagnosis on incomplete multimodal ultrasound data, *IEEE Trans. Med. Imaging* (2023).
- [121] Y. Li, M. El Habib Daho, P.-H. Conze, H. Al Hajj, S. Bonnin, H. Ren, N. Manivannan, S. Magazzeni, R. Tadayoni, B. Cochener, et al., Multimodal information fusion for glaucoma and diabetic retinopathy classification, in: International Workshop on Ophthalmic Medical Image Analysis, Springer, 2022, pp. 53–62.
- [122] L.S. Rodrigues, K. Sakiyama, E. Takashi Matsubara, J. Marcato Junior, W.N. Gonçalves, Multimodal fusion based on arithmetic operations and attention mechanisms, 2022, Available At SSRN 4292754.
- [123] F. Zhao, C. Zhang, B. Geng, Deep multimodal data fusion, *ACM Comput. Surv.* 56 (9) (2024) 1–36.
- [124] H.R.V. Joze, A. Shaban, M.L. Iuzzolino, K. Koishida, MMTM: Multimodal transfer module for CNN fusion, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 13289–13299.
- [125] Y. Qiu, F. Lin, W. Chen, M. Xu, Pre-training in medical data: A survey, *Mach. Intell. Res.* 20 (2) (2023) 147–179.
- [126] Y. Qiu, F. Lin, W. Chen, M. Xu, Pre-training in medical data: A survey, *Mach. Intell. Res.* 20 (2) (2023) 147–179.
- [127] A.A. Mukhlif, B. Al-Khateeb, M.A. Mohammed, An extensive review of state-of-the-art transfer learning techniques used in medical imaging: Open issues and challenges, *J. Intell. Syst.* 31 (1) (2022) 1085–1111.
- [128] C. Matsoukas, J.F. Haslum, M. Sorkhei, M. Söderberg, K. Smith, What makes transfer learning work for medical images: Feature reuse & other factors, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 9225–9234.
- [129] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., Imagenet large scale visual recognition challenge, *Int. J. Comput. Vis.* 115 (2015) 211–252.
- [130] A. Mumuni, F. Mumuni, Data augmentation: A comprehensive survey of modern approaches, *Array* 16 (2022) 100258.
- [131] I. Gat, I. Schwartz, A. Schwing, T. Hazan, Removing bias in multi-modal classifiers: Regularization by maximizing functional entropies, *Adv. Neural Inf. Process. Syst.* 33 (2020) 3197–3208.
- [132] W. Samek, K.-R. Müller, Towards explainable artificial intelligence, Explain. AI: Interpret. Explain. Vis. Deep. Learn. (2019) 5–22.
- [133] G. Joshi, R. Walambe, K. Kotecha, A review on explainability in multimodal deep neural nets, *IEEE Access* 9 (2021) 59800–59821.
- [134] C.M. Caruso, V. Guarrasi, E. Cordelli, R. Sicilia, S. Gentile, L. Messina, M. Fiore, C. Piccolo, B. Beomonte Zobel, G. Iannello, et al., A multimodal ensemble driven by multiobjective optimisation to predict overall survival in non-small-cell lung cancer, *J. Imaging* 8 (11) (2022) 298.
- [135] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2009, pp. 248–255.
- [136] V. Guarrasi, L. Tronchin, D. Albano, E. Faiella, D. Fazzini, P. Soda, Multimodal explainability via latent shift applied to COVID-19 stratification, *Pattern Recognit.* (2024) 110825.
- [137] A. Rofena, V. Guarrasi, M. Sarli, C.L. Piccolo, M. Sammarra, B.B. Zobel, P. Soda, A deep learning approach for virtual contrast enhancement in contrast enhanced spectral mammography, *Comput. Med. Imaging Graph.* 116 (2024) 102398.
- [138] H. Poon, Multimodal generative AI for precision health, *NEJM AI Spons.* (2023).