



HAL
open science

How will tomorrow's algorithms fuse multimodal data? The example of the neuroprognosis in Intensive Care

Théodore Soulier, Olivier Colliot, Nicholas Ayache, Benjamin Rohaut

► To cite this version:

Théodore Soulier, Olivier Colliot, Nicholas Ayache, Benjamin Rohaut. How will tomorrow's algorithms fuse multimodal data? The example of the neuroprognosis in Intensive Care. *Anaesthesia Critical Care & Pain Medicine*, In press, pp.101301. 10.1016/j.accpm.2023.101301 . hal-04227434

HAL Id: hal-04227434

<https://hal.sorbonne-universite.fr/hal-04227434v1>

Submitted on 3 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Title:

How will tomorrow's algorithms fuse multimodal data? The example of the neuroprognosis in Intensive Care.

Authors :

Théodore Soulier¹, Olivier Colliot², Nicholas Ayache³, Benjamin Rohaut^{1,4}

Affiliations:

¹ *Sorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, CNRS, Inserm, AP-HP, Hôpital de la Pitié Salpêtrière, F-75013, Paris, France*

² *Sorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, CNRS, Inria, Inserm, AP-HP, Hôpital de la Pitié Salpêtrière, F-75013, Paris, France*

³ *Inria, Epione Project-Team, Sophia-Antipolis, France*

⁴ *Department of Neurology, Groupe hospitalier Pitié-Salpêtrière, AP-HP, Paris, France*

Corresponding author: Théodore Soulier (theodore.soulier@aphp.fr)

ORCID numbers:

Théodore Soulier: 0009-0005-4726-493X

Olivier Colliot: 0000-0002-9836-654X

Benjamin Rohaut: 0000-0001-6752-8756

Acknowledgements:

The research leading to these results has received funding from the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute) and reference ANR-10-IAIHU-06 (Agence Nationale de la Recherche-10-IA Institut Hospitalo-Universitaire-6), and from ICM under the Big Brain Theory program (project IMAGIN-DEAL in MS). This work was supported by the Fondation pour la Recherche Médicale, grant number FDM202006011247 to Théodore. Nicholas Ayache's research is also partially funded by the 3IA Côte d'Azur (reference number ANR-19-P3IA-0002).

Competing interests:

Théodore Soulier, Nicholas Ayache and Benjamin Rohaut: nothing to disclose

Olivier Colliot. Competing financial interests related to the present article: none to disclose. Competing financial interests unrelated to the present article: OC reports having received consulting fees from AskBio and Therapanacea and that his laboratory has received grants (paid to the institution) from Qynapse. Members from his laboratory have co-supervised a PhD thesis with Qynapse. OC's spouse is an employee and holds stock-options of myBrainTechnologies. OC holds a patent registered at the International Bureau of the World Intellectual Property Organization (PCT/IB2016/0526993, Schiratti J-B, Allasonniere S, Colliot O, Durrleman S, A method for determining the temporal progression of a biological phenomenon and associated methods and devices) (2017).

Word count: 2831

Keywords:

Disorders of Consciousness, Neurological Prognosis, Multimodal Data, Artificial Intelligence

Disorders of Consciousness:

Disorders of Consciousness can affect any brain-injured patient in the intensive care unit, and correspond to the loss, at least in part, of behavioral signs of consciousness (1). Consciousness can be defined as the ability to be aware of mental representations (e.g. external stimulus, memory (2)). It requires wakefulness which depends on the ascending activating reticular formation but this is not sufficient, thus a patient may be awake but not conscious (3). Conscious processing allowing awareness is usually probed through subject's reportability: a conscious representation can be reported to oneself and another person. The classification of Disorders of Consciousness (4) distinguishes the Vegetative State (VS), corresponding to observable wakefulness (spontaneous eyes opening) without signs of awareness (also called Unresponsive Wakefulness Syndrome), from Minimal Conscious State and Conscious State (MCS). MCS corresponds to a minimal stage of consciousness, characterized by wakefulness with behavioral signs suggestive of conscious processing (awareness) of external stimuli, including at most language processing (MCS+, (3)), but no ability to report mental representations to the observer that belong to Conscious State.

Being able to classify a patient between those categories helps for the long-term functional neuroprognosis and thus, also has immediate implications in the ICU such as adjusting the goal-of-care (continuation, withholding or withdrawing life sustaining therapies) to patient and family wishes (1).

Modalities used for Disorders of Consciousness evaluation:

Methods for investigating the state of consciousness in Disorders of Consciousness patients can be separated into two categories that address different questions, whatever the modality. The first, corresponding to a necessary condition, is to probe the structural integrity of the cerebral network that is required for conscious processing. The second, which is a sufficient condition, is to functionally test this network by demonstrating a conscious response to a given stimulus.

The assessment of consciousness and neuroprognosis is multimodal as it encompasses clinical, biological, brain imaging and electrophysiological investigation:

- The first investigation is clinical (4). The most basic clinical assessment is to test the response to simple and complex commands, enabling physicians first to verify the ability to integrate stimuli from the outside world and then to produce a non-reflex response in return. Certain reflexes can also be assessed, reflecting either the integrity of the anatomical structures required for conscious information processing. Today, this clinical assessment is structured by scales like the Full Outline of UnResponsiveness (FOUR) score and the Coma Recovery Scale-Revised (CRS-R). Those ratings may lack reproducibility (5).

Behavioral assessment is constantly improving. The ability to inhibit reflexes can be informative reflecting the ability to integrate a repetitive pattern from the environment and top down inhibition capacities (6). In addition to structured scores or paraclinical assessment, it has been shown that caregivers' unstructured overall impression ("gut feeling") can also be useful (7).

Despite incontestable progresses, clinical assessment can be flawed and a now classical caveat of an exclusive behavioral approach is the case of Cognitive Motor Dissociation, in which cerebral processing of the external stimulus does exist but without any behavioral translation, the incidence of which is not marginal (3), and that is distinct from locked-in syndrome in which oculomotricity is partially preserved.

The automatic collection of vitals on the scope of patients admitted in ICU has recently shown (8) its relevance in terms of predictive capacity for predicting confusion and could probably be useful as a clinical tool in Disorders of Consciousness in the future as it could help detecting Cognitive Motor

Dissociation. Finally, metrics such as the Neurological Pupil Index obtained with a pupilometer can improve clinical evaluation and make it more quantitative, that has already shown its utility for the prognosis after cardiac arrest (9).

- Biological paraclinical assessments. Blood Neuron Specific Enolase (NSE) concentration, as a marker quantifying neuronal suffering, has been shown to be a key neuroprognostic marker in cardiac arrest (10). Similarly, the assessment of blood levels of tau and long-chain neurofilaments looks promising (11).

- Imaging-based paraclinical assessments. While cerebral Computed Tomography (CT) scan is easily accessible and can be used to assess the extent of lesions in hemorrhagic cases (primary or secondary to head trauma), cerebral Magnetic Resonance Imaging (MRI) has demonstrated its superiority for neurological prognosis thanks to a more detailed assessment of the brain parenchyma. Structural MRI, thanks to advanced sequences enabling studying the integrity of underlying white matter bundles, e.g. through the measure of water diffusion anisotropy (FA), has been shown to perform well in post-cardiac arrest (12) (13). Functional MRI (fMRI), based on the indirect measurement of the local increase in Cerebral Blood Flow in response to the solicitation of a relevant brain area, enables to measure the response to an external task. In this way, fMRI has made it possible to improve the classification of patients with a good prognosis, mistakenly stamped Vegetative State on their clinical evaluation (14). More recently, Positron Emission Tomography (PET) has proven to give valuable information for Disorders of Consciousness, by measuring glucose consumption and thus local energy activity in the cortex (15).

- Electrophysiological examinations. Electroencephalography (EEG) is a minimally invasive, relatively inexpensive technique, accessible from the patient's bed in the ICU, with a temporal resolution of the order of a millisecond, compatible with the measurement of the electrical activity of large neural networks. In addition to diagnosing confounding factors that can alter consciousness (encephalopathy, epileptic seizures), resting EEG measurements can be used to extract different frequencies from different regions (central, temporal, frontal, occipital), corresponding to cerebral rhythms. These rhythms and their spatialization provide information on their own; indeed, the presence of a posterior alpha rhythm, for example, is correlated with a high state of consciousness (16). The study of the synchronization of these rhythms between different regions also makes it possible to estimate the functional connectivity between distant brain areas (16) implicated in conscious processing (2). EEG recordings can also be used to assess the electrical response to certain stimuli. In this respect, a particularly elegant paradigm is the Event Related Potentials (ERPs) of Bekinstein *et al.* to check the ability to detect the presence of a deviant sound (or sequence of sounds) in a series of repetitive sounds (or sequences of sounds) (17), and to see whether this discordance generates an automatic “local” detection of deviant sounds or a conscious “global” detection of deviant sequences across the different EEG areas.

Why fusing modalities in Disorders of Consciousness?

To provide a prognosis, neurointensivists have to perform a multiclass classification i.e. to associate a given patient with a probability of belonging to the class of Conscious State, Vegetative State, or Minimal Conscious State. A way to ensure the reliability of this classification is to gather complementary information provided by several sources. In today's multi-disciplinary staff meetings in specialized centers, the neurointensivist works as a multimodal integrator. She or he gathers, weighs and synthesizes information from different clinical, imaging and electrophysiological modalities.

Each modality can itself be pre-converted by supervised learning algorithms into a probability of belonging to the Vegetative State, Minimal Conscious State or Conscious state class (18). Using a

training database, supervised learning learns the transformation that associates input features with an output, in this case, a probability of belonging to a prognostic class (19). The features given as input may themselves correspond to an advanced extraction from the raw information acquired (average FA, requiring processing steps compared with the raw diffusion MRI acquisition), or to the raw information (such as NSE blood concentration (10)). Support Vector Machine (20) and logistic regression (12) are for example already available as supervised learning methods published for unimodal prognosis of Disorders of Consciousness.

Deep learning, which has been booming in the biomedical field since 2010s (19), refers to the methods used to achieve learning, based on connections between unitary computational entities called artificial neurons. These artificial neurons can receive several inputs, give several outputs and be organized in layers by analogy with biological neurons (21). During learning, the first layers specialize in extracting low-level features from the input, and in the deeper layers, combinations of combinations of neurons extract so-called higher-level features. Besides the fact that they do not require a priori pre-selection of input features, the advantage of neural networks is that they allow multiple combinations of input features, thus accessing a so-called latent representation of these input features, not immediately accessible on the input but only via the interplay of different connections during the learning process. The latent representation is a set of variables that is more compact (fewer variables) and powerful (each variable is more relevant to the classification and/or the prognosis) representation than the original set of input variables. For instance, variables representing different EEG frequency bands can be combined into a single variable in the latent representation, this variable being more predictive than a single frequency band alone.

Challenges in multimodal fusion for Disorders of Consciousness:

Whatever the supervised learning method, the integration of multimodal input data is always a challenge (19), and is not specific to biomedical problems (22). The taxonomy of different multimodal fusion methods in the existing literature classifies the methods according to the type of representation (joint or coordinate), translation, registration, fusion timing and operation, and choice of co-learning (23). As far as representation is concerned, we can imagine that, for our Disorders of Consciousness problem, each modality, for example EEG or MRI is immediately merged with the other (joint representation), or processed in parallel (coordinated representation). Similarly, the choice could be made to translate from one modality to the other (e.g. translate the lesional slow waves in EEG that a parenchymal lesion would eventually give to imaging) to have a "reference" modality into which we bring the other modalities by translation (24). Indeed, generative deep learning networks have given excellent results in biomedical inter modality translation (25). This question has been taken one step further in the work from Antelmi *et al.* (26) that demonstrated the possibility, for dementia paraclinical evaluation, to resynthesize all different modalities like medical records, MRI or PET, from one to another by using a common joint latent representation. The question of registration between modalities of very different nature, structure and dimensionality (e.g. between functional MRI and the concentration of a serum biomarker) is also crucial. To be processed together, the data scientist can either perform normalization of the different modalities before they are entered as input into the model, or take advantage of the model's internal transformations to realign them with each other. Probably the most crucial question is that of the merging process itself. First of all, the operations used to merge the different inputs can either be based on models operating as small integrating funnels, which are not necessarily based on deep learning methods (27). Within deep learning methods (28), the actual merging operation may be a simple concatenation, or may be based on a predetermined mathematical operation (tensor based) (29), or it may be guided by attention processes (30). In neural networks, the attention mechanism generally refers to a local connection module that focuses the learning and

feature extraction process on certain parts of the input. In image-based deep learning networks, this could for instance ensure that the network focuses on the parenchyma on a brain MRI, rather than the jaw.

Early, late and intermediate fusion – pro and contras in Disorders of Consciousness:

More than the fusion operation itself, one of the main challenges for future algorithms will be to find the right place within the model to perform this multimodal fusion. Such issue has already been reviewed in oncology (31). A distinction is made between early, intermediate and late multimodal fusion models (22). For example, the deep-layer architecture found in deep learning networks enables connections to be made between parallel branches, for example from different modalities, to connect them together **[FIGURE 1]**. From then on, the question arises of where to connect the parts of the neural network together to enable the best output prediction.

Late fusion. The more layers of neurons that independently learn input features before merging, the more likely it is that latent input-specific features will be extracted, resulting in a unimodal marginal representation (22). The neurointensivist's current work, during dedicated staff meetings in referral centers, consists today precisely in this late fusion (1). It combines the prognostic probability given by each modality, weighting them based for instance on quality criteria or on the amount of evidence for a given etiology, on a case-by-case basis. This late merging has the advantage of the explainability of each modality's weighting in the final decision, since they are only merged when the final prognosis is delivered. It is even possible to use a second simple supervised learning (e.g. support vector machine, logistic regression) model taking these unimodal probabilities as input to merge them, the study of which could be particularly interesting to assess the final weight of each modality on the final prognostic decision learned during the learning process.

Finally, this late fusion lends itself easily to transfer learning. In artificial neural networks, transfer learning corresponds to the operation of first training a model on an unspecific task for which large databases are available (e.g. natural image recognition) and then fine-tuning the model to perform the specific task of interest (e.g. prognostication of Disorders of Consciousness). It uses the fact that the first layers of neurons extract very general, low-level features from the input image, while the deeper layers extract high-level features, specific to the question posed during the learning process. When fine-tuning the model to perform the specific task of interest, the first pre-trained layers are "frozen", and the learning process focuses on the deepest layers, thus saving time and ensuring the generalizability and reliability of the low-level concepts learned in the first layers, thanks to the large databases used to train them. The benefit of this process has been demonstrated in medical image classification (32). For the case of Disorders of Consciousness with late fusion, we can imagine that the EEG branch would have its first layers pre-trained on thousands of EEGs unrelated to the Disorders of Consciousness problem, and the same for the MRI branch. Only the deeper layers of each branch would be trained on Disorders of Consciousness-specific databases. Taking a step further, if a new modality emerges for Disorders of Consciousness assessment, it could be trained independently in another model that could be merged as a new branch on this late fusion process, without having to retrain all pre-existing branches.

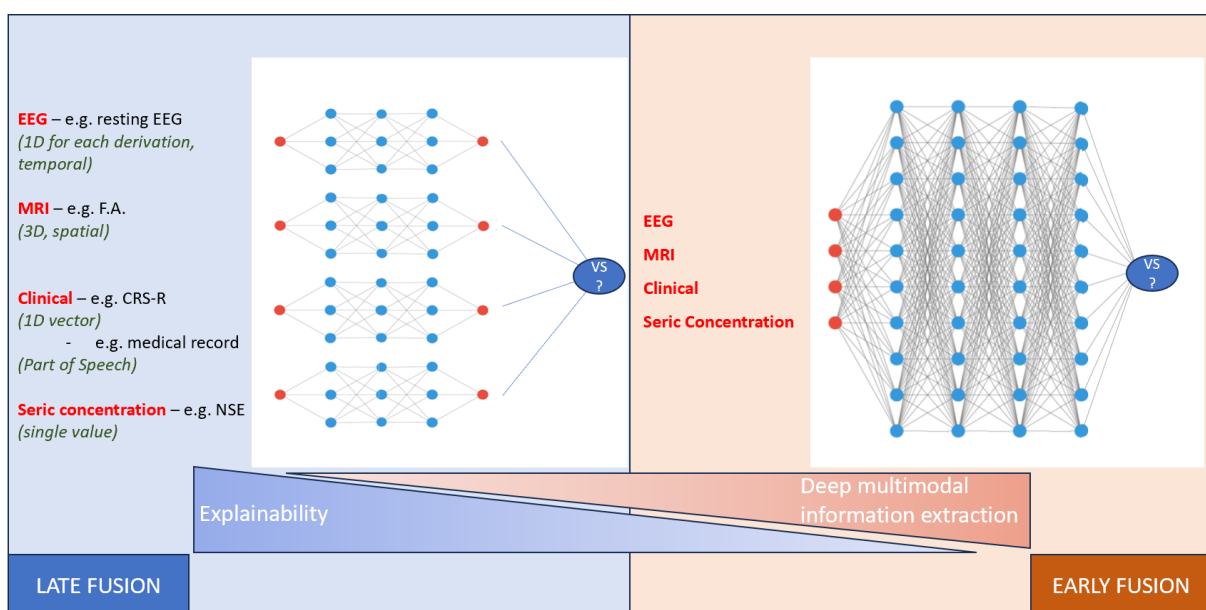
Early fusion. While late fusion is very similar to the neurointensivist's current work, an earlier fusion of the different modalities would have many advantages. It enables to learn complex combinations of the different modalities from the beginning thanks to combinations between neurons, and thus to have several deep layers that will simultaneously process these different modalities and create combinations of combinations between them. This early fusion enables the creation of a multimodal combined latent space, instead of marginal representations (22). The gain in terms of extraction of general

characteristics linked to neurological prognosis, independently of each modality, made possible by this multimodal latent space, is at the same time offset by the lack of explainability concerning the final weight of each modality on the outcome. Moreover, one of the main difficulties of an early fusion lies in the problem of data commensurability. Indeed, the heterogeneous nature (language, image, time series), their dimensionality and the order of magnitude of the various modalities relevant to neuro-prognostication (e.g. language on the medical record, imaging, EEG) hinder the information synthesis in a single input vector. This lack of commensurability of data from each modality may render impossible the standard normalization procedures known to improve the quality of machine learning (19).

All in all, as late fusion is the easiest and the most explainable, and the early fusion has more potential to extract deep latent multimodal common representation linked to Disorders of Consciousness prognostication, the multimodal deep learning networks of the future will probably operate an intermediate multimodal fusion. In this partial multimodal fusion, each modality will have been initially processed in an intermediate way by dedicated networks, resulting in latent marginal representations specific to each modality. These marginal latent representations can then be more easily connected together on subsequent deep layers, thanks to the possibility of forcing their commensurability on previous layers. In this way, prognosis-specific multimodal latent features could be extracted simultaneously from the following layers using nested combinations.

Artificial intelligence (AI) has the potential to improve classification, prognostication and ultimately care for patients with Disorders of Consciousness. To address these challenges, a close multidisciplinary collaboration between professionals from various disciplines (neurologists, electrophysiologists, radiologists, engineers, mathematicians). Importantly, cross-training should be a key element of curricula: physicians need to be trained in AI while engineers need to learn about the specificities of this medical field.

Figure 1: Pros and cons of early and late fusion for multimodal evaluation of Disorders of Consciousness prognosis. VS = Vegetative State, EEG = Electro Encephalography, MRI = Magnetic Resonance Imaging, FA = Fractional Anisotropy, CRS-R = Coma Recovery Scale-Revised, NSE = Neuron Specific Enolase.



Bibliography:

1. Rohaut B, Eliseyev A, Claassen J. Uncovering Consciousness in Unresponsive ICU Patients: Technical, Medical and Ethical Considerations. *Crit Care*. 2019 Dec;23(1):78.
2. Dehaene S. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*. 2001 Apr;79(1-2):1-37.
3. Edlow BL, Claassen J, Schiff ND, Greer DM. Recovery from disorders of consciousness: mechanisms, prognosis and emerging therapies. *Nat Rev Neurol*. 2021 Mar;17(3):135-56.
4. Giacino JT, Ashwal S, Childs N, Cranford R, Jennett B, Katz DI, et al. The minimally conscious state: Definition and diagnostic criteria. *Neurology*. 2002 Feb 12;58(3):349-53.
5. Wannez S, Heine L, Thonnard M, Gosseries O, Laureys S, Coma Science Group collaborators. The repetition of behavioral assessments in diagnosis of disorders of consciousness: Repeated CRS-R Assessments for Diagnosis in DISORDERS OF CONSCIOUSNESS. *Ann Neurol*. 2017 Jun;81(6):883-9.
6. Hermann B, Salah AB, Perlberg V, Valente M, Pyatigorskaya N, Habert MO, et al. Habituation of auditory startle reflex is a new sign of minimally conscious state. *Brain*. 2020 Jul 1;143(7):2154-72.
7. Hermann B, Goudard G, Courcoux K, Valente M, Labat S, Despois L, et al. Wisdom of the caregivers: pooling individual subjective reports to diagnose states of consciousness in brain-injured patients, a monocentric prospective study. *BMJ Open*. 2019 Feb;9(2):e026211.
8. Gong KD, Lu R, Bergamaschi TS, Sanyal A, Guo J, Kim HB, et al. Predicting Intensive Care Delirium with Machine Learning: Model Development and External Validation. *Anesthesiology*. 2023 Mar 1;138(3):299-311.
9. Sandroni C, D'Arrigo S, Cacciola S, Hoedemaekers CWE, Kamps MJA, Oddo M, et al. Prediction of poor neurological outcome in comatose survivors of cardiac arrest: a systematic review. *Intensive Care Med*. 2020 Oct;46(10):1803-51.
10. Cronberg T, Rundgren M, Westhall E, Englund E, Siemund R, Rosen I, et al. Neuron-specific enolase correlates with other prognostic markers after cardiac arrest. 2011;
11. Mattsson N, Zetterberg H, Nielsen N, Blennow K, Dankiewicz J, Friberg H, et al. Serum tau and neurological outcome in cardiac arrest: Serum Tau in Cardiac Arrest. *Ann Neurol*. 2017 Nov;82(5):665-75.
12. Velly L, Perlberg V, Boulier T, Adam N, Delphine S, Luyt CE, et al. Use of brain diffusion tensor imaging for the prediction of long-term neurological outcomes in patients after cardiac arrest: a multicentre, international, prospective, observational, cohort study. *Lancet Neurol*. 2018 Apr;17(4):317-26.
13. Luyt CE, Galanaud D, Perlberg V, Vanhauzenhuysse A, Stevens RD, Gupta R, et al. Diffusion Tensor Imaging to Predict Long-term Outcome after Cardiac Arrest. *Anesthesiology*. 2012 Dec 1;117(6):1311-21.
14. Owen AM, Coleman MR, Boly M, Davis MH, Laureys S, Pickard JD. Detecting Awareness in the Vegetative State. *Science*. 2006 Sep 8;313(5792):1402-1402.

15. Hermann B, Stender J, Habert MO, Kas A, Denis-Valente M, Raimondo F, et al. Multimodal FDG-PET and EEG assessment improves diagnosis and prognostication of disorders of consciousness. *NeuroImage Clin.* 2021;30:102601.
16. Engemann DA, Raimondo F, King JR, Rohaut B, Louppe G, Faugeras F, et al. Robust EEG-based cross-site and cross-protocol classification of states of consciousness. *Brain.* 2018 Nov 1;141(11):3179–92.
17. Bekinschtein TA, Dehaene S, Rohaut B, Tadel F, Cohen L, Naccache L. Neural signature of the conscious processing of auditory regularities. *Proc Natl Acad Sci.* 2009 Feb 3;106(5):1672–7.
18. Claassen J, Doyle K, Matory A, Couch C, Burger KM, Velazquez A, et al. Detection of Brain Activation in Unresponsive Patients with Acute Brain Injury. *N Engl J Med.* 2019 Jun 27;380(26):2497–505.
19. Colliot O, editor. *Machine Learning for Brain Disorders* [Internet]. New York, NY: Springer US; 2023 [cited 2023 Jul 30]. (Neuromethods; vol. 197). Available from: <https://link.springer.com/10.1007/978-1-0716-3195-9>
20. Sitt JD, King JR, El Karoui I, Rohaut B, Faugeras F, Gramfort A, et al. Large scale screening of neural signatures of consciousness in patients in a vegetative or minimally conscious state. *Brain.* 2014 Aug;137(8):2258–70.
21. Gao J, Li P, Chen Z, Zhang J. A Survey on Deep Learning for Multimodal Data Fusion. *Neural Comput.* 2020 May 1;32(5):829–64.
22. Stahlschmidt SR, Ulfenborg B, Synnergren J. Multimodal deep learning for biomedical data fusion: a review. *Brief Bioinform.* 2022 Mar 10;23(2):bbab569.
23. Baltrušaitis T, Ahuja C, Morency LP. *Multimodal Machine Learning: A Survey and Taxonomy* [Internet]. arXiv; 2017 [cited 2023 Aug 12]. Available from: <http://arxiv.org/abs/1705.09406>
24. Bahdanau D, Cho K, Bengio Y. *Neural Machine Translation by Jointly Learning to Align and Translate* [Internet]. arXiv; 2016 [cited 2023 Aug 12]. Available from: <http://arxiv.org/abs/1409.0473>
25. Wei W, Poirion E, Bodini B, Durrleman S, Ayache N, Stankoff B, et al. Predicting PET-derived demyelination from multimodal MRI using sketcher-refiner adversarial training for multiple sclerosis. *Med Image Anal.* 2019 Dec;58:101546.
26. Luigi Antelmi, Nicholas Ayache, Philippe Robert, and Marco Lorenzi. Sparse Multi-Channel Variational Autoencoder for the Joint Analysis of Heterogeneous Data. *Proceedings of Machine Learning Research*, (97):302-311, 2019.
27. Bach FR, Lanckriet GRG, Jordan MI. Multiple kernel learning, conic duality, and the SMO algorithm. In: *Twenty-first international conference on Machine learning - ICML '04* [Internet]. Banff, Alberta, Canada: ACM Press; 2004 [cited 2023 Aug 12]. p. 6. Available from: <http://portal.acm.org/citation.cfm?doid=1015330.1015424>
28. Pei X, Zuo K, Li Y, Pang Z. A Review of the Application of Multi-modal Deep Learning in Medicine: Bibliometrics and Future Directions. *Int J Comput Intell Syst.* 2023 Mar 29;16(1):44.

29. Shetty S, S. AV, Mahale A. Multimodal medical tensor fusion network-based DL framework for abnormality prediction from the radiology CXRs and clinical text reports. *Multimed Tools Appl* [Internet]. 2023 Apr 21 [cited 2023 Aug 18]; Available from: <https://link.springer.com/10.1007/s11042-023-14940-x>
30. Hori C, Hori T, Lee TY, Sumi K, Hershey JR, Marks TK. Attention-Based Multimodal Fusion for Video Description [Internet]. arXiv; 2017 [cited 2023 Aug 18]. Available from: <http://arxiv.org/abs/1701.03126>
31. Lipkova J, Chen RJ, Chen B, Lu MY, Barbieri M, Shao D, et al. Artificial intelligence for multimodal data integration in oncology. *Cancer Cell*. 2022 Oct;40(10):1095–110.
32. Kim HE, Cosa-Linan A, Santhanam N, Jannesari M, Maros ME, Ganslandt T. Transfer learning for medical image classification: a literature review. *BMC Med Imaging*. 2022 Dec;22(1):69.